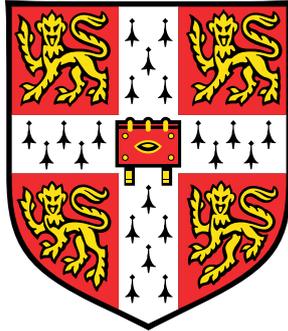


Soft Morphological Computation



Luca Scimeca

Department of Engineering
University of Cambridge

This thesis is submitted for the degree of
Doctor of Philosophy

Robinson College

October 2020

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other University. This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration, except where specifically indicated in the text. This dissertation contains less than 65,000 words including appendices, bibliography, footnotes, tables and equations and has less than 150 figures.

Luca Scimeca
October 2020

Abstract

Soft Robotics is a relatively new area of research, where progress in material science has powered the next generation of robots, exhibiting biological-like properties such as soft/elastic tissues, compliance, resilience and more besides. One of the issues when employing soft robotics technologies is the soft nature of the interactions arising between the robot and its environment. These interactions are complex, and their dynamics are non-linear and hard to capture with known models. In this thesis we argue that complex soft interactions can actually be beneficial to the robot, and give rise to rich stimuli which can be used for the resolution of robot tasks. We further argue that the usefulness of these interactions depends on statistical regularities, or structure, that appear in the stimuli. To this end, robots should appropriately employ their morphology and their actions, to influence the system-environment interactions such that structure can arise in the stimuli. In this thesis we show that learning processes can be used to perform such a task. Following this rationale, this thesis proposes and supports the theory of *Soft Morphological Computation (SoMComp)*, by which a soft robot should appropriately condition, or ‘affect’, the soft interactions to improve the quality of the physical stimuli arising from it. *SoMComp* is composed of four main principles, i.e.: Soft Proprioception, Soft Sensing, Soft Morphology and Soft Actuation. Each of these principles is explored in the context of haptic object recognition or object handling in soft robots. Finally, this thesis provides an overview of this research and its future directions.

Acknowledgements

I would like to thank all members of the BIRL laboratory, with whom discussions often lead to interesting research. I would like to especially thank Fumiya Iida, and Perla Maiolino for the guidance throughout the projects I have worked on during my PhD. I would like to thank my family, without whom none of this would have been possible. Finally, I would like to thank the Agriculture and Horticulture Development Board, which has funded the work I have carried out so far, and sustained the expenses of many events contributing to my development as a researcher.

Preface

All Chapters in this thesis, excluding the Introduction (Chapter 1) and Conclusion (Chapter 11), have been published or are forthcoming publication. Chapter 2 is a forthcoming book chapter in the book “Cognitive Robotics Handbook”, to be published by MIT press [224]. Chapter 4, 5, 7 are respectively based on [229], [98], [231] and have been published in international Journals. Chapter 8 is currently under review in the Soft Robotics journal. Chapter 3, 6 and 10 are respectively based on [230], [227] and [232] and have been published on international conferences. Other work, not incorporated within this thesis, includes 3 additional publications in international journals [86, 99, 101], and an additional publications in international conference proceedings [233]. All chapters based on reference articles report the details of the collaborations for the article, and my contributions within the research.

Table of contents

Nomenclature	xv
1 Introduction	1
1.1 Embodied intelligence and Soft Robotics	1
1.2 Soft Morphological Computation (<i>SoMComp</i>)	4
1.3 Contribution	6
1.3.1 A Use Case of Complex Soft Interactions	7
1.3.2 Model-Free Soft Proprioception in Continuum Soft Bodies	7
1.3.3 Soft Morphology to Structure Tactile Stimuli	8
1.3.4 Soft Actuation to Structure Tactile Stimuli	8
1.3.5 Complex Soft Actuation to Improve Tactile Perception	9
1.3.6 Soft Morphology & Action Co-Optimization	9
1.4 Structure of this Thesis	10
2 Background: Developmental Soft Robotics	13
2.1 Introduction	14
2.2 Bio-inspired Soft Robotics	14
2.2.1 Soft Materials and Soft Actuation	15
2.2.2 Soft Robot Control, Simulation and Learning	16
2.3 Developmental Robotics	19
2.3.1 Facets of Development	19
2.4 Developmental Soft Robotics	22
2.4.1 Design Principles	23
2.4.2 Ontogenetics and Adaptivity	27
2.4.3 Learning Through the Body	29
2.5 Challenges and Perspectives	31
2.5.1 Evolutionary Robotics	31
2.5.2 Complexity and Scalability	32

3	A Case Study of Soft Interactions: Robot Maturity Assessment	35
3.1	Introduction	36
3.1.1	Mango Ripeness, a Case Study	36
3.1.2	Approaches for Ripeness Assessment	37
3.2	Methods	38
3.2.1	Stiffness Model	38
3.2.2	Probing Gripper Mechanism	39
3.2.3	Tactile Sensor Technology and Data Acquisition	39
3.2.4	Robot and Experimental Set-Up	41
3.3	Results	43
3.3.1	Sensor Force Calibration	43
3.3.2	Mango Penetrometer Testing	43
3.3.3	Stiffness Measurement Analysis	45
3.4	Conclusions	48
4	Proprioception in Soft Continuum Bodies	51
4.1	Introduction	52
4.2	Methods	54
4.2.1	Tactile Sensing Technology	54
4.2.2	Soft Continuum Body Design and Sensor Embedding	55
4.2.3	Sensor Pressure Distributions and Visual Tracking	56
4.3	Results	57
4.3.1	Soft-Body Configuration Learning from Visual Tracking	57
4.3.2	Tactile Proprioceptive analysis	58
4.3.3	Tactile Proprioceptive Work-space Exploration	61
4.4	Discussion and Conclusions	61
5	The Importance of Morphology: Achieving Robotically Peeled Lettuce	65
5.1	Introduction	66
5.2	Methods and Implementation	67
5.2.1	Vision and Detection	67
5.2.2	Lettuce Orientation	71
5.2.3	Leaf Removal	71
5.3	Results and Discussion	72
5.3.1	Vision and Detection	72
5.3.2	Lettuce Orientation	72
5.3.3	Nozzle Design	74

5.3.4	Leaf Removal Radius	75
5.3.5	Leaf Trajectory	75
5.4	Demonstration	76
5.5	Conclusion	76
6	Soft Morphology to Structure Information	80
6.1	Introduction	81
6.2	Autonomous Category Formation	82
6.3	Tactile Sensor Technology and Data Acquisition	84
6.4	Methods and Experimental Set-Up	84
6.5	Experimental Results	89
6.5.1	Task Optimization	89
6.5.2	Autonomous Category Formation variations	90
6.5.3	Spatial Resolution Influence	91
6.6	Discussion	91
6.7	Conclusion	92
7	Soft Actuation to Structure Information	94
7.1	Introduction	95
7.1.1	Robotic Medical Palpation	95
7.1.2	Chapter Structure	97
7.2	Methods	97
7.2.1	Soft Phantom and Robot Set-Up	97
7.2.2	Tactile Sensor Technology and Data Acquisition	99
7.2.3	Probing Strategies	99
7.3	Analytical Framework	101
7.3.1	Task and Physical Interactions	101
7.3.2	Categorization	102
7.3.3	Motion Strategy Scoring	103
7.3.4	Experimental Procedure	104
7.4	Results	105
7.4.1	Sound Dimensionality Reduction	105
7.4.2	Information Structure and Silhouette Coefficient	108
7.4.3	Motion influence on Cognitive Maps	110
7.4.4	Categorization and Similarity Abstractions	112
7.5	Palpation Test Case	115
7.6	Conclusion	117

8	Action Complexity: Soft-Body Palpation	122
8.1	Introduction	124
8.2	Materials and Methods	126
8.2.1	Phantom Development	126
8.2.2	Sensor Technology	127
8.2.3	Robot Control Experimental Set-up	128
8.2.4	Bayesian Framework	129
8.3	Results	134
8.3.1	Exploring Action Complexity in Robot Medical Palpation	135
8.3.2	Bayesian Approaches for Confident Abnormality Detection	140
8.3.3	Online, Rapid Bayesian Exploration: Pruning the Search Space	141
8.4	Discussion	144
9	Action Complexity: Expressive Piano Playing	148
9.1	Introduction	150
9.2	Methods	152
9.2.1	Learning Framework	152
9.2.2	Experimental Set-Up	155
9.2.3	Robot Control	156
9.3	Results	157
9.3.1	Robot key-press control to sound feedback	157
9.3.2	Gaussian Process based Framework analysis	160
9.3.3	Human vs Robot Piano Playing	166
9.4	Discussion and Conclusion	168
10	Morphology-Action Co-Optimization: Complex Tactile Object Feature Recognition	171
10.1	Introduction	173
10.1.1	Morphology-Action Co-Optimization	173
10.1.2	Bayesian Exploration for Morphology-Action Co-Optimization	174
10.2	Methods and Experimental Set-up	175
10.2.1	Morphological Bayesian Exploration framework	175
10.2.2	Experimental Set-Up	179
10.2.3	Sensor Technology	181
10.3	Results	181
10.3.1	Morphology and Action for Object Classification	181
10.3.2	Morphological Bayesian Exploration	182

10.4 Discussion and Conclusion	184
11 Conclusion and Future Work	185
11.1 Conclusion	185
11.2 Future Work	187
11.2.1 Ubiquitous Proprioception	188
11.2.2 A Comprehensive Theory of Robotic Learning	189
11.2.3 Limitations and Future Remarks	190
11.2.4 Additional Work Beyond This Thesis	190
References	192

Chapter 1

Introduction

1.1 Embodied intelligence and Soft Robotics

Since the term Artificial Intelligence was coined in 1956, the research field of intelligence was initially dominated by the “computational paradigm of intelligence” (traditional cognitivism). In this context, intelligence was regarded as a computational processes, where symbolic operations were of central interest without explicitly considering what the symbols actually meant. At the time, a strong connection was conjectured between the idea of “intelligence”, the power of symbolic representation (for example in the brain), and the possibility for a system to change from a state to another [67, 175]. In this context, an individual would create a symbolic representation of the world by means of sensory perception, then a process akin to rule-based symbol manipulation would allow them to exhibit intelligence [150, 190]. While the computational paradigm of intelligence has given significant impact mainly in cyber-space, there have been a number of aspects of intelligence which cannot be fully explained in this framework.

Traditional cognitivism discards the importance of sensory perception for learning in the real world. Sensory perception is the ability of individuals to get, interpret, select, and organize physical stimuli. Learning, on the other hand, is the ability of individuals to use sensory perception to evolve and improve over time. Because symbols must represent entities or concepts in the world, and because the entity representations must in some way be task independent, there is an issue with how exactly sensory perception can lead to a symbolic representation of the sensed entities in the world, whether these are physical entities, emotions or even concepts. Without a representation of entities in the real world learning is impossible. This problem is also known as the symbol grounding problem. The physical grounding hypothesis arises to match the need for an agent to have its representation grounded in the physical world. In this context, the world becomes its own best model, and appropriate

interactions modulates the behavioural intelligence, as eloquently explained in an early article titled "Elephants Don't Play Chess" [25].

In contrast to the traditional cognitivist view of intelligence, the "embodied cognition" paradigm sees cognition as deeply dependent on the characteristics of the physical body of an agent [272]. In this view, agents are not passively exposed to sensory perception, but instead actively interact with their surrounding environment [150]. Therefore agents should strive for their action and motor control to contribute to the improvement of perceptual abilities. The embodied view of intelligence has shown that physical system-environment interactions cannot be overlooked to understand intelligent adaptive systems, especially in the physical real-world environment. For example, in the book "Vehicles: Experiments in Synthetic Psychology" [24], Valentino Braitenberg describes a series of thought experiments in which vehicles can exhibit complex and meaningful behaviours through increasingly more complex system-environment interactions. In his thought experiments, the vehicles' steering is connected to sensor outputs in various ways. By an increasingly more complex network of inhibitory or excitatory connections, and some analog sensors and actuators, the vehicles can steer towards a light, they can avoid each other, or group in various ways and exhibit basic behaviours of fear, aggression, liking and love, at least in the eye of the observer. The message in his book was key in showing how intelligence can be explained in a bottom-up fashion, and how intelligent behaviour can emerge from "simple" system-environment interactions.

Embodied intelligence has changed the field of Robotics, aiding in the search for those elements meant to achieve intelligence and adaptability in machines. Under the view of embodied intelligence, the importance of the characteristics of the physical body of an agent and its behaviour has been recognized for robotics systems. This change has fuelled a new found effort to endow robots with new capabilities. It is in this context that a field known as "Soft Robotics" has pushed forward robotics technologies by employing soft materials to mimic and learn from soft biological organisms [120]. Progress in Soft Robotics has been spurred by advancements in material science, which have benefited several areas including the development of functional materials, rapid prototyping technologies, like 3D printing, sensor development and much more. Soft Robotics has revoked the rigidity constraints of the industrial robots established in the earlier century, and has allowed the study and development machines which can bend, stretch, adapt, morph, grow, and achieve bio-inspired capabilities beyond the rigid machines capabilities. In manipulation, for example, the "universal gripper", a soft gripper capable of particle jamming through vacuum pressure control, has been shown to be able to grasp a diverse range of objects, previously impossible with any rigid solution [26]. Another branch of research as allowed to study and mimic animals such worms,

caterpillars, octopuses and more besides, furthering our knowledge of biology, as well as our ability to endow robots with functionalities akin to those of animals. These functionalities have been shown to be applicable in a number of scenarios [131, 140, 236]. A comprehensive overview of these fields will be provided in Chapter 2.

Soft robots often have large (or infinite) degrees of freedom, which makes them hard to control. Moreover, the softness of their body makes their motions and body dynamics complex and hard to predict, while the environment itself may not be rigid. Sensory stimuli (especially those from cameras, force, tactile, or other positional sensors) are also affected by these complex motions and dynamics, which makes it difficult (if not impossible) for robots to perceive their own state, as well as the state of their environment. Both perception and learning are influenced by the quality of the physical stimuli arising from the system-environment interactions, and as such, they are subject to the difficulties observed in soft robots/environments.

The goal of this thesis is to explore general principles for soft robots to exploit system-environment interactions for better perception and learning in the real world.

We hypothesize that the complexity in the system-environment interactions can actually be useful for perception. This hypothesis is based on the assumption that the complexity in soft robots/environments enhances the richness of the physical stimuli arising from their interaction. We argue that these enriched stimuli can be directly useful to the robot to improve perception and learning, and that to be useful these stimuli must be “structured” or otherwise present some statistical regularities.

The structures of sensory stimuli in soft robotics are not direct outcome of neuronal circuits or programmes in isolation, but that of complex interplays between morphology, soft action, soft sensing, and proprioception, on top of which learning occurs. These interplays arising from the embodiment of soft robotics are complex in nature and depend on the tasks and environment around the robot. These interplays are explored in this thesis as the framework of Soft Morphological Computation (*SoMComp*).

1.2 Soft Morphological Computation (*SoMComp*)

SoMComp is the framework by which the soft interactions arising between a robot in its environment can be purposefully used to improve the robot's perception and its ability to learn. This improvement is due to the *conditioning* of the stimuli arising from these interactions, through the robot's appropriate use of its Soft Morphology and Soft Actuation. In this context both Soft Proprioception and Soft Sensing are necessary for purposeful action and morphological adaptation.

This section introduces the four principles which constitute *SoMComp* (Fig. 1.1), i.e.: Soft Morphology, Soft Actuation, Soft Sensing and Soft Proprioception. It is important to notice that although the focus of this thesis is on soft robotic interactions, many of the concepts proposed are also relevant for more traditional robotic paradigms.

Soft Morphology: The physical stimuli of an agent acting in its environment are heavily dependent on the properties of the sensory system which captures them. For example, the elasticity of our dermis, the morphology of our fingers and our hands capabilities influence the perception of tactile stimuli. Similarly, the soft morphology of a robot can be used to influence the system-environment interactions.

In Fig. 1.1 Soft Morphology is placed as the main medium between the environment and learning processes. The morphology can, in fact, filter and select relevant physical stimuli from the environment by exploiting soft interactions. This process can be such that appropriate structure emerges in the information retrieved, and thus the information can be useful for learning. In other words, within *SoMComp* the robot should explore ways in which to change its morphology, so to improve the quality of the sensed stimuli, such that perception and learning can be improved. In later chapters we will explain how these changes can be achieved by the robot in an evidence-based manner.

Soft Actuation: Similar to the morphology of the body, our actions shape a robot's perception. In Developmental Robotics, the process by which the coordinated robot actions can influence the physical stimuli is known as Sensory-Motor Coordination (SMC)[189]. This will be further explained in Chapter 2. In *SoMComp* we take advantage of SMC. However, SMC differs slightly. The law of requisite variety states that "if a system is to be stable, the number of states of its control mechanism must be greater than or equal to the number of states in the system being controlled" [6]. The complexity of soft robots is usually such that an infinite number of states are possible. We argue that because of this, a similar complexity

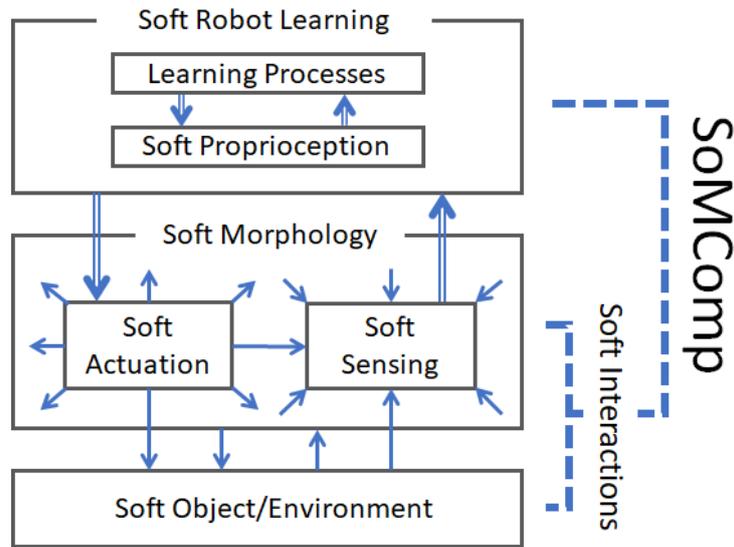


Fig. 1.1 The Figure shows a diagram of Soft Morphological Computation (*SoMComp*).

in the action may also be necessary. We call this Soft Sensory-Motor Conditioning (SSMC). In SSMC these complex actions can improve the richness of the sensed stimuli by inducing complex soft dynamics in the interactions. If appropriate actions are used, these rich stimuli will present statistical regularities, or structure, and can be directly useful for perception and learning.

As shown in Fig. 1.1, the actuation influences the interactions by acting on the body of the robot, and the physical stimuli are affected according to the changes induced by the actuation. A robot should thus explore ways in which to actively use its own body to improve the quality of the stimuli arising from interactions, by using appropriate actuation. Similarly to soft morphology, in later chapters we will show how this can be achieved in an evidenced based manner.

Soft Sensing: Soft Sensing refers to the ability of a robot to retrieve stimuli arising from system-environment interactions. In Soft Sensing it is important that the stimuli perceived depend, or can be influenced by, the robot. For example, a camera mounted on the robot, or one which is facing the robot's workspace, can perform soft sensing. This is because the sensor response can be influenced by the interactions of the robot and its environment. Even a sound-sensor can perform Soft Sensing, as shown in later chapters. On the other hand, a camera not mounted on the robot, and whose field of vision does not include the robot's body or its workspace cannot perform soft sensing, even if the visual feeds are used by the robot for decision making. This can, for example, be the case for an autonomous vehicle

using satellite images for route prediction. That is because the robot interactions with the environment do not generally affect the physical stimuli captured by the camera, and thus *SoMComp* is impossible. Both the Soft Actuation and Soft Morphology principles depend on soft sensing. Without a medium to retrieve physical stimuli from the environment, in fact, it is impossible to reason about the quality of said stimuli, and use them to explore ways to improve the robot's morphology or actions.

Soft Proprioception: In this thesis, we will refer to proprioception as both the ability of the robot to sense its own state, as well as the knowledge of its own physical properties. For example, without inverse kinematics knowledge there cannot be any conditioning of the stimuli via actuation. In Fig. 1.1, soft proprioception is tightly related to learning processes. This relationship is due to the fact that learning processes need knowledge of the body and the robot's capabilities to explore these appropriately for perception, while learning processes should be used to improve the robot's understanding of its own body and capabilities. Note that proprioception is not a prerequisite for learning, as learning processes can function without proprioception. However, it is a prerequisite for *SoMComp*, as no appropriate action or morphological changes can be achieved without knowledge of the body and robot's capabilities. Moreover, proprioception can be learned, as shown in later chapters.

Several of the concepts developed through this thesis are inspired by older concepts within the visionary field of *Developmental Robotics*. These concepts are presented in the review chapter of this thesis, and within the theory of *Soft Morphological Computation* they are revisited and changed, in the light of new developments in material science and the advent of the area of Soft Robotics, sensing and artificial intelligence. The employment of *SoMComp* can change the way we perceive robotics systems, where adaptation arises via learning *through the body*. Depending on the task at hand and the robot capabilities (e.g. object grasping, object detection etc.) it may be useful for the robot to explore changes within its body, as well as changes in the way it interacts with the objects, to influence the physical stimuli in ways that make learning simpler, faster or more robust. If the physical stimuli are appropriately influenced, it may even be possible to solve a task as complex as object recognition by simple unsupervised methods, like shown in later chapters.

1.3 Contribution

In this section we will give an overview of the contributions in this theses in the context of each chapter and their content.

1.3.1 A Use Case of Complex Soft Interactions

As previously mentioned, soft system-environment interactions are characterized by the complex dynamics between the robot (or a biological organism) and its environment. This complexity is due to several factors, including the many or infinite degrees of freedom of soft bodies, the hard-to-capture collision dynamics of soft objects and more besides. As previously explained, this is the principal motivation behind the need of employing *SoMComp*. The majority of the interactions arising in real world settings are, in fact, soft in nature. This is the case for the human hand manipulation of objects, for example, due to the softness of our dermis and often the softness of the objects we manipulate.

In Chapter 3 we demonstrate the complexity of soft interactions in a real world scenario. We consider the area of agri-food robotics, and more specifically the problem of estimating fruit ripeness by touch, via a robotics platform. In the chapter we propose a model-based approach to use non-destructive light touch experiments to estimate ripeness with a capacitive tactile sensing array. As shown in the chapter, the complex dynamics arise from the complexity of the structural properties of the fruit considered, including the soft properties of both the skin and pulp of the fruit, which influence the tactile sensor stimuli. These complex dynamics make it hard to achieve accurate ripeness estimations with simple model-based approaches. This chapter further highlights the pro and cons of purely based model-based approaches, and shows their limitation in capturing the complexity of soft interactions.

1.3.2 Model-Free Soft Proprioception in Continuum Soft Bodies

After the complexity showcase, we dive into each of the principles in *SoMComp*. We start from the principle of Soft Proprioception. As previously explained, Soft Proprioception is a core part of *SoMComp*, as only the knowledge of the robot's bodily properties and skills can allow it to appropriately manipulate these to its advantage.

We focus on continuum soft robots which, due to their softness, are hard to control, and their dynamics hard to predict. In Chapter 4 we propose a sensorization method to achieve proprioception in continuum soft robots with capacitive tactile sensing arrays. The method proposed is model-free, as opposed to the previous Chapter, where a model-based approach was used to capture the complex interactions. In this work, no knowledge of the body is needed a priori, and only the careful placement of sensors is necessary. The proprioception is achieved through movement akin to 'twitching' in infants, where random jerky motions reinforce the robot understanding of its own body given sensor evidence. Moreover, the robot interacts with entities around it, and these interactions are faithfully captured by model-free

approaches. The chapter provides a rationale, together with advantages and limitations of these approaches, with respect to the model-based approach shown in Chapter 3.

1.3.3 Soft Morphology to Structure Tactile Stimuli

The Soft Morphology principle in *SoMComp* revolves around the concept that the morphology of a soft robot can influence the physical stimuli, and that this influence can be useful for computation. This concept is inspired by the Developmental Robotics concept of “Morphological Computation”, but it is revisited in the context of soft interactions.

In Chapter 5 we first demonstrate the importance of morphology for soft interactions in a real-world practical scenario. The task set for the robot is that of removing the outer leaves of a lettuce, a soft and delicate object which requires careful manipulation. We show how the resolution of the task depends on the properties of the robot manipulator, which in turn influences the soft interactions between the robot and the soft produce.

In Chapter 6 we take this one step further, and introduce Soft Sensing. The robot task is object feature recognition, and the robot is equipped with a tactile sensor, as well as changing soft morphology. In the chapter we use unsupervised methods to show how the soft morphology can be used to influence the soft interactions, such that structure arises in the physical stimuli, and learning is possible. These morphological changes allow the robot to discriminate amongst different objects based on their physical properties with simple unsupervised methods. This in contrast to common approaches in the field, where more and more complex machine learning techniques would normally be used to solve complex problems. Complexity in the computation is here off-loaded to the body itself, while brain would instead be free to resolve higher cognitive tasks.

1.3.4 Soft Actuation to Structure Tactile Stimuli

Besides Soft Morphology, the principle of Soft Actuation is the other mean by which the robot can influence the physical stimuli. The principle of Soft Actuation is one by which the physical actions are used to appropriately affect the soft interactions such that structure can emerge in the physical stimuli.

In Chapter 7 we treat the principle of Soft Actuation. We focus on the task of robotic palpation, where a robot has to perform tactile discrimination of areas in a silicon phantom containing hard spherical inclusions, and areas without. Once a morphology has been established, we must reason about the type of actions necessary to influence the soft interactions appropriately, such that structure emerges in the stimuli. We propose an unsupervised mathe-

mathematical framework to achieve this. We show that by using this framework, it is possible to identify actions which can aid learning processes to perform tactile discrimination.

1.3.5 Complex Soft Actuation to Improve Tactile Perception

As previously mentioned in the Soft Actuation principle of *SoMComp*, we argue that complexity in the actions of the robot may be required to achieve appropriate structuring of the physical stimuli. In Chapter 8 and Chapter 9 we focus on the aspect of complexity in Soft Actuation. We employ large physical experimentation with a robot to: first, assess whether complexity in the robot actions can be beneficial; and second, to find those actions which may influence the soft interactions appropriately so that structure emerges in the stimuli.

In Chapter 8 we go back to the topic of robotic palpation, and we devise a control strategy to generate more complex multi-axis robotic palpations. We show that these trajectories induce complex dynamics in the soft interactions between the phantom and the robot. We also show that these dynamics improve the robot's ability to find inclusions, and that complex palpations can therefore be beneficial to achieve discrimination.

In Chapter 9 we move away from classification to show instead a regression case. Here, a robot is made to perform complex key-strokes of piano keys with an elastic finger, and optimize its performance over trials via sound feedback. A Gaussian Process based approach is developed, where the robot is made to approximate the complex, non-linear, relationship of keystroke patters to sound outputs. The complexity of the action space allows the robot to appropriately interact with the piano instrument, so to be able to play comparably to a human player over 10 different playing styles.

1.3.6 Soft Morphology & Action Co-Optimization

Finally, in Chapter 10 we put together the concepts of Soft Morphology and Soft Actuation. If it is true that a soft robot should be able to change its own bodily properties according to the task to solve, then its actions must also change with the changing body, else the improvement achieved through the changes in one may be hindered by the inflexibility of the other. Moreover, these actions can be complex, to account for the complexity in the system-environment interactions, and enable the retrieval of appropriate physical stimuli. It is here that we propose a probabilistic approach to efficiently reason about changes in sensor data due to both changes in the morphology and the action. In this chapter we focus on tactile perception for the discrimination of objects based on salient features. The chosen features are shown to be undetectable unless the robot chooses an appropriate morphology and action, and efficiently explores both.

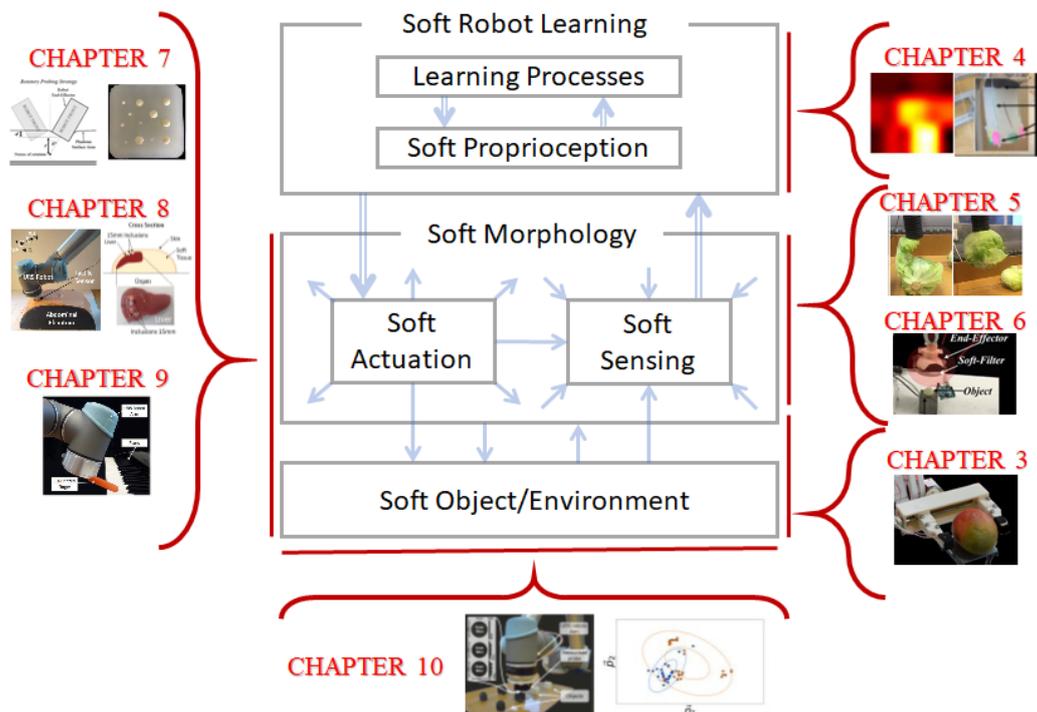


Fig. 1.2 The Figure shows a the chapter contributions within Soft Morphological Computation (*SoMComp*).

1.4 Structure of this Thesis

As summarized in Fig. 1.2, this thesis is structured as follows. Chapter 2 provides a state of the art review on fundamental topics to be discussed in this thesis. It introduces the areas of Soft Robotics and Developmental Robotics, and finally discusses the emergence of a new field, Developmental Soft Robotics, which re-establishes important concepts from the developmental sciences to be seen under the new context of Soft Robotics, and advancements in areas such as machine learning, sensing and actuation and material science. Throughout the remaining chapters in this thesis we will discuss each of the topics exposed in Section 1.2, and present a series of frameworks through which a robot would be able to change its bodily morphology, as well as its own action control, to improve the sensory perception derived from interaction with the environment. In Chapter 3 we present a case study of complexity in a real world robot manipulation problem, an important motivation behind the need for *SoMComp*. Chapter 4 treats the important principle of proprioception in Soft Robotics, and how to achieve a proprioceptive understanding of a continuum soft body with model-free methods. In Chapter 5 we move on to the Soft Morphology principle, and show the importance of morphology in a complex soft manipulation task. In Chapters 6, we show how we can appropriately make use of the soft morphology of a robot's tactile apparatus to

induce changes in its sensory perception. These changes can be purposeful and useful, and we show how they can aid a robotic system to perform classification of complex objects by touch. In Chapters 7, 8 and 9 we show how appropriate robot actions and complex interactions can similarly influence physical stimuli, and so help in complex object haptic perception and object handling scenarios. In Chapter 10 we put together the concept of soft Morphology and Soft Actuation to propose a unique, evidence-based, approach to jointly improve sensory perception via appropriate changes in both action and morphology. At last, in Chapter 11 we draw the conclusions of the work, summarizing the findings, and finally discussing the remarks of this research and its future in the robotics field.

Related Research Publications

Book Chapters

- 1 Hughes, J., Birell, S., **Scimeca, L.** & Iida, F. (forthcoming). Field robotics for harvesting. Elsevier.
- 2 **Scimeca, L.** & Iida, F. (forthcoming). Soft robotics: A developmental approach. MIT Press. **(CHAPTER 2)**.

Journal Articles

- 1 He, L., Herzig, N., de Lusignan, S., **Scimeca, L.**, Maiolino, P., Iida, F. & Nanayakkara, T. (forthcoming). An abdominal phantom with tunable stiffness nodules and force sensing capability for palpation training. *IEEE Transactions on Robotics*.
- 2 **Scimeca, L.**, Hughes, J., Maiolino, P., He, L., Nanayakkara, T. & Iida, F. (forthcoming). A bayesian framework for multi-axis soft-body palpation. *Soft Robotics*. **(CHAPTER 8)**.
- 3 **Scimeca, L.**, Ng, C. & Iida, F. (2020). Gaussian process inference modelling of dynamic robot control for expressive piano playing. *PLOS ONE*, 15(8), 1–17. **(CHAPTER 9)**.
- 4 **Scimeca, L.**, Maiolino, P., Bray, E. & Iida, F. (2020). Structuring of tactile sensory information for category formation in robotics palpation. *Autonomous Robots*, 1–17. **(CHAPTER 7)**.
- 5 Hughes, J., Gilday, K., **Scimeca, L.**, Garg, S. & Iida, F. (2019). Flexible, adaptive industrial assembly: Driving innovation through competition. *Intelligent Service Robotics*, 1–10.
- 6 **Scimeca, L.**, Hughes, J., Maiolino, P. & Iida, F. (2019). Model-free soft-structure reconstruction for proprioception using tactile arrays. *IEEE Robotics and Automation Letters*, 4(3), 2479–2484. **(CHAPTER 4)**.
- 7 Hughes, J., **Scimeca, L.**, Ifrim, I., Maiolino, P. & Iida, F. (2018). Achieving robotically peeled lettuce. *IEEE Robotics and Automation Letters*, 3(4), 4337–4342. **(CHAPTER 5)**.
- 8 Stone, T., Webb, B., Adden, A., Weddig, N. B., Honkanen, A., Templin, R., Wcislo, W., **Scimeca, L.**, Warrant, E. & Heinze, S. (2017). An anatomically constrained model for path integration in the bee brain. *Current Biology*, 27(20), 3069–3085.

Conference Proceedings

- 1 **Scimeca, L.**, Maiolino, P. & Iida, F. (forthcoming). Self-supervised learning through scene observation for selective item identification in conveyor belt applications, In *Annual conference towards autonomous robotic systems*. Springer.
- 2 **Scimeca, L.**, Iida, F., Maiolino, P. & Nanayakkara, T. (2020). Human-robot medical interaction, In *Companion of the 2020 acm/ieee international conference on human-robot interaction*.
- 3 **Scimeca, L.**, Maiolino, P. & Iida, F. (2020). Efficient bayesian exploration for soft morphology-action co-optimization, In *2020 3rd ieee international conference on soft robotics (robosoft)*. IEEE. **(CHAPTER 10)**.
- 4 **Scimeca, L.**, Maiolino, P., Cardin-Catalan, D., del Pobil, A. P., Morales, A. & Iida, F. (2019). Non-destructive robotic assessment of mango ripeness via multi-point soft haptics, In *2019 international conference on robotics and automation (icra)*. IEEE. **(CHAPTER 3)**.
- 5 **Scimeca, L.**, Maiolino, P. & Iida, F. (2018). Soft morphological processing of tactile stimuli for autonomous category formation, In *2018 ieee international conference on soft robotics (robosoft)*. IEEE. **(CHAPTER 6)**.

Chapter 2

Background: Developmental Soft Robotics

Reference Publication

This chapter was adapted from a book chapter titled “**Soft Robotics: a Developmental Approach**” [224], written in collaboration with Dr. Fumiya Iida, and forthcoming in the book “Cognitive Robotics Handbook” to be published by MIT Press. The chapter reviews the topics and Soft Robotics and Developmental Robotics, before introducing a new area titled “Soft Developmental Robotics”, where the two fields merge into a unique new area spurring further research. The contribution of the chapter included the writing of all sections, with guidance by Dr. Fumiya Iida on style, content and layout.

2.1 Introduction

In this Chapter we will first briefly introduce and review the area of Soft Robotics, with emphasis on how compliance and *softness* have changed the robotics landscape in the past two decades. We then introduce the area of Developmental Robotics, and the key ideas which are fundamental for understanding the relationship between biological systems and artificial systems. Finally, we discuss how the developmental sciences and Soft Robotics are irrevocably linked, into what we have chosen to call “Developmental Soft Robotics”. Here, in fact, the two fields can be merged into one where the developmental sciences can aid in the design and make of soft robots, which can then be used as platforms to better understand biological systems. We will also discuss how the biological concepts of phylogenetic development, ontogenetic development and short-term adaptation are indeed naturally suited to be embedded within a “soft” robotic context, and conclude with the concept of “Soft Morphological Computation” within this new area.

2.2 Bio-inspired Soft Robotics

Deformation is a fundamental characteristic of biological systems. Almost 90% of the human body is composed of soft tissue; Many vital functions such as heart, lung, muscles, eye lenses, and more besides, depends on deformation of materials. In bipedal walking, for example, evidence has shown how the soft tissue of the body might not only cushion impacts on every stride, but also both save muscles the effort of actively dissipating energy, and perform a considerable amount of the total positive work per stride by soft tissue elastic rebound [280].

In the past few decades there has been an unprecedented advancement in material sciences and manufacturing techniques, furthering our knowledge of functional materials and empowering artificial systems with newfound abilities. These advancements, together with the better understanding of biological systems, gave rise to the era of Soft Robotics, where bio-inspired robotics platforms make use of soft and deformable materials to achieve more flexible, adaptable and robust behaviours [97, 120].

Since the dawn of Soft Robotics, the application of material science and soft-body compliance has changed the robotics landscape. In manipulation, for example, the “universal gripper”, a soft gripper capable of particle jamming through vacuum pressure control, has been shown to be able to grasp a large number of objects [26]. Other solutions for grasping and manipulation range from tentacle-like systems [131] to pneumatic soft grippers [276] and human inspired soft-robotic hands [94] (Fig. 2.1).

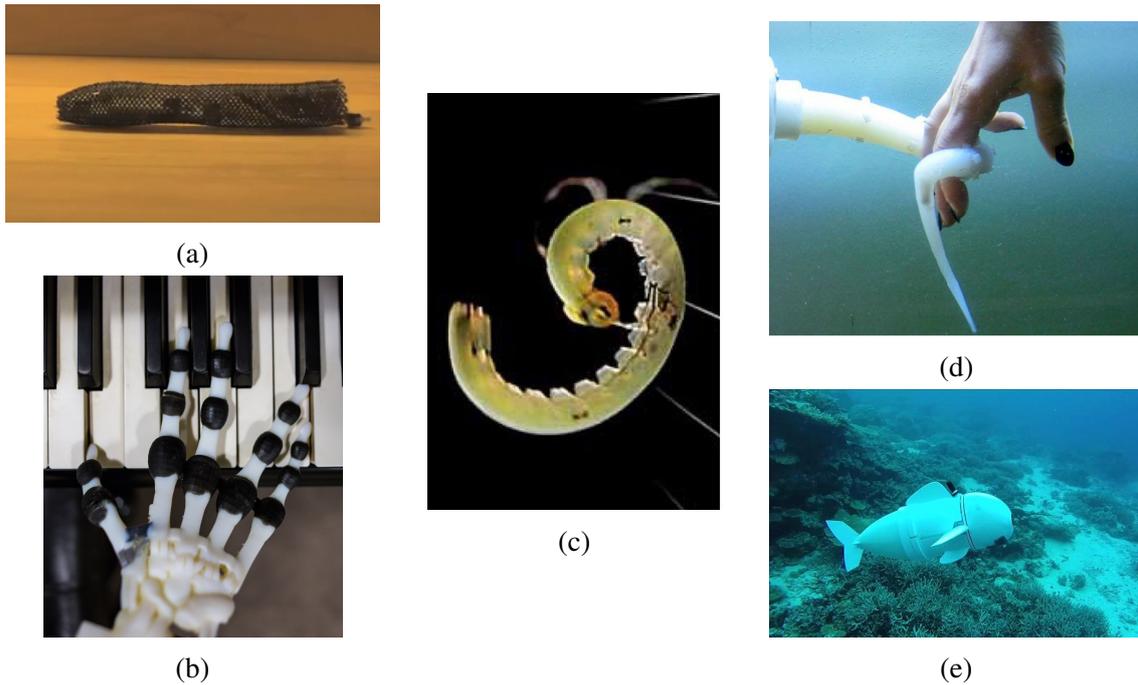


Fig. 2.1 Bio-inspired Soft Robot examples. (a) Worm-inspired soft robot [236]. (b) Human-inspired soft passive hand [94]. (c) Caterpillar-inspired soft robot [140] (d) Octopus-inspired tentacle [34]. (e) Fish-inspired soft robot [119].

Animal-inspired soft-robots are amongst the most developed sub-areas of Soft Robotics, where the robot platforms range from worms [236] or caterpillars [140], to octopuses [131], fish [119] and others besides (Fig. 2.1). In worm-like soft-robots for example, akin to their biological counterparts, the contraction of longitudinal muscles followed by the contraction of circumferential muscles simulates a travelling wave along the body, generating locomotion [263]. In caterpillars, motion is generated by coordinated control of the time and location of the prolegs attachment to the substrate, together with waves of muscular contraction [14].

The ability to mimic these unique systems makes soft robots an exciting new field, where the limits of the (rigid) robots of the past century can be overcome with new-found solutions.

2.2.1 Soft Materials and Soft Actuation

The area of Soft Robotics is inevitably connected to the field of material science, where new discoveries in the latter facilitate progress in the former. For a soft robot to be able to use material compliance to aid in robotics tasks, it is necessary for the make of the robot to be, at least in part, deformable. Elastomeric (polymer) materials, like EcoFlex or DragonSkin [241] have been at the center of researchers' attention for several years, with new materials composite materials being discovered every year. Moreover, the advent of 3D printing

technology has induced robot design and testing operations to be much faster than before, facilitating rapid and cheap prototyping in Soft Robotics.

Actuation poses one of the biggest challenges in Soft Robotics. In many animals, the co-action of a large number of muscles distributed over their body is capable of generating relatively high forces, facilitating coordinated and robust action. Replicating this ability is no easy feat, as the majority of the robotics solution lack the ability to generate forces comparable to the *industrial* robots of the past.

Four main soft actuation techniques exist: tendon driven, pressurized air or fluids, dielectric elastomeric actuators, or DEAs, and shape memory alloys, or SMAs [120]. Tendon driven actuation mimics biological musculoskeletal systems, where actuation is achieved through the pull and release of tendons, via the appropriate control of motors (Fig. 2.2a). Although a powerful and widespread actuation technique, a large number of tendons is usually necessary to achieve complex behaviors and control complexity increases with the number of motors necessary to control the tendons. For softer robots, like continuum soft robots, this type of actuation usually does not scale. Pressurized air and fluids are one of the most powerful actuation techniques for soft robots, capable of generating high forces and displacements. The actuation usually consists of varying the pressure inside pre-designed chambers within the body of the robot, to achieve their expansion and contraction and generate motion or morphological changes (Fig. 2.2b). However, these actuation systems are usually bulky, heavy and require high power sources, making it unsuitable for untethered robotics systems [130]. DEAs are made of soft materials that can be actuated through electrostatic forces (Fig. 2.2c). DEAs have been shown to have high strain/stress and mass-specific power, however, the need for DEAs to be pre-strained imposes rigid constraints on the robots' design [181]. Finally, SMAs with the most common nickel-titanium alloys, can generate force through a change in shape due to a rise or fall in temperature of the material (Fig. 2.2d). Temperature change control, however, is a challenge. High voltages are usually required to achieve temperature changes, and robustness over varying temperatures in the environment is still an issue to be overcome [208]. Other methods exist; it is possible for example to induce pneumatic contraction by evaporating ethanol via resistive heating [163], or achieve bending through combustion [259]. Other issues, like reduced output force or slow speed however come into play [204]. Soft Robotics actuation and material sciences are still an ever changing field, with new solutions being expedited by fast prototyping and iteration.

2.2.2 Soft Robot Control, Simulation and Learning

Soft robotic control poses several challenges and opportunities. Here, the “degree of softness” matters. Take, for example, a rigid robotic hand, where the palms and fingertips are covered

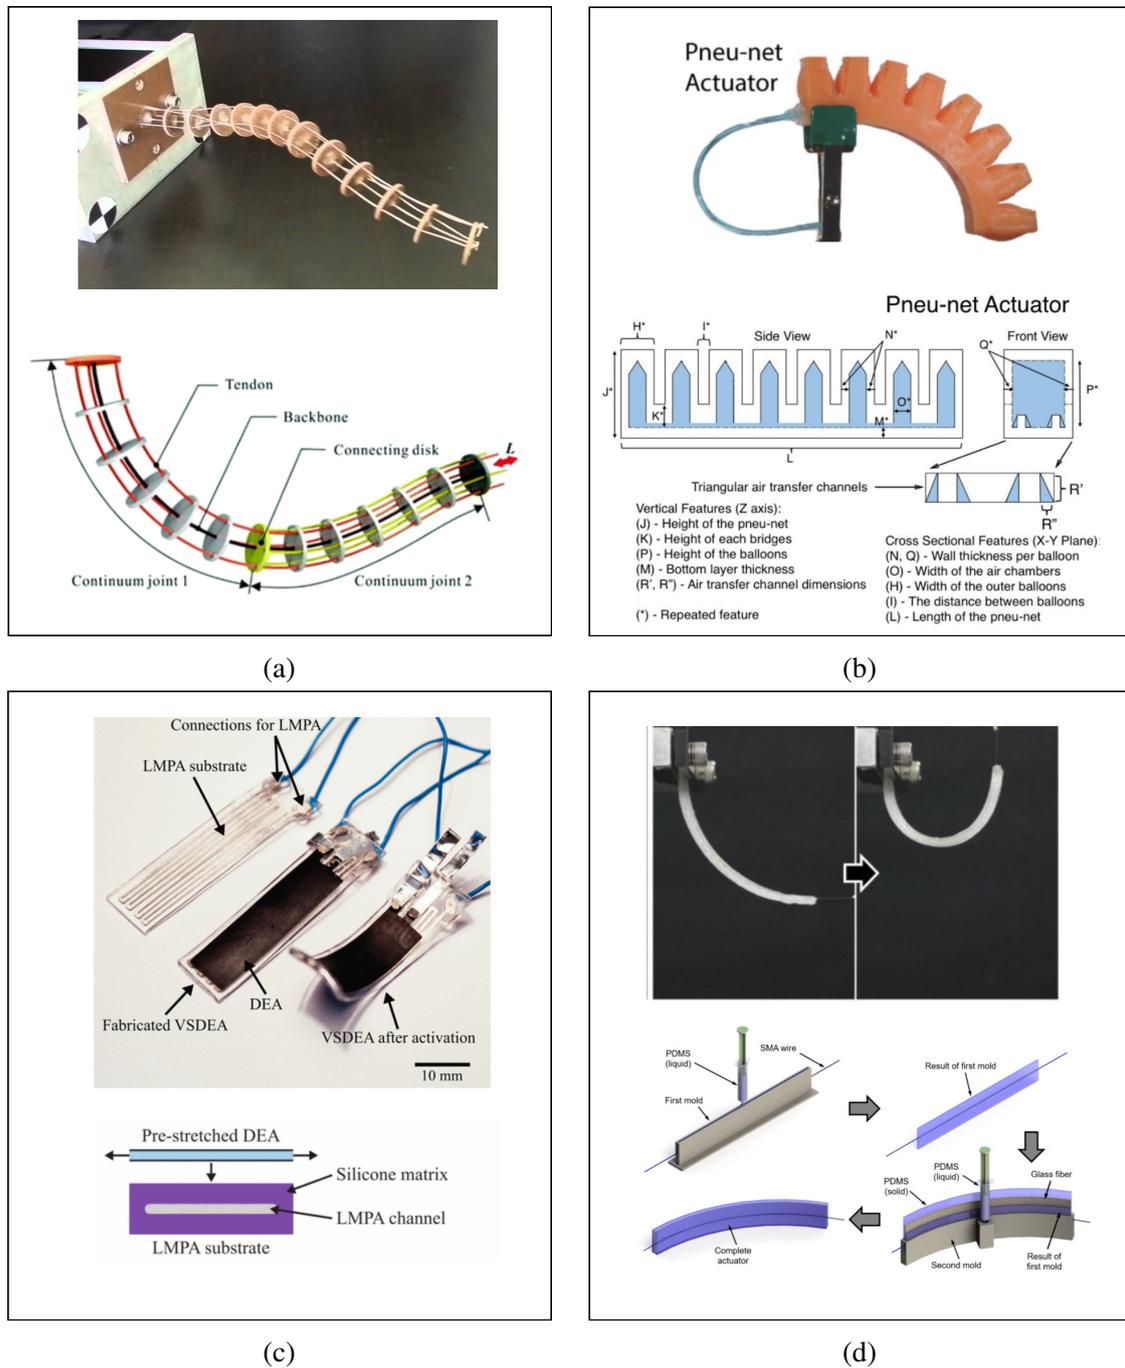


Fig. 2.2 Examples of some of the main actuation mechanisms used for soft-robotic systems. (a) Tendon driven continuum robot [213] and model [73]. (b) Pneumatic soft actuator and [279]. (c) Variable stiffness dielectric elastomer actuator [239] (d) Curved shape memory alloy-based soft actuator [209].

with an elastomeric material. The control of the hand may be achieved with classical methods (i.e. inverse kinematics), however, the complexity of the control may be reduced, as it may here be possible to make use of the mechanical passive dynamics of the soft-fingers to achieve a desired grasping behavior, averting the need for sub-millimeter precision in the robot control [105, 193]. On the other hand, as the “degree of softness” in the body increases, new challenges arise.

A robot made entirely of elastomeric materials, for example one simulating the tentacle of an octopus or the trunk of an elephant, can not be controlled classically; moreover, proprioception and simulation become problematic. As opposed to the hard links with sliding or rotational joints in classical robots, the continuity and softness of the body makes the control and simulation of continuous soft robots much harder. Novel actuation methods aid robotics researchers in their endeavours to achieve desired robot control (Section 2.2.1), and new sensing and control methods are discovered on a daily basis [214]. Achieving a proprioceptive understanding of the robot’s configuration is necessary to be able to control continuum soft robots accurately and repeatably, thus making the appropriate sensorization of soft-bodies fundamental.

Much effort has been put in the sensorization of soft robots. The most common soft sensors are perhaps strain sensors, which are soft deformable sensors capable of sensing body deformations through stretching. It is thus possible to embed such sensors into the (soft) body of a robot, without influencing its ability to deform. Other sensors have been used, based on resistive [90] or capacitive [153] technologies. Recently, work in [69] and [229] have shown how it is possible to achieve a high fidelity proprioceptive understanding of a continuum soft body through its sensorization via Fiber Optic and capacitive tactile sensors respectively.

In the context of control, and simulation, Learning plays a fundamental role. With the infinite degrees of freedom posed by continuum soft body, for example, precise control via classical methods is hard, and usually does not scale. Model based solutions based on the piece-wise constant curvature assumption have been shown to work for small tentacle-like robots [52]. However, the error in the controller always increases with the increase in the number of soft segments within the robots. The models, in fact, are usually too simplistic to accurately capture the complexity of continuum soft robots. Learning in this case, has been shown to be useful in compensating for the lack of knowledge or model complexity [227].

2.3 Developmental Robotics

Cognitive Developmental robotics (CDR) is an area of research where robotics and the developmental sciences merge into a unique field; one that seeks to better robotics with insights from developmental sciences, and further our understanding of developmental sciences through the use of robotics platforms [150]. The need for Cognitive Developmental Robotics to be a research area on its own, arose at the dawn of the 21st century from the need to understand, not only the cognitive and social development of individuals (as explored in the area of epigenetic robotics [282]), but also how the acquisition and development of motor skills, as well as morphology, influence the development of higher order cognitive functions [4, 5, 150]. In this context, robots can be used as experimental subjects, where developmental models can be implemented in robotics platform and scientists can gain insights from behavioural analysis, an approach known as synthetic methodology ([219, 246]).

As mentioned in Chapter 1, in stark contrast to the traditional cognitivist approach, in developmental robotics there is no clear separation between the physical body, the processes that determine reasoning and decision making (cognitive structure) and the symbol representation of entities in the world. Rather, these processes influence each other and intelligence emerges from their interaction.

2.3.1 Facets of Development

In biology, ontogeny can be defined as the development of an organism, usually from the moment it is conceived and thereby throughout its lifespan. Ontogenetic development thus can be seen as the evolution of an organism throughout its life, as dictated by a co-active action of internal (endogenous) and external (environmental) factors to the organism itself [112, 150]. Within ontogeny there is dissent on the role of physical development, through maturation and growth, and the role of learning. Although initially ontogenetic development was seen as tightly coupled to physical development, and learning would only occur as a consequence of development itself [186], several subsequent views thought to break the boundary between learning and development. One such view would see learning and development influencing each other bi-directionally, thus learning, modulated by physical developmental processes, could actively advance development itself [126]. Others thought dynamics processes to be at the base of development and learning, breaking the boundaries between development and learning altogether in what is known as “Dynamic Systems Approach to Development” [68]. Understanding ontogeny has the potential to give us invaluable insights to better understand biological system and thus build better artificial systems. A few components, or facets, of

ontogenetic development are key to that end, and their understanding will be important for the purpose of the concepts proposed throughout the chapters in this thesis. Therefore, we will briefly report some of these facets in the following subsections.

Degrees of Freedoms, Freezing and Freeing

The degrees of freedom problem was first introduced by Russian physiologist N. Bernstein [15], and it refers to the ability of many biological organisms (humans in particular) to achieve highly controlled and coordinated behaviors, despite the non-linearity of their muscular-skeletal system and their high number of muscular degrees of freedom. Part of a solution to the degrees of freedom problem was proposed by Bernstein himself, through what he called the principle of freezing and freeing degrees of freedom. By freezing, or tightly coupling the peripheral joints, a body can reduce its degrees of freedom so that learning is possible. Freeing, or weakening the coupling at the peripheral joints at a later stage would instead allow the body to re-claim its degrees of freedom, and learn more complex motion patterns [250]. In this process, the emergence of coordinated motion from a dynamic interaction with the environment is fundamental for any organism to achieve controlled behaviors.

Self-exploration and Spontaneous activity

At infancy, the inquiring of one's surroundings through physical exploratory action, and the perceptual consequences of said explorations, have been shown to play a crucial role in an infant's "sense of bodily self" [207, 264]. The concept of self-exploratory action is tightly coupled with the concept of spontaneous activity. Infants, for example, explore their physical constraints through coordinated behavior emerging from spontaneous neural and motor activity (e.g. kicking or suckling), useful to create or reinforce joint coupling and motor synergies.

Intrinsic Motivation

When infants explore the surrounding environment and their own physical constraints, not all their actions may be strictly goal oriented. Instead, an intrinsic value mechanism guides the motor actions to achieve specific motion patterns. Intrinsic motivation has a key role in learning and development, and, in infants, it has been linked to the concept of curiosity [184]. In artificial systems, such a mechanism could allow learning without explicit teaching or supervision [11, 12].

Categorization

Categorization is one of the most fundamental abilities for the majority of living organisms, without whom there would be no chance for an organism to distinguish any entity within its environment, including food sources, dangers, peers and so on [150]. The vast majority of the organisms, thus, are capable of categorising and discriminating between a wide range of sensory stimuli [56]. This is not usually a passive process, but rather one where the active exploration [150, 199] and the motor-coordinated behaviour of an organism would purposefully influence the sensory stimuli so that structure arises, and categorization is both possible and simpler [255].

Sensory Motor Coordination

In the context of embodied cognition, it becomes key to consider and understand the interactions arising when an individual learns and adapts its motor skills while interacting with entities in the world it lives in. The sensory-motor coordinated behaviours link action and perception in a cohesive loop to better discriminate known or unknown entities in the world, these are also known as “Sensory Motor Contingencies” (SMC) [189]. Under the light of SMCs, the perceptual experience of an individual is no more thought of as a consequence of brain computations, but rather a synergic interaction of skilful actions and sensing, both influencing each other to improve the perceptive experience [27, 183, 189].

Morphology and Morphological Computation

Akin to the role of sensory motor coordinated behaviour, the morphology of the body plays a fundamental role in the developmental process, and the acquisition of motor skills [16, 197]. Within the idea behind Morphological Computation, some complexity due to the interaction of an agent with its environment can now be outsourced to the body, leaving more canonical computational frameworks to serve higher level functions [192]. It is here that the role of morphology, materials and mechanical intelligence can be understood.

Body schema

Body schema, also known as forward models, are simulators of the musculo-skeletal system and the environment [37, 78, 273]. These simulators confer the ability to an organism to understand both their body and the environment surrounding them, and thereby predict the consequences of their actions. Systems possessing body schema are capable of predicting future states, given their body configuration and other sensory inputs. In humans, instances of such systems have been previously hypothesized to exist in the cerebellum [162].

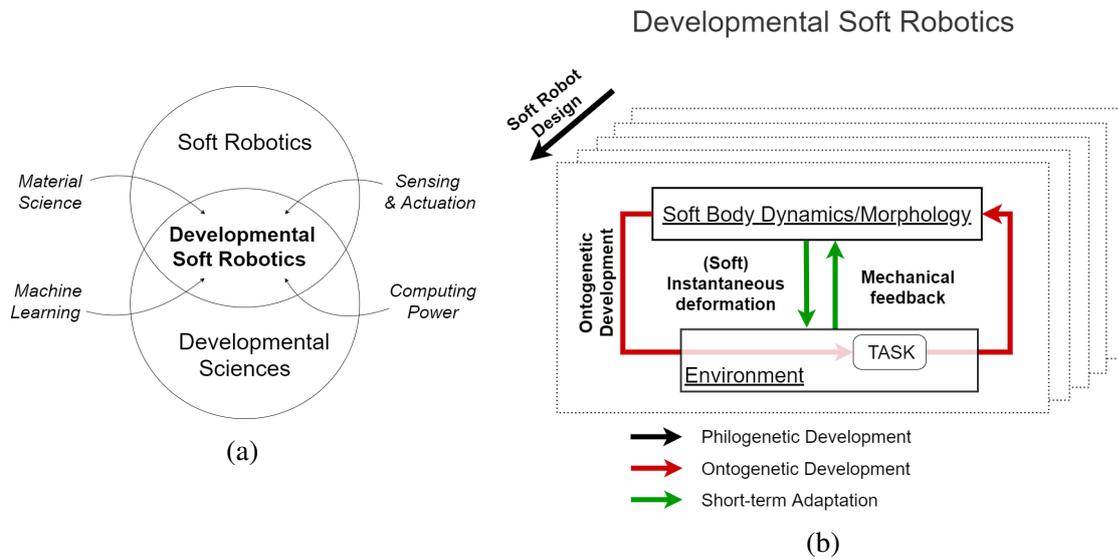


Fig. 2.3 Developmental Soft Robotics. (a) Conceptual Map of Developmental Soft Robotics. (b) Diagram of Developmental Soft Robotics across timescales.

2.4 Developmental Soft Robotics

One of the most difficult tasks in modern day robotics is to achieve an appropriate robot design for a robot to perform certain tasks in the world. As previously mentioned, the advent of Soft Robotics, if anything, has increased the complexity of robots, revoking the rigidity constraints established in the earlier century, and bringing about a new era. In this new era, robot design is driven by factors much like biological systems, where functional morphology, co-ordinate sensory-motor action, physical adaptation and embodiment all contribute to the “robot’s survival” in the world, and to its ability to see a task to completion. Developmental Soft Robotics aims at bringing together the areas of Soft Robotics with that of Developmental Robotics and the developmental sciences. These, in fact, are irrevocably linked, as we will later show (Fig. 2.3a).

Within the developmental sciences, in its simplest form, the development of a biological organism can be distinguished on three different scales: Phylogenetic, Ontogenetic and short-term.

In biological organisms, *Phylogenetic Development* has the largest time-scale, where changes happen at the level of groups of organisms, over many generations, and where processes like natural selection are responsible for certain “traits” to survive and evolve, while others to become extinct. Akin to Phylogenetic development is Soft Robotics design, where the design of robots is adaptive and ever-changing, to comply and conform to the task the robot has to achieve. Currently, the majority of the adaptation is due to human design and

biased by human skill and experience. However, new methodology for autonomous designed are a hot research topic, and processes like evolutionary algorithms have shown promise in the past [55, 179].

Ontogenetic development, like previously explained, concerns changes throughout and within the lifespan of an organisms, and include growths and bodily adaptation. The ability of robots to “morph” throughout their lifespan to achieve desired behaviour has been one of the key advantages of soft robots, as opposed to their rigid counterparts of the previous century. Robots navigating through growth, like fungal hyphae [84], elongating their bodies due to pressure, as well changing their stiffness to change their body dynamics and achieve different behaviors [36] are all examples of such adaptability.

Short-term adaptation refers to the shortest adaptive and developmental time-scale of all, where adaptation needs to be achieved instantaneously. Short-term adaptation is perhaps the most naturally suited to be discussed in a soft setting. In the past, this type of adaptation needed to be actively achieved at the control level, where real-time control would allow short-time adaptive behavior through mechanical or sensory feedback. Within the Soft Robotics framework, much like biological organisms, the short time adaptation is just a consequence of the soft instantaneous deformation of the soft body itself. When we delicately slide our finger through a ridged surface, for example, the need for complex and precise control is void by the ability of our dermis to deform and conform to the surface under touch. Much like the illustrated example, the compliance and softness of materials, in soft robots, can achieve short-term adaptation. The mechanical feedback becomes only a physical consequence of contact, and compliance can naturally suppress the need for complex controllers.

Fig. 2.3b illustrates the diagram of Developmental Soft Robotics within the timescales presented. For the remainder of the chapter, we will highlight some of the design principles to achieve short term and ontogenetic adaptation, as well as briefly explain evolutionary algorithms on a phylogenetic time-scale interest. Finally, we discuss some of the challenges and perspectives for the future.

2.4.1 Design Principles

In the next subsections we will discuss a collection of design principles which can aid in the development of adaptive soft machines, based on current literature.

Functional Morphology and Morphological Computation

As explained in Section 2.3.1, the morphology of the body plays a fundamental role in living organisms, one which influences their learning and developmental process, and aids in the

ability of said organisms to perform every-day tasks. When designing robotics systems, if shape was initially the most salient of morphological features, with the advent of Soft Robotics this may no longer be the case. Materials at different levels of elasticity have been shown to be able to perform “computation” [57, 227]. In Chapter 6, for example, we will show how complex haptic information can be used to classify objects based on different properties, simply by unsupervised clustering methods. This is achieved by exploiting the soft characteristics of elastomers, which can influence the response of tactile sensors so as to simplify object discrimination by a robot.

A paradigm trying to make use of the complex body-environment interactions is the “reservoir computing” framework of computation. The original idea behind reservoir computing begins with network computation, where an input is fed to a network, which computes a corresponding output. In reservoir computing a fixed random dynamical system, also known as reservoir, is used to map input signals to a higher dimensional space. The “readout” final part of the network, then, is trained to map the signals from the higher dimensional space to their desired output. As previously mentioned, soft robots, as well as biological organisms are usually made, at least in part, of soft materials. The body dynamics of soft robots are thus very complex, highly non-linear and high dimensional, making control hard. Through the reservoir computing paradigm it is possible to capitalize on the complexity of such system by exploiting the soft body as a computational resource, using the body dynamics to emulate non-linear dynamical systems, and thus offloading some of the control to the body itself [169, 172]. Work in [170], for example, has shown how it is possible to control a complex continuum soft arm, inspired by the tentacle of an octopus, in close loop without any external controller, by using the body of the robot as a computational resource. Under this light, high non-linearity and complexity may be a desirable property of the body, and design might have to be thought of accordingly.

An additional property which allows soft bodies to be used as a computational resource is memory. The soft body dynamics of soft robots, in fact, can exhibit short-term memory, allowing robots to emulate functions which require embedded memory [170]. When under-actuating a continuum soft robot, for example, it may be the case that control mechanism is not deterministic with respect to the behaviour of the robot. In these cases the behavior of the robot may depend, not only on the induced control and its current state, but also on the history of the previous robot states. This may be the case when actuating a soft tentacle arm via only moving one of its extremities.

Soft System-Environment Interactions

At the dawn of the 21st century, the concept of “morpho-functional machines” was proposed. Morpho-functional Machines were defined as machines which were adaptive by being able to change their morphology as they performed tasks in the real world [82]. In this context changes at different timescales were already argued to be important, i.e.: short-term, ontogenetic, and phylogenetic, or evolutionary. It is important to note that the adaptation and the resolution of the task is here achieved not at the control level, but at the morphological level.

As advocated by the Developmental Robotics paradigm intelligence and coordinated action are the results of complex interactions between the body, the mind and the environment. The latter, in fact, plays an important role in determining the behaviors of the artificial or natural organisms living within it.

Following this trend, one of the most influential experiments of the last two decades was the “dead fish experiment”, performed in a collaboration between Harvard and MIT in 2005 [13]. In the experiment, a dead fish was shown to be able to swim up-streams when no control impulse was clearly being sent by the brain. Upon further studies it was apparent how the streamlined body of the fish, passively oscillating, was capable of turning the surrounding energy into mechanical energy, and thus propel itself forward passively. Although the morphology and make of the body allowed the dead fish to transduce the surrounding energy, the environment was the enabling factor. The vortexes created by water streams were key in the experiment, as they generated the energy to be transduced and recreated the conditions for the body to manifest its propelling abilities. The interaction between the body and the environment were, in fact, the decisive factors in determining the observed behavior. A similar influential experiment was the passive dynamic walker. In the experiment, a walking robot was built without a controller. The make of the robot, with kneecaps, springs and pendulum-like leg swings was capable of stable, human-like and low energy, bipedal locomotion. However, the walking locomotion was initiated and stabilized by the environment itself, as it manifested when the robot was placed on a downward slope [39], thus the potential energy could be skillfully be turned into kinetic energy.

In robot design, it is therefore always necessary to take the environment into account. Much like the examples previously mentioned, the body and the brain are often not enough to achieve useful objectives. Things in the world exist to affect and change their surroundings, and live within the environment they are situated in [156]. In this context, it is in the interplay of the body and the environment that intelligent, situated, behavior can be observed, and that morphology can be empowered and purposefully adapted.

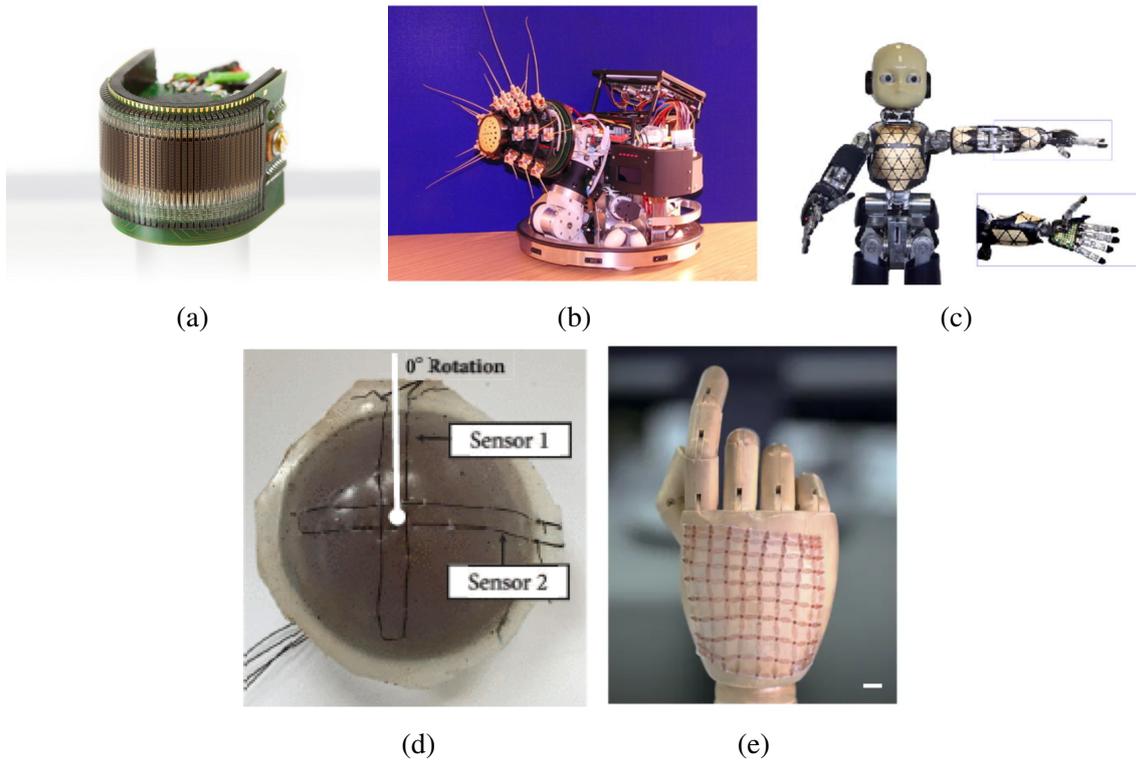


Fig. 2.4 Bio-inspired flexible and soft sensing examples. (a) Artificial compound eyes [66], (b) Robotic tactile vibrissal sensing [187], (c) I-cub robot with large-area flexible capacitive tactile skin [89], (d) Conductive Thermoplastic Elastomer's sensorized universal gripper [95], (e) Stretchable and conformable sensor for multifunctional sensing [92].

Sensor Morphology and Soft Perception

In nature, morphology plays a fundamental role within the sensing landscape, mechanically converting, filtering and amplifying sensor stimuli from the outside world, to make sense of the surrounding environment, or internal states [106, 261]. In rats and mice, for example, vibrissae, or sensitive tactile hairs, have been known to confer these mammals specialized tactile capabilities, aiding them in a number of sensory discrimination tasks [200]. In a similar light, most mammals have evolved to mediate vision through compound eyes, compromising resolution for larger fields of view and high temporal resolution, enabling fast panoramic perception [129]. Within the biomimetic robotics field attempts have been made to endow robotics systems with the capabilities of organisms observed in nature, haptic robot perception through whiskers [187] and compound vision [66] are two such examples (Fig. 2.4).

Soft sensing is one of the most popular fields within the Soft Robotics landscapes. Augmenting Soft Robotics system with the ability of sensing the environment can enable

robots to react to unknown events, to improve their control and morphology over time, and capture information or reason about entities in the world. Sensorizing soft robots is no easy task. One of the goals within this field is to devise sensors which themselves exhibit some “soft” behavioural characteristics, usually flexibility (i.e. can be bent), and stretchability [146] are desirable. Currently, approaches to achieve stretchable electronics include wavy circuits [155, 210] and conductive liquids [31]. One of the most widespread soft sensors are strain sensors, which have been shown to be highly elastic [167]. New embedding methodologies have also shown the possibility of embedding strain sensors within elastomers through 3D printing techniques. Other flexible sensing technologies, like capacitive tactile sensing [154] and optic fibers [69] have been used within Soft Robotics systems.

As previously mentioned, sensory-motor coordination and morphology can enhance the sensing capabilities of robotics systems. Sensors thus, should not be thought of simply as independent and self-sufficient technologies, but instead, it is fundamental to think of sensor technologies as apparatuses which reside within a body. The body dynamics derived from its morphological properties, coupled with the environment the robotic system is situated in, should all contribute to the sensor morphology, its characteristics and its perceptual capabilities. The appropriate coupling of these factors has been shown to be able to improve the sensing capabilities of robotic systems [107]. In [95], for example, the sensorization of a universal gripper was achieved with a pair of Conductive Thermoplastic Elastomer (CTPE) strain sensors 2.4 (d). Differential sensing was then used to compute deformations within the soft body. Morphology, however, was key. By weaving the strain sensor in different patterns within the soft gripper, information regarding the magnitude, orientation or location of a deformation could be detected. Because the sensing is also inescapably linked to motor control, the mechanical dynamics and the objectives of the robotic system, the concept of “adaptive morphology” has recently been proposed [106], where the iterative design, assembly and evaluation of sensor mythologies attempts to explain the adaptive nature of the perceptual abilities of living organisms.

2.4.2 Ontogenetics and Adaptivity

Adaptation and Growth

The principles previously discussed encourage a different approach to design, in line with endowing robots with the ability to adapt to ever changing environments, and indeed make use of the environment as a means to solve the tasks given to them. Besides design principles at a phylogenetic scale, and instantaneous deformation on the short term scale via material properties and design, another important factor is ontogenetic change and adaptation. Plants,

for example, are capable of continuously changing their morphology and physiology in response to variability within their environment, in order to survive [160]. Inspired by the unique abilities of plants to survive in diverse and extreme environments, a stream of researchers has more avidly tried to reproduce some of their adaptivity in robotics system. Plantoids, or robotic systems equipped with distributed sensing, actuation, and intelligence to perform soil exploration and monitoring tasks, have started to gain traction in this direction [160]. Root-like artificial systems in [215] and [216], for example, have been shown to be able to perform soil explorations through novel methodologies simulating growth via elongation of the tip. Other plant-inspired technologies in biomimicry and material sciences include Velcro, from the mechanisms behind the hooks of the plant burrs [267], bamboo-inspired fibers for structural engineering materials [135], or novel actuation mechanisms in [249] based on reversible adsorption and desorption of environmental humidity, and in [159] based on the osmotic principle in plants.

Another important factor in ontogenetic adaptivity is the ability for organisms to mend their own tissue over their lifespan. Endowing artificial systems with self-healing abilities has recently become of primary importance, setting the landscape for untethered robots to “survive” for longer periods of time in more uncertain and dynamic task environments. Self-healing of soft materials is typically achieved through heat treatments of the damaged areas, which allow some polymers to re-connect and retrieve most of their structural properties. In [254], for example, a soft gripper, a soft hand, and artificial muscles were developed with Diels-Alder materials [220]. In the developed systems, the Diels-Alder were shown to be reversible at temperatures of 80°C, recovering up to 98 to 99% of the mechanical properties of the polymers post damage.

Tool Use and Extended Phenotype

In Biology, the phenotype is known to be the set of observable traits of an organisms, including its morphology, developmental process and physiological properties. The idea of extended phenotype was first introduced by Richard Dawkins in [50], where he argued that the concept of phenotype might have been too restricted. In fact, the effects that a gene may have are not limited to the organism itself, but to the environment the organism is situated in, through that organism’s behaviour. The coupling of an artificial agent and its environment was discussed in Section 2.4.1. The extended phenotype notion, however, extends to even more radical concepts.

One of the most fascinating examples of this is found in primates, corvids and some fish, which have been found to be able to purposefully make and use ‘tools’ to achieve goals within their environments, such as acquiring food and water, defence, recreation or construction

[240]. Extending the phenotype concept, the observable traits of the organisms in this case should be augmented to include the extended functionalities, behaviors and morphology derived from the tool under use. When a primate is holding a small branch, for example, the physical characteristic of the primate are undeniably changed, its reach is longer, its weight and morphology is affected, as is the stance to keep balance on two or three limbs, or the ability to affect the environment around them. Under the extended phenotype concept, these as well as many other changes need to be captured within the phenotypic traits of the organism.

In the context of Soft Developmental Robotics, the ontogenetic development of robotics systems should include their ability to adapt to their environments over their lifespan (physical adaptation), and indeed the ability to augment their functionality by the active creation and use of tools, initially excluded from their phenotypic traits. This ability was previously investigated in [88] and [168] where it was obvious that at the foundation of the idea of tool use there was the concept of body schema, previously mentioned in Section 2.3.1. The body schema in this scenario requires adaptability and alterability throughout ontogenetic development, to cope with the changes in one's body, including growth, as well as with the extended capabilities conferred by the use of tools. In [168], the temporal integration of multisensory information was argued to be a plausible candidate mechanism to explain tool use incorporation within the body schema. Another core component in this context is proprioceptive sensing, or the ability to sense self-movement and body position. Previously discussed in Section 2.2.2 to be important in Soft Robotics, proprioception also plays a significant role in the Perception/Action model of body representations [51].

2.4.3 Learning Through the Body

The advancements in AI in the last two decades have begun a scientific revolution, endowing machines with the possibility of achieving super-human performance levels in several different fields, like image based object detection [222], virtual agent control [164] or haptic texture identification [64]. In robotics, machine learning has extensively been used both on the perceptual side, like object detection and recognition, and on the control side, like robot trajectory planning and motor control.

The most powerful machine learning algorithms make use of supervision, or the knowledge of target labels, to improve performance over time or trials. Broadly speaking, from the machine learning point of view, it is common to try and fit the best function to some collected data, to be able to achieve good behavior in future instances of similar data. The data could, for example, be streaming images from a camera mounted on an indoor mobile robotic platform, and the supervised machine learning module could have learned when and

how to turn the wheels left and right, based on collected and labelled visual feed in a similar indoor environment.

Throughout the sections in this chapter we have treated the concepts of morphology, with the repercussions of what is known as morphological processing, sensory-motor coordinated behavior and environment. In cases such as the one mentioned above, it is common that this interconnection of mind, body and environment is neglected. In fact, in robotics the data is usually perceptual information collected by the robot itself. As such, the perceptual information is subject to influences from the both the way in which the robot interacts with entities in the world (sensory-motor coordinated action), and the morphology of the robot's body itself. The robot can be thus be seen as a reality filter, which can act in its environment and affect the information the way that is most appropriate for learning. In this context, not only the information can be structured so to be rendered suitable for learning, but the structure information itself can guide both the morphology and the control of the robot, creating a sensory-motor and morphological adaptation loop capable of intrinsically drive the robot's behavior. 'Soft Morphological Computation' refers to the ability for the robot to understand how its own body and actions filter the information retrieved from the world, and change its configuration and interactions so to optimize information retrieval. This simplification can then drive learning and further the adaptive capabilities of autonomous robotics systems.

The concept by which information can be structured by the body has not passed unnoticed. The term 'Information self-structuring' is decades old, and advocates how that statistical structure and regularities arise in organisms as a consequence of effectively coordinated motor activity [149, 193]. Previous research has shown robots to be capable to purposefully affect the information gathered from its environment through both morphological processing, and sensory-motor coordination [189, 192]. The concept of *SoMComp* shares a similar intuition with three main differences. First, in contrast to previous research, *SoMComp* considers the soft context of the robot body or the environment, and thus concerns itself with 'soft interactions'. This is possible thanks to recent advances in material science, Soft Robotics, and our understanding of physical processes. Second, it puts together the previously self-standing concepts of morphological processing and sensory-motor coordination into a unified theory of Soft Morphological Computation. Finally, it links information conditioning to learning, and shows how learning processes can make use of embodied processes to improve baseline performance.

The ability of robotics systems to purposefully shape the sensory information through their actions, or morphology, and learn from the induced structure, has the potential to change the learning landscape within robotics systems. The reminder of the chapters will further develop this idea with specific frameworks and applications.

2.5 Challenges and Perspectives

Through this chapter we have been treating the various aspects bio-inspired robotics with emphasis on Soft Robotics and the idea that intelligence exhibits as an interplay, and reciprocal dynamical coupling, of the brain, the body, and the environment. The concept of Developmental Soft Robotics was introduced in this context, where some design principles can be established on three different time-scales, aiding and enabling roboticists and researchers to develop systems for the new generation of robotics. Many enabling technologies for sensing and actuation have driven progress in the past few decades, and have allowed robots to pass from rigid, and industrial, to soft and human-friendly. These robots have been shown to achieve locomotion, to pick and manipulate objects, to be able to safely interact with humans and much more. However, many challenges still await this field, as the road to the ultimate goal of creating machines with abilities akin to those of biological organisms is only at its early stages.

2.5.1 Evolutionary Robotics

On the phylogenetic time-scale, the question of how to achieve complex embodied behavior has been answered by nature for a very long time. The concept of evolution in biological organisms is fairly straightforward, where evolution is thought of as the change in inheritable characteristics of populations over successive generations [81]. Due to various sources of genetic variation, new generations have increasingly different traits, and by a mediating process like that of natural selection some traits will ensure higher or lower chances of survival [234]. Eventually the surviving population has all the different traits that we can now see in the immense variety of living organisms in our planet, which have adapted to use a plethora of different methodologies and techniques to ensure their survival.

The field of phylogenetics is tightly coupled with this concept, and consequently this field has a major impact in emergent design and control in robotics. In the area of “evolutionary robotics”, evolutionary computation is used to develop physical design or controllers for robots. Evolutionary computation takes inspiration from biological evolution. In robotics, for example, it is possible to create an initial set of candidate robots, and encode their physical and or control characteristics numerically. By testing the robot population against a specific task, it is then possible to identify which combination of morphology and control performed better. The encoded characteristics of the best performing robots can then be perturbed and used to create a new generation of robots which can now be tested again. The iteration of this process for thousands of iterations has been shown to be able to achieve robust controls [65, 158] and designs [142, 148, 191].

One of the biggest limitations of evolutionary algorithms lies with the resources and time necessary to be able to achieve good controllers or designs. Because iteration of robot design or robot control and robot evaluation are very time consuming, it is generally not feasible to apply evolutionary algorithms in very complex problems, by starting from a generic, non-bounded, robot characteristics' encoding. The world of simulation has historically been more suited for evolutionary algorithms [142, 158, 178] given the ease with which populations can be created, tested and iterations achieved. The controllers and designs found, however, are usually not robust real world solutions, as simulation environments are still very limited, and the solutions found within them do not necessarily correspond to solutions in the real world [108]. Moreover, depending on the complexity of the problem, computational resources are still an issue.

In Soft Robotics, given the complexity of the bodies, and the interactions emerging from them, design and control pose one of the biggest problems. Evolutionary algorithms find themselves suited as a candidate solution, but the limitations previously mentioned still apply. Further advancements in virtual reality engines, new manufacturing methods for fast prototyping, advancements in material science and the ever increasing computing power, however, may solve some of the mentioned limitations in the near future.

2.5.2 Complexity and Scalability

As of today, the robots we see still “feel” unnatural, they move slowly and sluggishly, humanoid robots still do not possess the ability to walk, run or move the way humans do, they can not reason about the world the same way we do and get confused when unknown events occur [195]. One of the several reasons contributing to this fact is complexity. The amount of actuators and distributed sensors present in humans is much too high to be replicated by motors and standard sensors in machines. This complexity poses a problem, as it does controlling the coupling of a high number of motors and sensors. Even when dealing with sub-problems, like humanoid hands, the complexity may very well be already too high to try and tackle with standard methods. Some attempts to replicate complexity have already been made, for example, by replicating in a robotic manipulator many of the degrees of freedom present in a human hand [265]. This approach, however, did not give the results many were hoping for, as complexity in the body was coupled with complexity in the control, and achieving adaptable, smooth grasp and manipulation behavior was no easy task. Recent advances have shown how under-actuated or even passive hands can achieve complex behaviors if the interactions with the environment is appropriately exploited [94, 97]. It is here that complexity can be avoided, since complex behaviour can emerge from simple design, when appropriate interaction takes place.

Within this framework, many questions still remain. It is, in fact, unclear how design should be achieved so that complexity can be avoided or exploited. Exploiting environmental constraints is no easy feat, as the constraints to be exploited are also tightly linked to the task at hand. In Soft Robotics, the softness of the robot themselves induces highly non-linear behaviors and complex dynamics. Paradigms like that of reservoir computing can capitalize on the complexity of such structures, using them as a computational resource and thus making complexity a desirable feature. Control, however, is still hard to achieve, and mathematical models fail to comprehensively account for dynamical interactions when the complexity of the body becomes too high. This complexity presents infinite challenges and opportunities, and the ever changing landscape of robotics will have to face many of them in the near future.

Chapter 3

A Case Study of Soft Interactions: Robot Maturity Assessment

This chapter reports a case study to show the complexity of soft interactions in a real world robot manipulation task. The chapter further highlights both the pro and cons of common model based approaches to capture these complex interactions for haptic perception.

Reference Publication

This chapter was adapted from an article titled “**Non-Destructive Robotic Assessment of Mango Ripeness via Multi-Point Soft Haptics**” [230] published in the International Conference on Robotics and Automation (ICRA) in 2019 [230]. The article was written in collaboration with Dr Perla Maiolino and Dr Fumiya Iida and proposes a novel model based tactile procedure to assess mango ripeness in a non-destructive manner. I am first author in the article and my contribution includes the design and application of the experiments, data analysis and writing of the article. Dr Fumiya Iida contributed with the conceptualization of the topics, experiment design and paper writing. Dr Perla Maiolino contributed with the conceptualization of the topics, the design of the mechanical model for the sensor, the sensor technology and paper writing.

3.1 Introduction

3.1.1 Mango Ripeness, a Case Study

With ever increasing demand for high quality horticultural products, and an increase in their acceptable marketing standards, there is a need to find fast, reliable and autonomous processes which can provide these guarantees [115]. The quality of crop is dependent on several pre-harvest factors, among which weather conditions, growing land, irrigation patterns, chemical treatments and others still [217]. After harvesting, an important characteristic determining quality, and which has a direct impact in the marketability of the produce, is ripeness [114]. Besides appropriate harvesting time, determining the ripeness of horticultural produce is useful for classification, transportation, handling and the security of its quality.

In crop such as bananas, strawberry, watermelons and tomato, visual cues have been shown to be sufficient in assessing ripeness and classifying produce accordingly [58, 166, 206]. Other fruits like kiwis, blueberry or mango, however, do not provide useful visual diversity between ripe and unripe units. In this paper we focus on ripeness assessment of mango, a high value agricultural and food product, which shows different ripeness visual cues depending on its variety. Given the large variety of mangoes (over a 1000 only in India), assessing ripeness by machine vision is discouraged.

The current approach for ripeness assessment of mango and similar fruits is by measuring their firmness through a penetrometer instrument [1, 20]. A penetrometer is a pressure tester with a cylindrical head, which is usually inserted in the pulp of test fruit at a set depth, and a speed approximately controlled by an operator. The firmness values of fruit tested by a penetrometer may largely vary depending on the instrument's user [23]. Moreover, penetrometer testing is a destructive procedure, as the post-measurement fruit damage is irreversible, and the tested fruit must therefore be discarded. Finally, the firmness distribution of the pulp of fruit across its surface may largely vary, and only successive single penetrometer tests in different locations can insure a thorough firmness assessment.

In this paper, we propose a method for testing the ripeness of mango by means of touch. The proposed method is non-destructive, and allows the user to test multiple surface areas with a single touch. The method designed is possible given the use of capacitive tactile sensing technology, endowing end-effectors with the ability to retrieve multiple contact pressure readings in relatively small areas. We make use of a custom-made probe equipped with a capacitive tactile sensor array in order to palpate the fruit and thereby assess its ripeness. We model the mango's pulp and skin behavior when deformed through a simple spring system and thus retrieve fruit firmness as a stiffness measure.

This paper is structured as follows. Section 3.1.2 presents a brief review of the current technologies available for ripeness assessment of horticultural produce. In Section 3.2 the methods are explained, including the theoretical framework for ripeness assessment. In Section 3.2.1, the gripper design in Section 3.2.2, and the tactile sensor technology used in Section 3.2.3. In Section 3.3 the results are reported and finally a conclusion is given in Section 3.4.

3.1.2 Approaches for Ripeness Assessment

In the past few decades, with advancements of machine learning and vision, systems have been made showing it is possible to assess the ripeness of horticultural produce by visual cues [58, 166, 206]. These systems, however, are limited to produce which show differences in chromacity at various ripeness stages. For mangoes specifically, systems using machine vision have previously been explored [173]. Given the large variety of mango types, however, these solutions were only limited to specific families.

Recently, spectral techniques have been used to assess the quality of post-harvested produce [104, 110, 243]. More specifically, Raman imaging, Fluorescence imaging, Laser backscattering imaging, Hyperspectral imaging and Nuclear magnetic techniques have been shown to be able to classify produce lacking chromaticity differences into various stages of ripeness. The equipment required for said methods, however, is usually bulky and the information processing often computationally intensive, making it hard to create solutions which can be exported in the field, or do not require the transportation of produce to appropriately equipped areas.

When assessing ripeness, consumers use a combination of tactile sensing and visual cues. The physical probing of produce is indeed one of the oldest modalities for ripeness assessment, and brought the advent of penetrometer measurements [20], non the least because of the instrument's ease of use and transportability. Recent years have seen a rapid development of robotics technology in the context of agriculture, mainly related to transport and harvesting [62, 85, 98, 176]. Moreover, technological advances in tactile sensing and perception [43, 87] have changed the landscape for tactile based inference procedures [196, 227]. In this context, however, robotics solutions for post-harvest quality assessment remain a largely unexplored area.

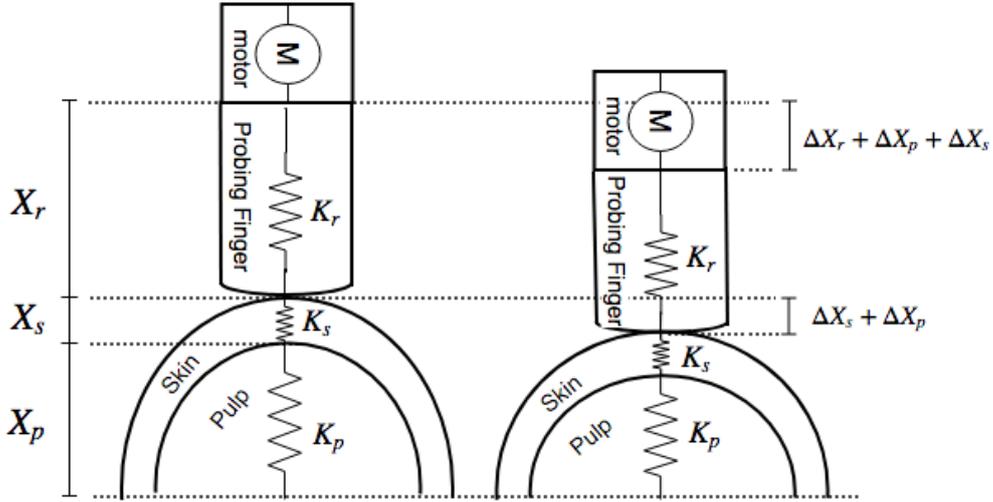


Fig. 3.1 Stiffness model.

3.2 Methods

3.2.1 Stiffness Model

We propose a palpation procedure to assess the ripeness of mango, and devise a simplified system to model the fruit's pulp and skin behavior throughout palpation. In the model we exemplify the scenario where a finger, equipped with a force sensor, is in contact with the surface of a mango. The finger is actuated by a motor, and its displacement is known by means of a motor encoder. Fig. 3.1 shows the modelled elastic response of the probing finger, and the object's surface, as system of springs. We choose a linear model as the simplest mechanical model of the fruit, and make the simplifying assumption that each spring is constant. The probing finger has a spring stiffness constant of K_r , while the mango can be seen as a two layered structure, the first layer of which consists of the spring response of the skin, with a spring constant of K_s , and the second the spring response of the pulp, with a spring constant of K_p . The lengths of each are also respectively X_r , X_s and X_p . The estimation of the produce's stiffness is equivalent to retrieving the elastic constant K_e . The motor generates a torque capable of directly influencing the distance between the finger and the produce. At equilibrium, the forces generated by the probing fingers F_r equate the reacting forces from the produce's surface F_e , thus $F_r = F_e$, i.e.:

$$F_r = \left(\frac{1}{K_s} + \frac{1}{K_p} \right)^{-1} (\Delta X_p + \Delta X_s). \quad (3.1)$$

As the skin of mango is much stiffer than its pulp, when applying a small displacement through the motor, it is useful to make the simplifying assumption that $K_s \approx \infty$, therefore $\Delta X_s = 0$ and

$$F_e = K_p X_p . \quad (3.2)$$

Finally, the motor displacement as computed by the encoder corresponds to $\Delta X_m = \Delta X_r + \Delta X_p$. So from equation 3.2, and the simplifying assumption we have:

$$K_e(1 + \varepsilon) = \frac{F_r}{\Delta X_r + \Delta X_p} \quad (3.3)$$

$$\text{where, } \varepsilon = \frac{\Delta X_r}{\Delta X_r + \Delta X_p} \quad (3.4)$$

i.e. when the motor displacement is large and/or the compression of the sensor ΔX_r is much smaller than that of the mango ΔX_p , ε can be neglected.

3.2.2 Probing Gripper Mechanism

Fig. 3.2 shows the gripper used for the experiments. The gripper is composed by a main rectangular case containing a lead screw and two metallic rods. The chamber contains two opposite fingers, which remain parallel to each other throughout the gripper's actuation. We designed two fingers with flat surfaces at the extremities, capable of holding the referenced tactile sensor. A central actuation unit reduces the distance between the fingers by actuating one finger through a Micro Metal Gearmotor motor, with a 6:1 gear ratio and equipped with a rotary encoder. The rotational actuation movement is then transferred into a linear displacement by the lead screw and metallic rod. We control the motor via a TB6612FNG Dual Motor Driver Carrier controller. Each gripper component was 3D-printed, for fast prototyping.

3.2.3 Tactile Sensor Technology and Data Acquisition

The tactile sensor (CySkin) used for the experiments is described in [223].

The adopted sensing mode is based on the capacitive transduction principle. A capacitive transducer (i.e., a tactile element, or *taxel*) is organized in a layered structure: the lower layer consists of the positive electrode, which is mounted on a Flexible Printed Circuit Board (FPCB). The dielectric for the sensor is here fundamental. The deformation of a too soft dielectric layer, like air, may reach its saturation before inducing any deformation in the pulp of a mango. From equation 3.4 it is clear how the deformation of the mango surface must be

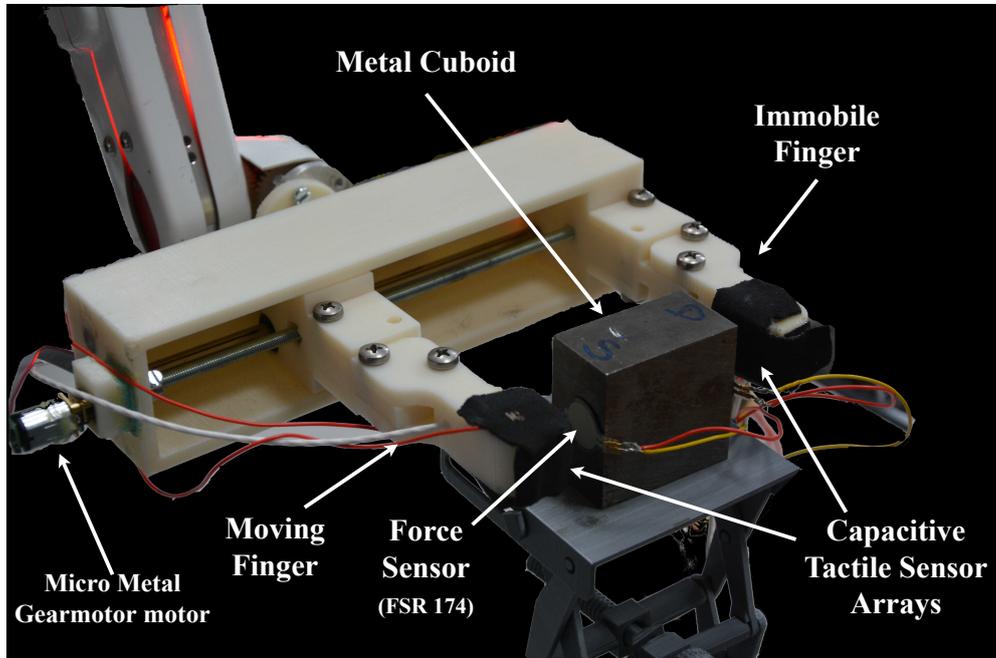


Fig. 3.2 Gripper, and force calibration set-up.

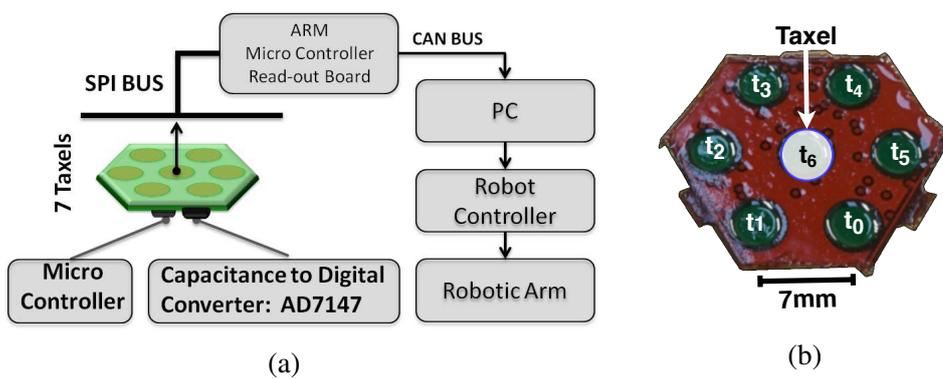


Fig. 3.3 (a) The CySkin technology architecture. (b) The CySkin patch used for the experiments.

grater than that of the sensor, for the assumptions to hold. Knowing typical mango firmness ranges², we choose a 3D-printed dielectric layer of $2mm$, composed of VeroBlack rubber with A-27 Shore coefficient.

In the experiments we use 2 hexagonal shaped modules, each placed in the inner flat extremity of a finger in the gripper (Fig. 3.2). Each module hosts 6 taxels (Fig. 3.3b), as well as the Capacitance to Digital Converter (CDC) chip (namely, the AD7147 from Analog Devices) for converting capacitance values to digital. The CDC chip can measure variations in capacitance values with 16 bits of resolution. All the modules are interconnected and communicate through an SPI bus to a read-out board which performs a preliminary processing of the tactile sensor data and send them to the PC through CAN bus (Fig. 3.3a) with a sensitivity of $0.32pF$. In this context, the normal forces exerted on the sensor produce variations in capacitance values reflecting the varied pressure over the taxel positions. A sensor reading, or tactile image, from the tactile sensors described is produced at $20Hz$, and corresponds to two 6-dimensional arrays, where each element contains the capacitance variation value of the corresponding taxel within each finger in the gripper. Here and for the remainder of the experiments, each taxel is considered a separate tactile sensing unit.

3.2.4 Robot and Experimental Set-Up

To perform the mango experiments the end-effector, coupled with the tactile sensor, was mounted onto an ST-Robotics R12/5 arm³. The R12/5 robotic arm was controlled open-loop in Cartesian coordinates. A teach-pendant was used to manually teach the robot the starting position with the arm and end-effector facing forward (Fig. 3.4a).

After the robot arm reaches the starting position, and a test produce is placed within the end-effector's reach, the probing experiment consists of three stages: a reaching, a probing and a release stage (Fig. 3.4b).

In the first phase, the gripper's moving finger is driven towards its immobile counterpart at $1mm/s$. When any taxel, in both fingers, reads values above 5% of their maximum calibration, a touch is detected and the gripper is stopped.

In the second phase, the mobile finger is further actuated to close the gripper until either of two conditions are met: first, the last touching taxel has moved of at least $1mm$ into the flesh of the fruit; second, the encoders do not change value over two consecutive readings (i.e. the motor is at equilibrium at its maximal gripping force). The slow motion induce quasi-static interaction validating the model's static assumptions. The depth of $1mm$ was

²https://www.mango.org/wp-content/uploads/2017/10/Mango_Maturity_And_Ripeness_Guide.pdf

³<http://www.robotshop.com/uk/st-robotics-r12-5-axis-articulated-robot-arm.html>

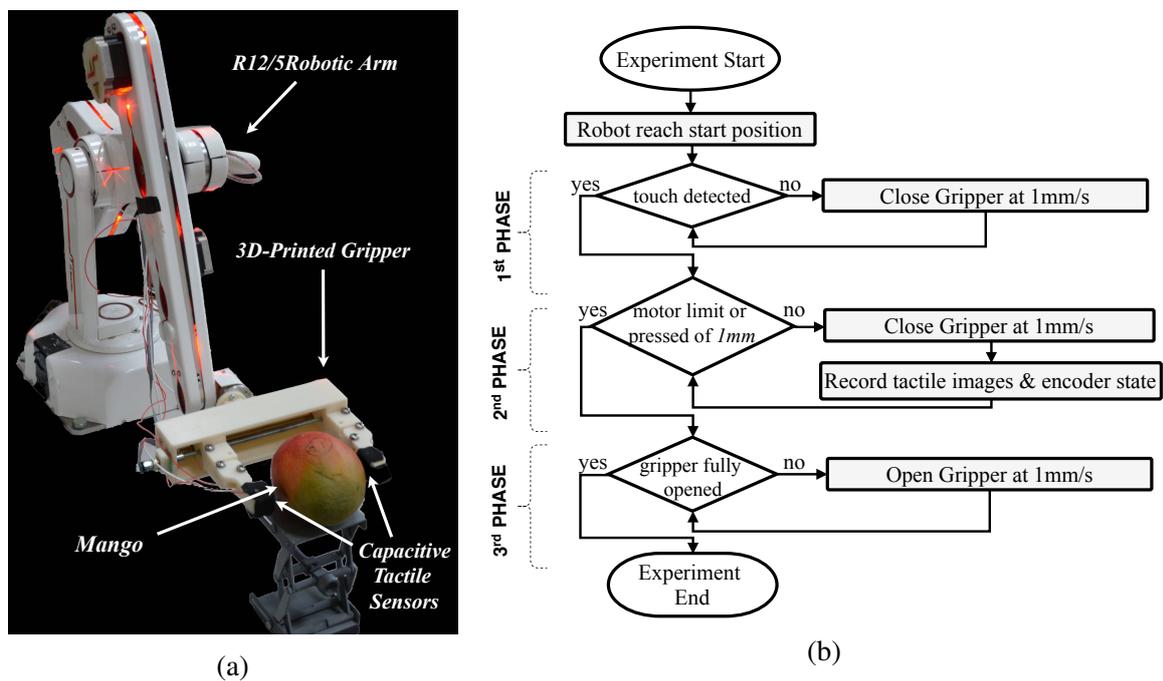


Fig. 3.4 Experimental procedure, including (a) the set-up for the experiments and (b) the Flowchart of the touch experiments.

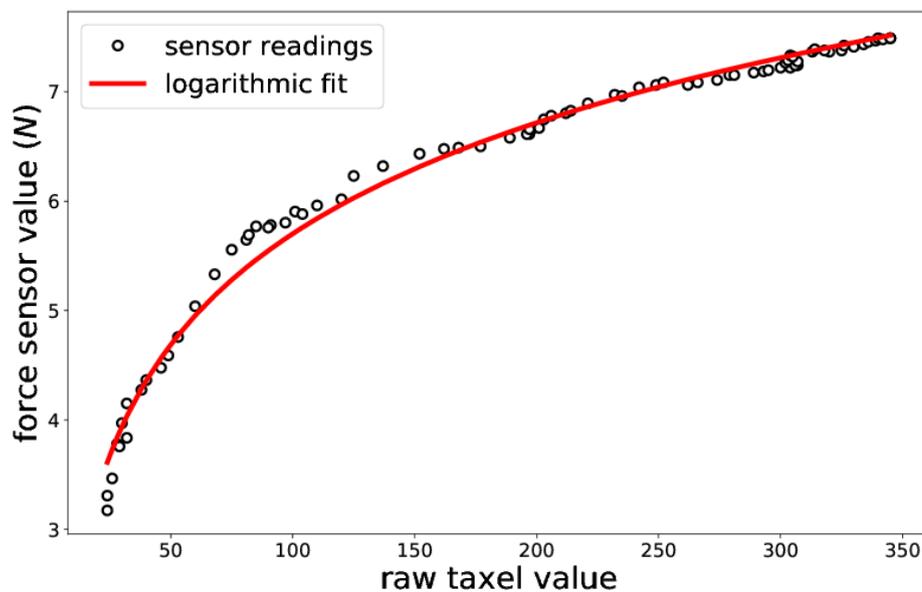


Fig. 3.5 Raw taxel value to force logarithmic fit on taxel 3.

chosen to induce enough deformation in the mango, while working within the linear range of the sensor (Fig. 3.5, force range $[0, 5.5]N$).

In the third phase, the gripper mobile finger is reversed at $1mm/s$ to the gripper's fully opened state (Fig. 3.4b).

The system has been implemented in MATLAB, synchronizing the gripper control and the sensors acquisition. The tactile images are thus recorded at $\approx 15Hz$ throughout the run of the second stage and later used to retrieve the stiffness of the touched produce.

3.3 Results

3.3.1 Sensor Force Calibration

To make use of the theoretical spring model we converted the tactile sensor pressure response, of each taxel, in force. To achieve accurate conversions, the end-effector coupled with the tactile sensor was made to close onto a stiff metal cuboidal object. The opposite surface areas of the cuboid, in contact with the end-effector's fingers, were covered by two force sensitive resistors FSR 174 sensors, previously calibrated to measure forces in $0 - 10N$ range with an accuracy of $0.01N$ (Fig. 3.2). In the experiment, the gripper was actuated to close at $1mm/s$ until motor torque limit. The force sensor response and corresponding taxel values were sampled at $\approx 15Hz$. Given prior knowledge of the dielectric layer deformation behavior [153], we fit a logarithmic curve of the form $f(x) = a \log(bx - 1) + c$, mapping the pressure response to the punctual forces registered during the calibration. The a , b , and c parameters were optimized by least squares. Fig. 3.5 shows an example force fit for taxel 3, all other taxels were similarly calibrated.

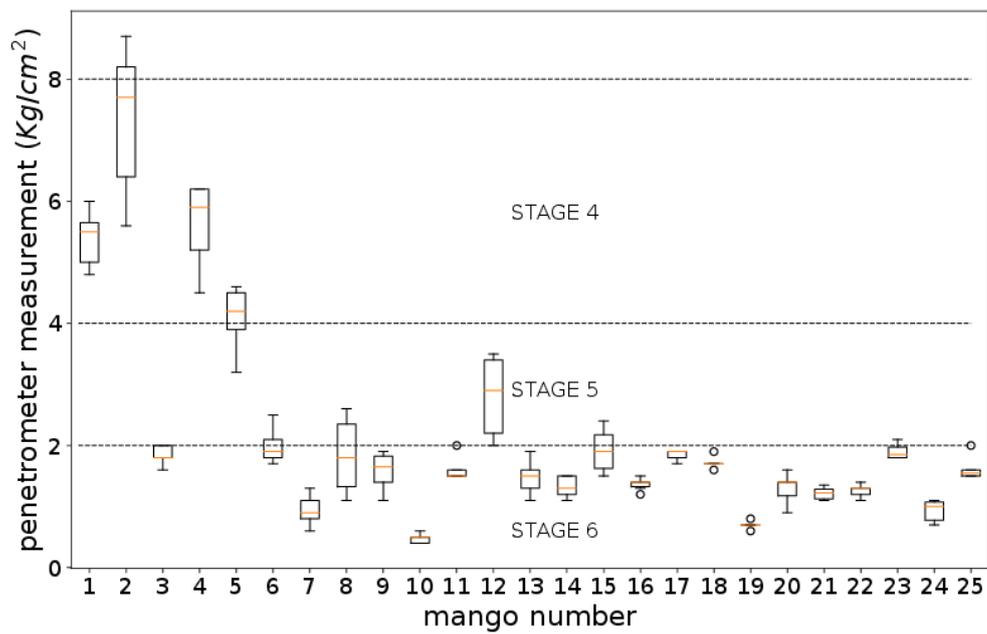
3.3.2 Mango Penetrometer Testing

A set of 25 mangoes of the Keitt variety were used for testing. The mango were divided in three subsets and made to ripen at room temperature for 1, 3 and 5 days, increasing the ripeness differences amongst subsets (Fig. 3.6a).

We use a penetrometer instrument to retrieve ground truth mango firmness measurements. Penetrometer tests were done following industrial standards. The skin of each mango was removed before the measurement, and pressure was applied to reach the penetrometer's head full insertion in approximately $2s$. A total of 10 measurements were done on each mango. Fig. 3.6b shows the penetrometer measurement values for each fruit in the test set. Mangoes 1-9 were tested on day one, mangoes 10-18 were tested on day two and mangoes 19-25 were tested on day five. Comparing the measured values to standard ripeness levels⁴ 20 mangoes



(a)



(b)

Fig. 3.6 (a) Mangoes used for the experiments at purchase time. (b) Penetrorometer measurements of each mango when tested at a distance of 1, 3 and 5 days from purchase.

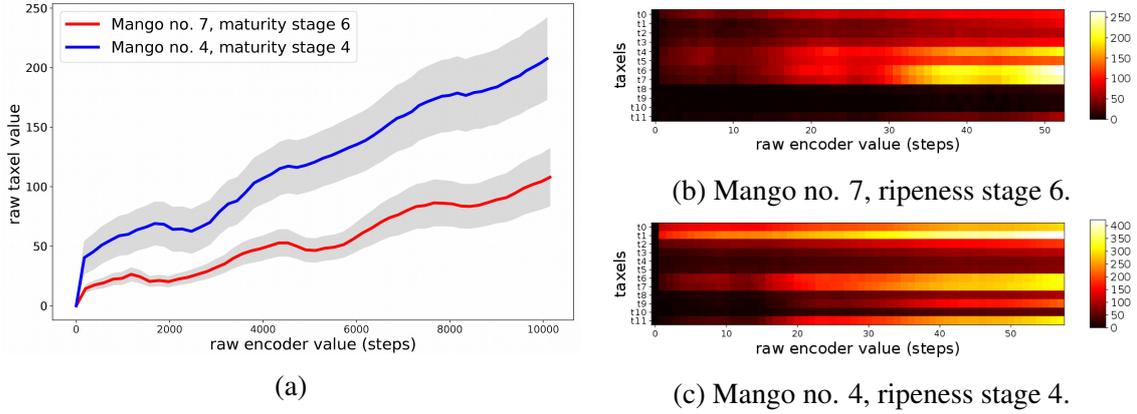


Fig. 3.7 (a) Mean and error of all raw tactile sensing units, (b) raw tactile sensor response, when performing the touching experiment on mango number 7, at ripeness stage 6, and (c) a mango number 4, at ripeness stage 4.

were found to be at ripeness stage 6 (very ripe), 2 mangoes at ripeness stage 5 (ripe) and 3 mangoes at ripeness stage 4 (non-ripe). Fig. 3.6b shows how penetrometer tests, even when the instrument is operated by the same user, are somewhat variable (average standard deviation of penetrometer measures $\approx 0.287Kg/cm^2$), and do not always clearly collocate a mango in a ripeness stage.

3.3.3 Stiffness Measurement Analysis

We analyze whether from the raw tactile information is possible to dissociate between ripe and unripe mangoes. Fig. 3.7 shows the raw tactile sensor response, when touching fruits at ripeness stage 4, and 6 (as determined by penetrometer testing). In Fig. 3.7a it is clear how the ripeness stage information is captured by the sensor response. Moreover, we observe how different taxels activate at different times and with different intensity depending on the mango. The variability is mainly due to the curvature of the fruit against the sensor's flat surface.

During the experiments with each sample, the recorded sequential sensor response was used, together with equation 3.3, to retrieve the stiffness K_e of the mango's pulp (Fig. 3.8). We compare the computed stiffness of each mango against average penetrometer measurements in two scenarios: one, as an average of the stiffnesses computed by all sensing units (Fig. 3.9a), and two, as an average of the four taxels registering the highest change in pressure over the course of each touch experiment (Fig. 3.9b). Fig. 3.9 shows how it is only possible for the sensor (y-axis in figure) to separate between ripe and non-ripe mangoes when given the opportunity to choose reliable tactile sensing units, depending on the touch experiment. The

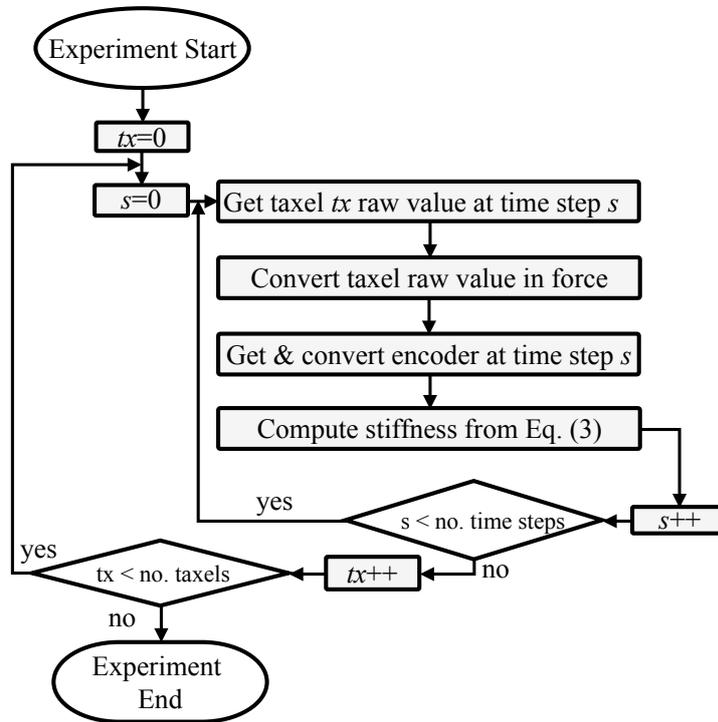


Fig. 3.8 Flowchart of the stiffness extraction processing.

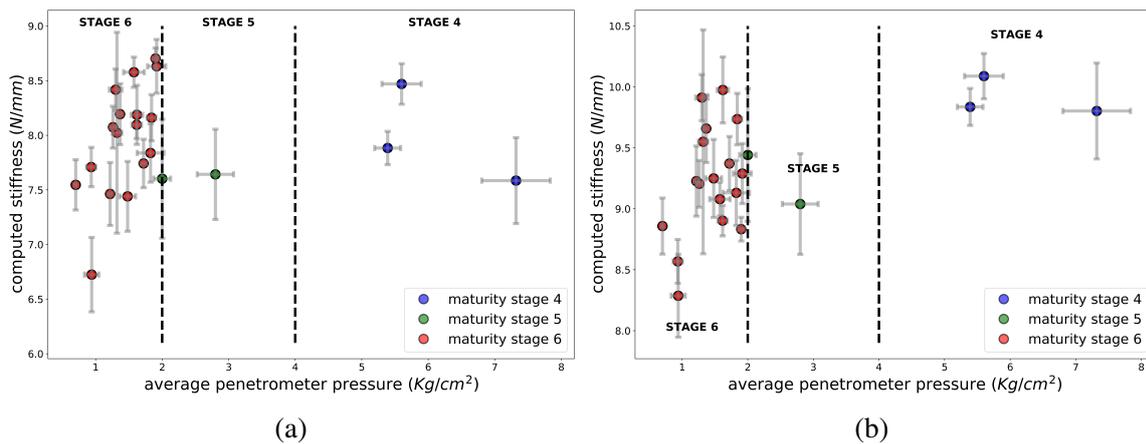


Fig. 3.9 (a) Computed stiffness averages from all sensing units and (b) stiffness averages of the 4 taxels recording the highest change in pressure throughout the experiment, against average penetrometer measurements. Each point in the plot is a different mango in the dataset.

Table 3.1 Results Summary

STAGE	No.	Correctly classified	Misclassified	Accuracy(%)
5&6 (ripe)	22	19	3	86
4 (non-ripe)	3	3	0	100
Overall Accuracy(%)				88

result emphasizes the need of sensors capable of drawing multiple measurements at once, as any one measurement might be unreliable in its ripeness estimation. Moreover, the sparse sensor response for mango at ripeness stage 6 (Fig. 3.9b), suggests that the skin of mangoes has a non-linear influence on the measured pressures. Given typical mango skin thicknesses, such influence is negligible for non-ripe mangoes, where the stiffness of the pulp dominates the sensor readings. For ripe and very ripe mangoes, however, the stiffness of the skin, at times, induces the stiffness of the whole fruit to be much higher than its pulp.

For the final estimation of mango stiffness after contact, we thus limit the computation on an average of the four taxels measuring the highest change in pressure over the course of the touch experiment's second stage. In this context, the use of multiple taxels for stiffness estimations allows for the dismissal of outliers generated by the curvature of the fruit. We consider non-ripe the mangoes whose stiffnesses is in the range $K_e > 9.7$, and ripe the mangoes those where $K_e < 9.7$. The ripeness threshold was chosen to maximize accuracy over the tested fruit. Table 3.1 reports the thresholding results. In particular, we find we can classify 88% of the tested fruit correctly, and accurately detect all the non-ripe samples in the tested mangoes.

Finally, with respect to mango ripeness estimation, it is noteworthy to notice the distribution of classes in Fig. 3.9b. In the figure, in fact, it is clear how when separating stage 4 and stage 5-6 mangoes, an equal number of mangoes (approximately three) falls within the "non-ripe" category. With the decision boundaries set as previously described, our approach misdetects three of the ripe mangoes as non-ripe. This, however, does not heavily impact the overall accuracy of the model, as the density of ripe mangoes below the chosen threshold is much higher than the ones above it. In other words, given the density distribution of mangoes at stage 6, it is less likely for these to fall in a range $K_e > 9.7$ than for those belonging to stage 4, and thus classification should be possible. It is however worth noticing that this statistical argument is not necessarily shown for the mangoes at stage 4, as we observe only a few elements belonging to this class. Further experiments are here encouraged for future directions and possible commercialization.

3.4 Conclusions

Given the lack of standard, non destructive and non-chemical tests for assessing the ripeness of fruit, we devise a ripeness testing method based on capacitive tactile sensing technology. We devise a custom made gripper, supplied with 12 capacitive tactile sensing units distributed homogeneously over two fingertips. We perform experiments by which the gripper close onto the flesh of test fruit until a pre-set depth is reached, while recording tactile image sequences. The tactile image sequences, together with a spring stiffness model, are used to retrieve the stiffness of the palpated fruit and assess its ripeness. We test the proposed method on a set of 25 mango fruit of the Keitt variety. We compare the whole fruit stiffness computations to pulp ripeness measurements based on a standard penetrometer instrument. Results show that the tactile based ripeness assessment method is capable of classifying mangoes into ripe or non-ripe, with accuracies increasing with the stiffness of the pulp. Moreover, since the proposed method, based on capacitive tactile technology, hinges on the relationship between flesh stiffness and ripeness of target produce, we argue the method is valid for other types of horticultural produce showing such relationship, e.g. tomatoes, grapes, apricot, cherries, kiwis and others besides, some of which may present difficulties for visual based ripeness assessment.

As the scale of penetrometer measurements may vary depending on the user, the proposed method presents clear transferability advantages for testing ripe over non-ripe mangoes. Moreover, the method is non-destructive, the sensor technology utilized can test several surface locations at once, the gripping technology is light thus can be mobile, and its usage does not require any specialized expertise.

Finally, in the context of this thesis, we have shown how model-based approaches can get far in real-world settings, by embedding knowledge of the world within the robotic solution of the task to achieve. Even so, there are several issues and limitations with these approaches when trying to approximate the soft interactions arising from the robot and the environment. A non-linear effect to the stiffness estimation, due to the intact skin of the mango, was in fact observed. The effect makes the distinction between stage 5 and 6 mangoes almost impossible within this framework, as may be the discrimination of several other consecutive stages. One of the reasons behind these results is the simplicity of the model, which does not capture the dynamics of the complex interactions between the robotic end effector and the produce. The skin layer, in fact, was assumed to be infinitely stiff. This however, is an oversimplification, and a more complex non-linear model should instead have been adopted to capture the stiffness of the skin. Issues, however, might occur even then, where the distribution of the stiffness across different sections of the fruit might be different, as well as the pulp's stiffness whose depth to the bone will change depending on the location and the application of the

experimental touch. These issues can be in part solved by increasing the complexity of the model and fitting some parameters through experimentation. This approach, however, is eventually bound to fail, where any model is but an approximation and errors can be minimized but hardly ever they can be entirely discarded. In capturing soft interactions, model free approaches have the advantage of basing the model itself on experimentation, and thus they are not bound by the limitation of human assumptions and design. These also present several issues, which usually discourage purely model-free approaches. These will be further discussed in the next chapter.

Chapter 4

Proprioception in Soft Continuum Bodies

This chapter presents a case study for the principle of proprioception. Within *SoMComp*, Soft Proprioception is the necessary condition for any robot to be able to purposefully make changes to its body morphology or action, so to achieve the conditioning of the soft interactions arising from the robot in its environment.

Reference Publication

This chapter was adapted from a published Journal article titled “**Model-free Soft-Structure Reconstruction for Proprioception using Tactile Arrays**” [229] published in the journal *IEEE Robotics and Letters (RA-L)* in 2019. The article was written in collaboration with Dr Perla Maiolino, Dr Josie Hughes and Dr Fumiya Iida and proposes a novel model-free sensorization method to achieve proprioception in continuum soft robots. I am first author in the article and my contribution includes the design and application of the experiments, the formulation of the learning framework, data analysis and writing of the article. Dr Fumiya Iida and Dr Perla Maiolino contributed with the conceptualization of the topics, design of the experiments and the writing of the article, Dr Josie Hughes contributed with the creation of the physical set-up and review of the article.

4.1 Introduction

The advent of Soft Robotics has changed the robotics landscape, enabling rapid and low cost prototyping and providing resilience to external disturbance and internal failures [120, 155]. One key remaining challenge is the extrinsic and intrinsic sensing of soft bodies to provide environmental information which is key for complex environmental interaction. Considerable work has focused on the development of tactile sensing approaches [96, 227] and soft manipulators for environmental exploration [97]. There has been minimal investigation on how the inherent compliance of soft structures can be exploited to achieve environmental awareness [107, 180].

The use of fully soft continuum body structures for manipulation has been demonstrated with the creation of octopus tentacle systems [29, 269]. These adaptable manipulators take advantage of the intrinsic compliance of soft structures, exploiting environmental interactions in the process. However, the potentially infinite degrees of freedom offered by soft continuum body structures make it challenging, and often impossible, to accurately determine their spatial configuration, or proprioceptive awareness [237]. For environmental interactions to be understood in this context, it is necessary to develop sensing techniques capable of reconstructing the configuration of a soft continuum body. Research on the configuration reconstruction of such bodies has mainly been driven by tactile [44, 147, 268] and medical applications [174].

Historically, exteroceptive sensing technologies, such as cameras, have been predominantly used for soft body shape reconstruction [54, 133, 270]. These methods, however, are inapplicable in scenarios where it is impossible or impractical to set-up sensing units external to the soft body. More recently, sensors which are capable of measuring curvature and bending have been used, with early work showing how fiber optic curvature technology can be used to sense bending and twist [46]. Electroactive polymeric sensors have also been used to sense bend angles and rates in prostheses [18], while in [127], fluid-resistive bending sensor were developed for flexible pneumatic balloon actuators. Additionally, work in strain sensor technology has shown how proprioceptive curvature information can be retrieved and used for partial reconstruction of soft continuum bodies [35, 43]. Despite the progress, integration into soft robot technologies is limited. Additionally, the technologies developed allow discrimination of different preset states, such as bending, twisting and pushing, but do not obtain sufficient information for the full configuration reconstruction of a soft body. Recent work has shown how flexible force sensors can be embedded in Soft Robotics manipulators, obtaining information useful in task such as grasping and object recognition [90, 91]. However, this force sensing technology allows the curvature in only one axis to be measured.

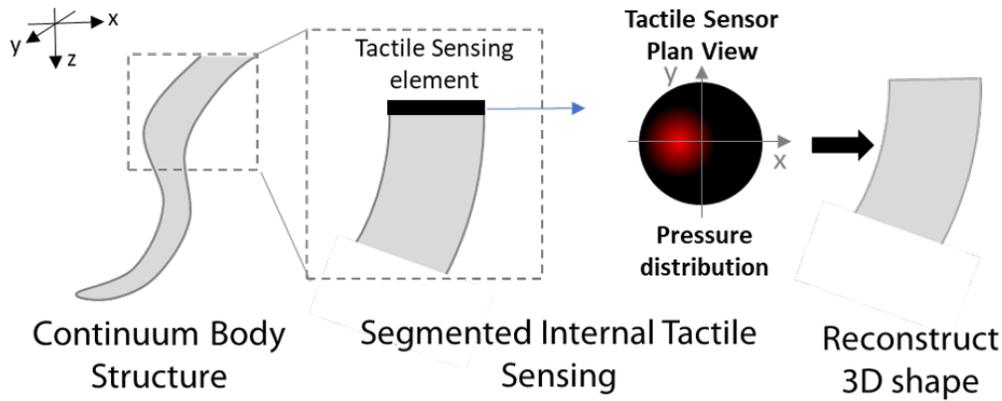


Fig. 4.1 Role of tactile sensing to allow reconstruction of passive systems. The deformation of the continuum body structure leads to internal pressure distributions which allow 3D reconstruction.

To solve the proprioceptive problem, it is necessary to devise a method capable of reconstructing the spatial configuration of continuum soft materials over all axis of deformation, while maintaining the soft body characteristics (e.g. stretch and bend). We propose that by integrating a tactile sensor array at the base of a tentacle-like continuum soft body, it is possible to use the distributed change in pressure over the surface of the sensor, induced by the change in posture of the body, to retrieve the 3D position of the tentacle end-point. Through this method, the position estimation can thus be used to reconstruct the shape and configuration of the soft continuum body. This is summarized in Fig. 4.1.

To demonstrate these hypotheses we sensorize a soft tentacle-like body segment using capacitive tactile sensing technology. Using a simple feed-forward Neural Network, the mapping from the spatial response of the tactile sensor and the deformation of the continuum body can be obtained. This allows deformations to be sensed along all axes in 3D space, whilst maintaining the *soft* properties of the body (e.g. bend and stretch). By understanding the body structure, we demonstrate how this can enable exploration and reconstruction of the system's work-space.

The ability to understand the shape of a continuum body structure has the potential to impact work including medical robotics, soft robotic exploration and enables control of continuum structure which was previously not possible. To the authors knowledge, this is the first soft-robotic sensorization method based on capacitive tactile sensing technology, which is capable of sensing accurate deformation along all three dimensional axis in space.

In this paper, Section 4.2 presents the methods developed, including the capacitive tactile sensing technology, continuum structure design and experimental set-up. The results are reported in Section 4.3 which characterize the performance of the system. Additionally, a

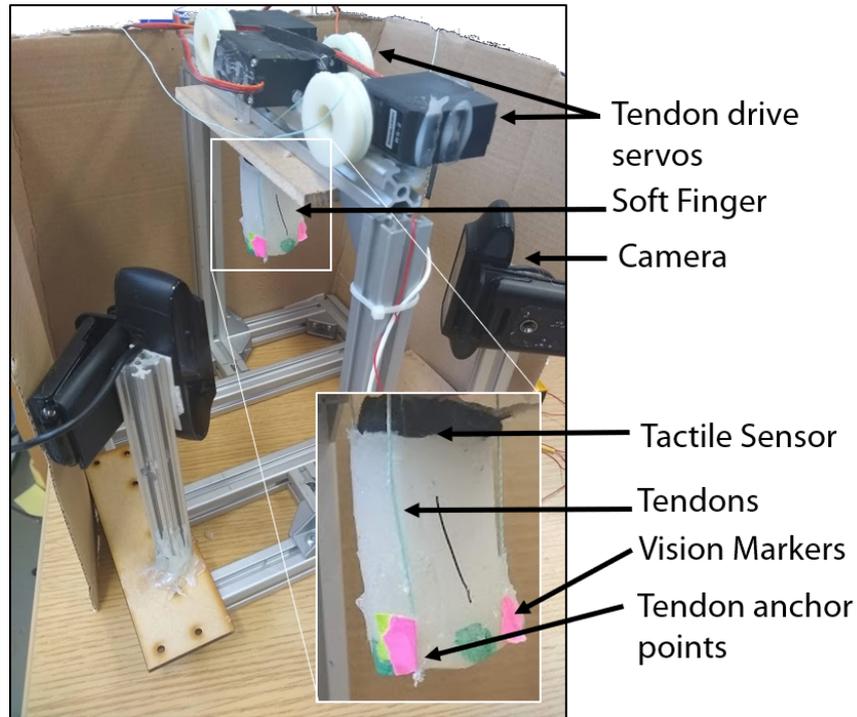


Fig. 4.2 Experimental setup showing the flexible soft body, location of the sensor, vision tracking system and the servo control and tendon system.

case study of environmental exploration through body posture inference is presented. The paper concludes with a discussion in Section 4.4.

4.2 Methods

4.2.1 Tactile Sensing Technology

The tactile sensor used for the experiments is described in Section 3.2.3. Differently than before, however, a small air chamber act as dielectric and the upper layer is a ground plane made with conductive Lycra. In the experiments we use an hexagonal shaped sensor module, hosting 7 taxels (Fig. 3.3b). Given the sensor design, a sensor reading, or tactile image, from the tactile sensors described is produced at $20Hz$, and corresponds to a 7-dimensional array, where each element contains the capacitance variation value of the corresponding taxel within the patch.

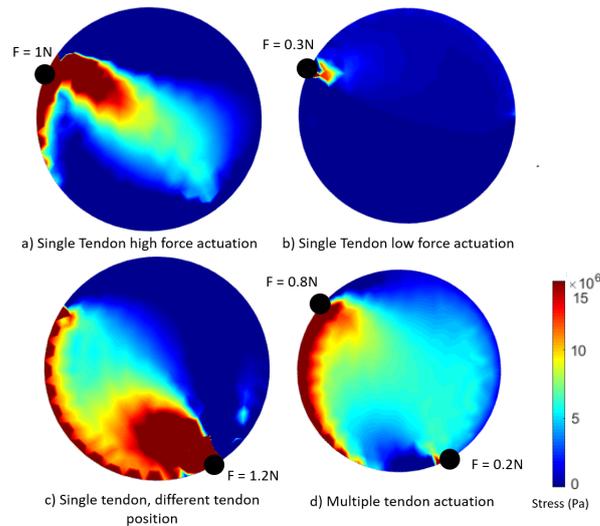


Fig. 4.3 Results from FEM experiments showing the stress distribution at the top of the sensorized body. (a), (b) and (c) show the resulting stress distributions when actuating a single tendon with a force of 1N, 0.3N and 1.2N respectively. (b) shows the stress distribution a two-tendon actuation, applying a force of 0.8N and 0.2N in two different locations.

4.2.2 Soft Continuum Body Design and Sensor Embedding

A continuum soft body segment has been developed by casting EcoFlex 00-30 silicone in a 3D printed mould. This soft ‘finger’ (height = 50cm, radius = 15mm) is controlled by three tendons equally distributed around the finger, and which allow full position control of the soft body.

The capacitive tactile sensor module described in Section 4.2.1 is placed at the base of the cylindrical finger. Thus, the capacitive sensor taxels are uniformly distributed along one of the circular surfaces of the elastomeric finger. The sensor placement allows pressure patterns to be sensed at the top of the finger when deformations are induced along its body.

To perform the experiments, we devised a set-up by which the soft continuum finger is suspended at the top of a cubical frame (Fig. 4.2). The hexagonal tactile sensing module is thus placed between the base of the finger and the top metal beam of the cage. Three servos are placed above the top beam, each connected to a tendon to allow the continuum structure to be deformed. Two cameras are mounted in adjacent corners of the frame. The cameras face the finger in orthogonal directions, and perform visual tracking of markers on the finger’s head, to reconstruct the posture of the continuum body (Fig. 4.2). Both visual tracking and tactile images are logged synchronously throughout the experiments.

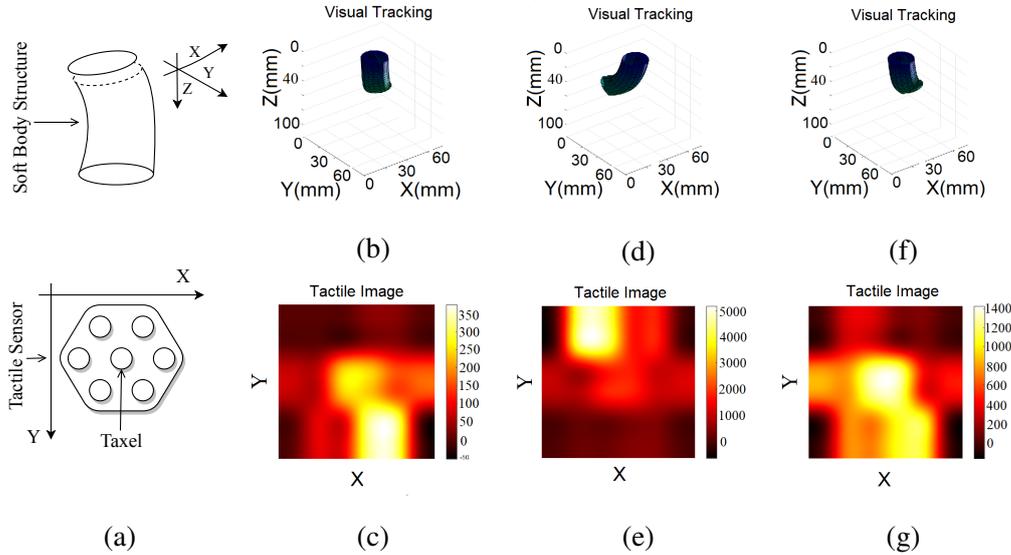


Fig. 4.4 Tactile image pressure sensing, induced by the change in posture of the soft body. (a) shows the coordinate system for each corresponding row in the figure. Figures (b), (d) and (f) show the visually tracked soft body pose in three different poses. Figures (c), (e) and (g) show the corresponding tactile images, where increasing brightness implicates higher sensed pressure.

4.2.3 Sensor Pressure Distributions and Visual Tracking

When the tendons are actuated by the servos, the finger's change in posture induces changes in the distributed pressure on the sensor's surface. To demonstrate how these vary for different poses of the soft structure, we have modelled the soft structure and pressure distributions through FEM. The soft continuum body structure has been modelled as a third order reduced polynomial with a modulus of $E=1.4$ MPa, which has been shown to provide the closest model to the true behaviour [59]. We use a Cosserat model to represent the interaction between the tendon and the soft body. The surface pressure for a given force applied to a tendon can be simulated. The simulation was performed through the MATLAB FEM toolbox. The force applied by a tendon changes the magnitude and area of the pressure distribution at that location (as shown by Fig. 4.3 (a), (b) and (c)). Similarly, by combining different tendon actuation, more complex pressure distributions can be observed (Fig. 4.3 (d)), reflecting the type of actuation both in magnitude and location. Following the hypothesis, the pressure distribution is here indicative of the posture of the soft body, determined by the specific tendon actuation.

This approach has also been shown experimentally (Fig. 4.4) where the tactile image represents the pressure distribution on the sensor's surface for different finger configurations, as retrieved and reconstructed by the camera tracking system.

The tracking system performs 3D tracking of markers placed at the head of the soft finger. Given the shape and make of the soft continuum body, the reconstruction of the body's configuration consists of a logarithmic interpolation between the base of the finger, sensed when the finger is at rest, and the tracked x-y-z positions thereafter.

4.3 Results

4.3.1 Soft-Body Configuration Learning from Visual Tracking

We perform experiments where the soft finger is actuated by random combinations of servo angles, reaching arbitrary points in space. The cameras, placed on the aluminum frame, perform visual tracking and retrieve the configuration of the 3D finger for each set of servo angles. Concurrently, tactile images from the capacitive tactile sensor array are recorded and stored for each configuration. The actuation of the finger is run autonomously for a total of 15000 random configurations. A Neural Network is used to map the tactile sensor responses directly to the finger head positions. Given the servo and tendon placement, stretching postures could not be achieved unless manually induced. As such, the network was not trained on stretch prediction.

The Neural Network is a fully connected feed-forward network with an input layer of 7 units, to read tactile image information, two hidden layers of respectively 60 and 30 units, and an output layer of 3 units, returning an xyz position in space, corresponding to the camera tracked outputs, and relative to the learned finger configuration (Fig. 4.5a). The non-linearity for all units is a \tanh , with a Glorot uniform initialization [75]. The design of the network is here secondary to the main research goal, with this specific implementation enabling testing, to identify whether it is possible to accurately determine the deformation of the finger via the capacitive tactile sensor array sensing. The number of layers and neurons was a design choice meant to confer the network enough flexibility (in terms of learning parameters) to learn 3D configurations, while keeping the architecture small enough for fast training in machines with low computing power.

The network is trained over 52 epochs, with the 15000 tactile images in input and corresponding x-y-z visually tracked positions for target outputs. 75% of the data was used for training while 15% was used for validation and 10% for testing. Fig. 4.5b shows the training error, validation and test performance of the network during learning. The network

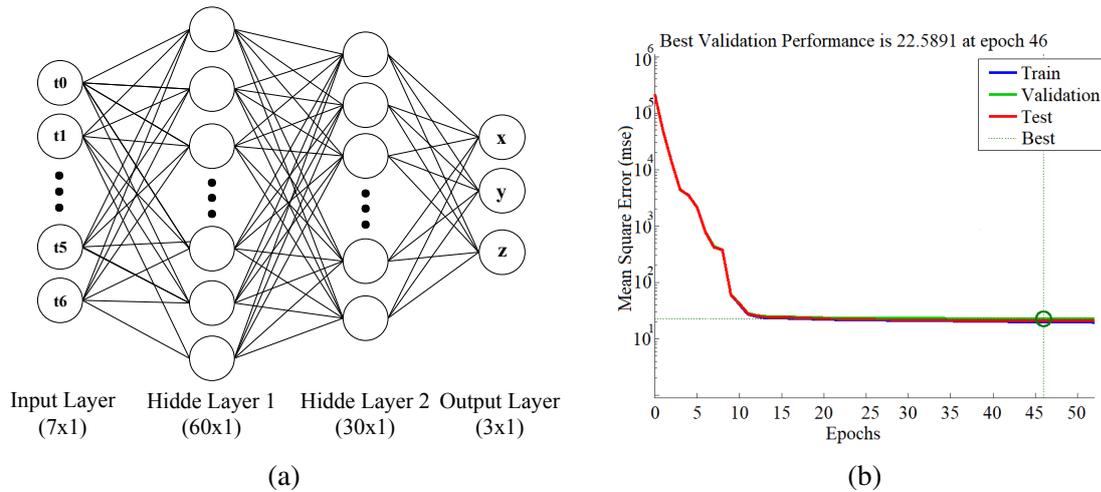


Fig. 4.5 The figure shows (a) the Neural Network for tactile pose estimation, and (b) the Neural Network training and validation curves.

trained over 52 epochs, before halting due to early stopping, and reaching a lowest validation error of $22.589px$ and test error of $22.732px$. Given cameras' placement with respect to the finger position, a pixel corresponds to $\approx 0.13mm$

4.3.2 Tactile Proprioceptive analysis

For testing purposes we retrieve 2000 previously unseen finger configurations, with corresponding tactile images and visually tracked head positions. Fig. 4.7 shows the compared finger deformation reconstruction of four different finger test configurations, based on the camera visual tracking and the proprioceptive tactile sensor response after learning.

The close correspondence between the true position and that estimated by the tactile sensors clearly demonstrates how the embedded capacitive tactile sensor is capable of matching the performance of the external camera tracking. In Fig. 4.6 we compare the error of the tactile sensor, over all axis within its work-space, to the ground truth retrieved by visual tracking. In a 'reachable' work-space of $\approx 40X40mm$ in x-y space, and $\approx 30mm$ in z (or height), from the figure it is possible to see how, on average, the tactile prediction after learning is within $1mm$ from its ground truth counterpart, retrieved from visual tracking. The experiments show how the tactile sensor array is capable of capturing the information relative to the deformation of the soft continuum fingers, with levels of accuracy near $1mm$ on average in all axis of deformation.

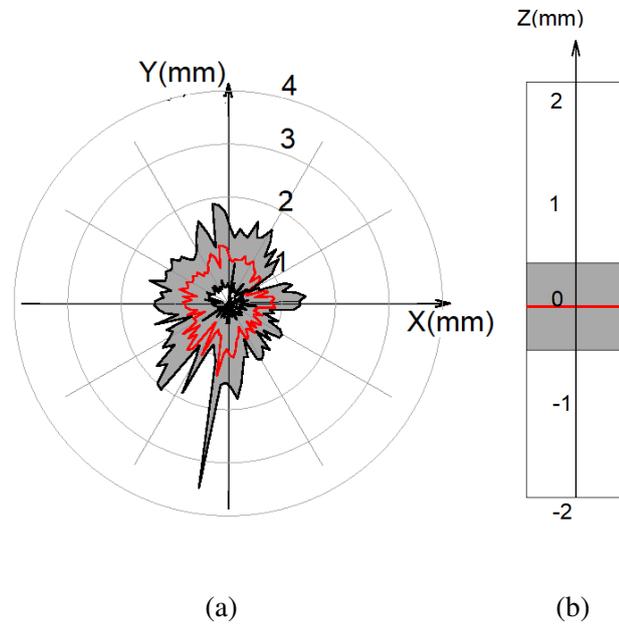


Fig. 4.6 The error (red line) and standard deviation (gray area) of the finger (a) in x-y coordinates, and (b) in z coordinates, as compared to the corresponding visually tracked positions.

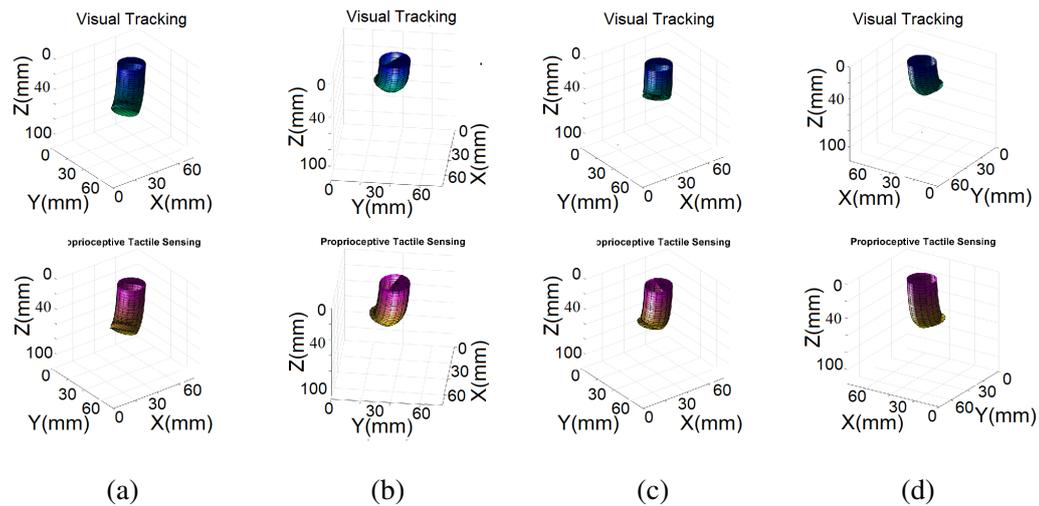
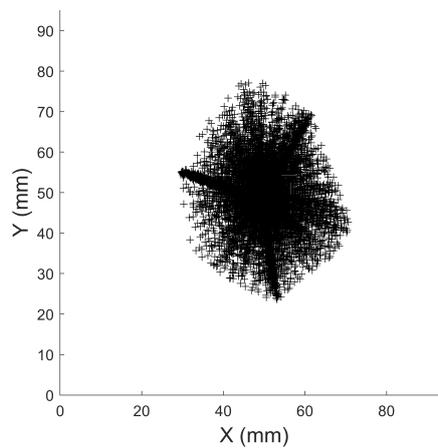


Fig. 4.7 Shape reconstruction based on the proprioceptive tactile sensing as compared to visual tracking.



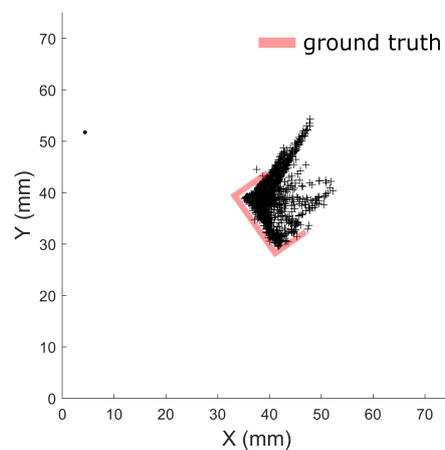
(a)



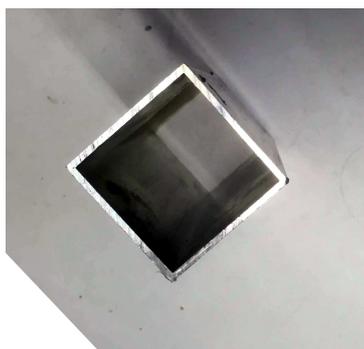
(b)



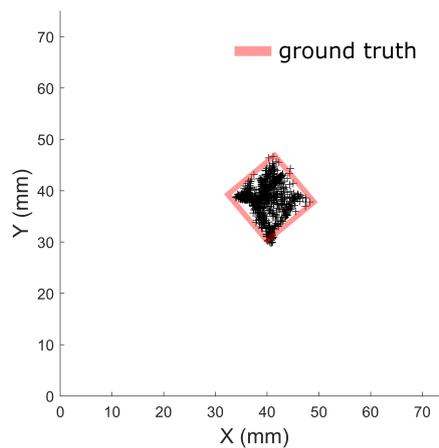
(c)



(d)



(e)



(f)

Fig. 4.8 Work-space reconstruction after random exploration. (a) shows the work-space explored when the finger was free to move over its reachable work-space. Figures (d) and (f) show respectively the reached positions when the soft finger was placed inside the cuboidal object shown in (c) and (e).

4.3.3 Tactile Proprioceptive Work-space Exploration

As previously mentioned, it is often the case that the conditions surrounding a deployed robotics system do not allow the installation of exteroceptive sensing mechanisms, like cameras. After learning, we halt the camera tracking system and undertake experiments where the sensorized finger is randomly actuated within unknown work-spaces. Initially, The finger is actuated over 2000 random servo angles, whilst free to move within its own environment. After, the soft continuum body is placed within a semi-closed and closed off environment, where vision based reconstruction methods are not possible. In the first instance, the finger is placed in a cuboid with three missing faces (Fig. 4.8c), and is further actuated over 2000 random configurations. In the second instance, the experiment is repeated with the finger placed within a space in the shape of a cuboid with two missing faces (Fig. 4.8e). Figures 4.8b, 4.8d and 4.8f show the x-y positions of the soft finger, as predicted by the Neural Network using only the sequential tactile images recorded. The finger explored spaces in Fig. 4.8f and Fig 4.8d are significantly smaller than the full work-space explored explored in Fig. 4.8b. Remarkably, the retrieved x-y positions accurately match the work-space explored by the sensorized soft structure in shape, as can be seen by comparing Fig 4.8c with Fig. 4.8d and Fig. 4.8e with Fig. 4.8f. The figure illustrates how through autonomous exploration and proprioception, it is possible to accurately retrieve the state of the work-space surrounding the soft finger.

4.4 Discussion and Conclusions

Retrieving the spatial configuration of soft continuum materials is currently a challenge. Over the past few decades, various methods have been devised, however, these methods are only capable of discriminating between preset states, or can sense deformations along one axis in space. We have devised a novel method to retrieve deformation information based on capacitive tactile sensing technology, where we embedded a capacitive tactile sensor array to read pressures at the base of a soft continuum cylindrical body, or finger. Experiments were performed, where the soft continuum body was deformed by actuating three attached tendons concurrently. The resulting material deformation allowed the finger's head to reach an arbitrary point in 3D space within its work-space. A camera tracking system was used to track the head of the finger synchronously to retrieving pressure patterns at its base, through the tactile sensor array.

By using the camera tracking system as a supervisor we have shown how it is possible to autonomously learn the body configuration and the work-space of the soft structure, through the random actuation of the finger. The capacitive tactile sensing technology is used to

achieve a proprioceptive kinaesthetic understanding of the soft structure, allowing a Neural Network to guess the head position of the cylinder within 1mm from its visually tracked position. Moreover, the system was shown to be reliable against external disturbances and interactions, as demonstrated during the random exploration of closed spaces, where the actuated body would unavoidable collide with the surface of the surrounding walls. We have shown how it is possible for the finger to perform autonomous explorations of work-spaces where cameras are impractical to use. Finally, the method of embedding the tactile sensor into the continuum body allows for the inherent soft properties of the continuum material to be maintained (e.g. stretch and bend). This work has the potential to enable and assist with many applications, such as medical exploration, e.g. provide postural feedback, or the resolution of exploratory and manipulation tasks.

The limitations of this approach lie with the range of physical modalities that are possible to recognize on the tactile sensor surface. The forces generated on the tip of the robot (via appropriate tendon pulls), in fact, generate a reaction force at the base of the robot, where the sensor is placed. The girth and length of the soft continuum body is such that these reacting forces almost always generate a unique bending configuration for the robot, in the form of compression and (single) bends. As such, it is possible to map the pose of the soft structure with one sensor, and a simple neural network. If the robot was to be such that these reacting forces could generate more than one unique configuration (e.g. allow for double bending, or buckling in different directions), then the final pose of the robot would not be unique to a pressure pattern, but would also depend on the “history” of bends of the soft structure. This, in turn, would not allow a simple neural network to uniquely identify the pose of the soft structure through the tactile information at the base alone. Multiple ways can be devised to obviate this issue. One, for example, would be the placement of sensors in multiple sections of the soft structure, such that for any two adjacent sensors, the assumption of unique bending modalities is held. Another, would be to allow for a model of the soft structure to account for the missing information, thus a hybrid model-based and learning approach could allow for accurate sensing of more complex structures.

In the context of conferring robots the ability to extend their bodily properties in order to cope with the uncertainty and unknowns of the real world, the ability to be able to perceive your own body is fundamental. Moreover, the ability to extend, and adapt, that perception over time becomes a necessity, when prior knowledge about the task to solve is not available a priori. This work marks a clear step towards conferring robots the ability to morph, and achieve a proprioceptive understanding of their own capabilities in a mostly autonomous manner.

Finally, this Chapter has shown the ability of model-free approaches to capture the complex interactions between a soft body and its surroundings. These have allowed the robot to autonomously achieve accurate proprioception of its continuum soft body, and to resolve a map reconstruction of the area surrounding it, through collisions. As detailed in Section 4.1, prior work has also applied model-based proprioception approaches to capture these interactions. One of the biggest limitations of model based approaches in the context of this chapter, is our limited ability to compute mathematically the complex soft physical interactions unfolding in the soft continuum robot when tendons are pulled, or external entities are collided with. Approximate models can give reasonable results, but slight changes in the body, or the environment, usually induce the model to move far away from the physical reality of the robot, and consequently, proprioceptive approximations to worsen. Although model free approaches can solve this problem by continuously updating the model through new sensor data, it is also true that they usually do not scale in complexity. If several sections of the sensorized soft robots presented in this chapter were to be connected together, such that complex twists or double bends were possible, thousands of additional data points would need to be generated by “random bubbling” before learning a proprioceptive mapping appropriately. One of the most important direction for future research should aim at combining model-free and model based approaches into one, such that scaling complexity is possible, but the models are not fixed a priori by human knowledge, rather they change with the changing of the physical properties of the robot. This topic will be expanded upon in Chapter 11.

Chapter 5

The Importance of Morphology: Achieving Robotically Peeled Lettuce

In this chapter we begin to contemplate the topic of Morphology. This chapter reports a case study of the usefulness of morphology and design to solve a real world robot task. The task is that of removing the leafs from a lettuce, a soft and delicate object.

Reference Publication

This chapter was adapted from a published Journal article titled “**Achieving Robotically Peeled Lettuce**” [98], in the journal IEEE Robotics and Automation Letters (RA-L) in 2018. The article was written in collaboration with Dr Perla Maiolino, Dr Josie Hughes and Dr Fumiya Iida and investigates the ability of robots to perform ‘leaf removal’ in lettuce produce for automation purposes. This paper highlights the importance of morphology in tactile tasks. Dr Fumiya Iida and Dr Perla Maiolino helped with the conceptualization of the topics, design of the experiments and the writing of the article. Dr Josie Hughes contributed by devising and carrying out the physical lettuce leaf removal experiments, and the writing of the manuscript. As co-first author in the article, my contribution includes the development of the vision based technologies in the project, design of the experiments as well as data analysis, and article writing.

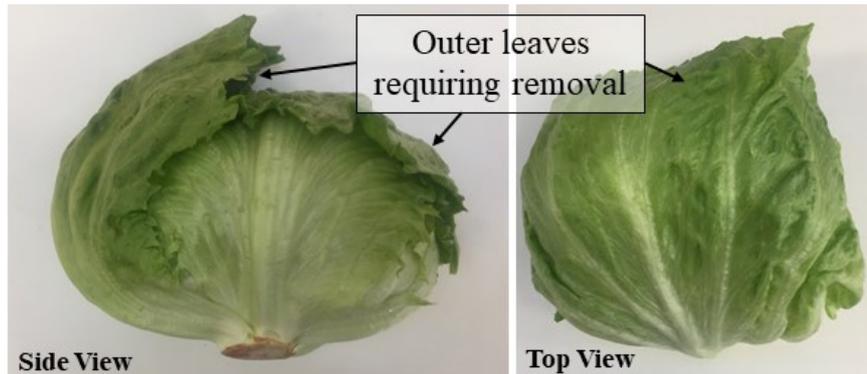


Fig. 5.1 Iceberg lettuce before processing showing the two outer leaves which require removal.

5.1 Introduction

The automation of agricultural systems presents many challenges due to the extreme variability in the environment e.g. varying light, terrain and weather conditions, unlike constrained industrial or lab environments. As such, the successful application of robotics systems to agriculture remains limited. However there is a growing need to develop automated robotic solutions due to the increasing demand for food, sub-optimal climate conditions and decreasing labor availability [71]. The majority of existing robotic agriculture research has focused on developing harvesting systems [62, 85, 128, 176]. However, often there is a crucial, labor-intensive post-processing step required after harvesting to produce supermarket ready produce [275, 278].

Previous work has seen the development of vision, learning and manipulation solutions for lettuce harvesting [32]. However, supermarkets require the complete removal of one or two loose outer-leaves from the stem after the harvesting process, with no bruising, damage or ‘browning’ to the lettuce [30, 70] (Fig.5.1). This is a challenging task requiring dexterous manipulation and robust vision as the leaves are fragile, tear easily and have limited identifiable features.

Manipulation of soft produce is challenging; dexterity is required and the produce must not be damaged. Although Soft Robotics has significant potential for such applications [211], current solutions provide limited dexterity and require complex control [19]. The vision problem is also challenging, especially when using only 2D stereo cameras. The variability in produce and the lack of rotational varying features makes lettuce a difficult object to interact with, and means existing approaches cannot be used [62, 128, 176, 248].

The problem addressed in this work is the removal of outer leaves, typically two, from a harvested lettuce. The produce, placed in an unknown pose on a flat environment, should

not be damaged by the peeling process. The proposed system uses a single vacuum suction point to grab a leaf and remove it from the main body of the lettuce by tearing. Suction is a method which has been applied successfully in other agriculture applications and reduces the complexity of the required control [9, 143]. This process has been optimized to maximize successful tearing.

To address this problem, firstly computer vision must be used to locate and determine the pose of the lettuce. A novel machine vision pipeline has been created for orientation detection which uses only a 2D web camera. The pipeline is robust, and unlike existing approaches it does not require depth information from RGB-D, TOF cameras or stereo vision [62, 128, 176, 248]. The lettuce must then be manipulated into a known state with the outer leaves exposed. The nozzle diameter, material properties, location and trajectory of motion must then be optimized for successful leaf removal.

In this chapter the specific methods and implementation are given in Section II, with corresponding results in Section III. Complete testing and demonstration of the leaf removal system is shown in Section IV, concluding with a discussion and review of the work presented.

5.2 Methods and Implementation

5.2.1 Vision and Detection

For a lettuce placed in a unknown location within a work space, machine vision must be used to determine the location and orientation of the lettuce. The former is possible through simple lettuce image segmentation. We propose a novel method for retrieving pose by detecting the lettuce stem. This method is necessary as the lack of depth information makes previously explored solutions unsuitable [116]. Moreover, existing approaches for 2D images use color thresholding techniques which are not applicable as the stem hue and saturation can vary significantly depending on freshness and growing conditions of the lettuce [79, 85].

Lettuce Segmentation and Position Estimation

The vision system uses a single *Carl Zeiss Tezzar HD 1080p* camera placed above the work space, where it is assumed a single lettuce is within the field of vision. The background is assumed to be distinguishable in color to the lettuce, thus, a combination of color-based thresholding and binary cleaning is used for robustness.

The approach adopted uses the HSV (Hue, Saturation and Value) format of the retrieved frame. Two noisy, binary masks are computed by thresholding the Hue channel for pixels $30 < p^h < 100$ and the Saturation channel, for pixels $p^s > 50$. These are combined with an



Fig. 5.2 The figure shows two sample images, (a) example masked frame, after segmentation. (b) Example lettuce under the A-channel of the LAB color space, showing homogeneous stem pixels.

element-wise AND operation. Hue thresholding is empirically set to retrieve pixels with a color pigmentation within the typical lettuce range. The Saturation channel thresholding removes the false positives due to no luminosity and discarding the brightness channel provides robustness to varying lighting conditions.

In the second step, we clean the mask by applying a morphological dilation followed by an erosion, with a circular disk small enough to connect only nearby pixels ($25px$ in radius). Finally, we find the contour of the largest connected binary area in the mask, set the enclosed pixels to 1, and extract the lettuce from the original frame by masking it with the computed mask (Fig. 5.2a). The center of the connected binary area in the final mask is assumed to be the center position of the lettuce.

Stem Detection

Stem detection is the first step towards retrieving a reliable orientation estimation of a lettuce. Due to radical variations in hue and saturation depending on the freshness of the produce, stem detection can not be performed using simple thresholding. A sliding window is used to perform a sequential search of the area within the detected lettuce. In the search we attempt to identify a set of features unique to the stem. The image is converted into the LAB color space, such that value of the pixels within the stem area will be homogeneous across different hues, saturation and lighting conditions. The A-channel spans from a minimum where pixels partake a green coloring, to a maximum where they show a red hue. As the stem never has the same coloring as the green outer leaves, the A-channel is suitable for the stem detection. Fig. 5.2b shows the A-channel of a detected masked lettuce picture after the LAB conversion.

The stem detection process starts by finding the expected size of the stem within the lettuce. Let \mathbf{M} be the previously computed binary mask. We first find the area of the lettuce a_l as the sum of the pixels in the mask. The radius of the circle whose area is the same as the detected lettuce is $r_l = \sqrt{\frac{a_l}{\pi}}$. Experimentally, the stem diameter has been determined to be $\approx \frac{r_l}{5}$, with an average error of $\pm \frac{r_l}{12}$.

The stem look-up is performed by searching pixels, sliding left to right and top to bottom within the area where the lettuce is contained. The stride length, ι , is tuned to a value small enough to fall within the expected area of the stem multiple times before the end of the search, i.e.: $\iota = \frac{r_s^e}{4}$, where r_s^e is the expected radius of the stem in the lettuce.

Let \mathbf{A} be the A-channel of the image in the LAB color space. For each set of indexes (i, j) , \mathbf{A} is thresholded around the value of $\mathbf{A}_{i,j}$, thus creating an initial masked window \mathbf{M}^f , where we compute the element in position (k, m) as:

$$\mathbf{M}_{k,m}^f = \begin{cases} 1 & \text{if } \mathbf{A}_{k,m} > \mathbf{A}_{i,j} - \varepsilon \text{ and } \mathbf{A}_{k,m} < \mathbf{A}_{i,j} + \varepsilon \\ 0 & \text{otherwise} \end{cases} \quad (5.1)$$

The ε determines the strength of the thresholding around the currently inspected pixel (Fig. 5.3b), where $0 < \varepsilon < 256$. We empirically set $\varepsilon = 7$, to ensure homogeneity amongst the thresholded pixels.

The binary image is de-noised by performing an *opening* morphological operation with a disk of $\frac{r_l}{8}$ followed by a closing operation with a disk of *disk radius* = $\frac{r_l}{3}$. The parameters were chosen to cope with the expected radius of a stem. Finally, we retrieve the largest binary area contour and set the enclosed pixels to 1 (Fig. 5.3c). We will refer to the clean mask as matrix \mathbf{M}^f . The binary area in \mathbf{M}^f is an estimate of the stem location in the image. To validate the stem we have devised four unit test: vicinity, comparative area, solidity and elongation.

For vicinity, we test if the (x, y) coordinate of the binary area found after the cleaning operation is within a distance equal to the expected radius of the stem. In the comparative area test, we retrieve the expected area of the stem as $a_s^e = r_s^{e2} \pi$, and validate the binary area a_s if $|a_s - a_s^e| < a_s^e$. In the third test, solidity is a measure of the convexity of the found area. Here we compute the solidity *sol*, as:

$$sol = \frac{a_s}{\Phi(\check{C}^s)} \text{ and } sol \in \{x \in \mathbb{R}_+ | x < 1\} \quad (5.2)$$

where \check{C}^s is the contour of the binary area in \mathbf{M}^f and $\Phi(\check{C}^s)$ is the area of the convex hull surrounding the contour \check{C}^s . In general we expect a stem to be approximately round (i.e. $sol = 1$), so we reject a candidate if $sol < 0.7$.

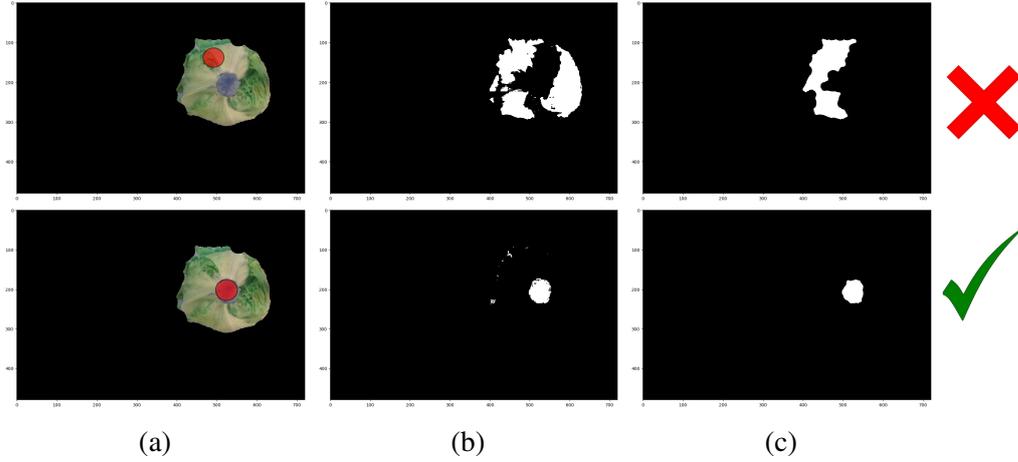


Fig. 5.3 The stem detection process: (a) search amongst selected pixels in the lettuce, (b) thresholding on the value of the looked up pixel, (c) cleaning the image. The stem candidate is accepted only after validation.

In the fourth and last test we require the stem to be circular, by assuming convexity, we can test this by $\frac{\Psi(\check{C}^s)}{\psi(\check{C}^s)} < \zeta$ where $\Psi(\check{C}^s)$ is the length of the major axis in \check{C}^s , $\psi(\check{C}^s)$ the length of its minor axis, and ζ calibrates the accepted error elongation margin (here $\zeta = 1.2$, to ensure approximately round binary areas).

Once all tests are passed, the (x, y) position in the frame will be returned as the location of the lettuce stem (Fig. 5.3). In this process, almost all tunable parameters are based on the previously found radius of the lettuce. Therefore, assuming that produce can be reliably detected this is a generic process for stem detection.

Lettuce Orientation Estimation

Given the shape, a lettuce can be found in one of two poses: stem facing downward or facing up. In the case where the stem cannot be found, an action can be taken to flip the lettuce over. If the stem is found, it is possible to find the pose as a 3D vector starting from the stem and pointing outward, towards the front of the lettuce. This vector is given as:

$$\begin{bmatrix} p_l[0] - p_s[0] \\ p_l[1] - p_s[1] \\ h_l^e \end{bmatrix} \quad (5.3)$$

where p_l is the estimated center of the lettuce, p_s the position of the stem and h_l^e the expected lettuce height. As the proposed solution is based on single 2D images, the height of the lettuce is expected to be the average computed height in a data-set of 10 iceberg lettuces ($\mu_l = 111.27mm \pm 3.51mm$). The radius can not be used due to the elliptical shape of lettuce.

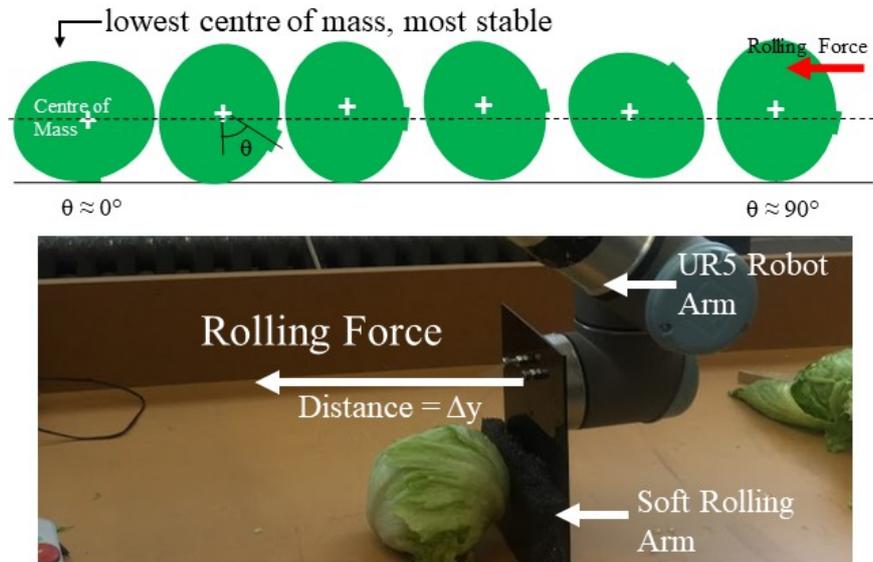


Fig. 5.4 Diagrammatic representation of the lettuce (top) and method for rolling the lettuce using a soft pad attached to a UR5 robot arm (bottom)

The leaf removal process is tolerant to changes in height, and thus height errors were found to have a negligible impact on the peeling process.

5.2.2 Lettuce Orientation

By applying a horizontal force and rolling the lettuce, a lettuce can be oriented with the outer-leaf on top and with minimal risk of damage. Modelling the lettuce as an ellipse, when a force is applied the lettuce will roll and then stop at a stable point where the centre of mass is at the lowest point, i.e. the top and bottom side of the lettuce. This corresponds to angles θ , between the stem and the normal, of $\theta \approx 0$ and $\theta \approx 180$ (Figure 5.4). The optimum distance to apply this force over to roll the lettuce such that it is in the top position ($\theta \approx 0$) must be found.

5.2.3 Leaf Removal

Nozzle Design

3D printed circular nozzles of varying inner diameters have been used, with a 3mm lip around the nozzle to allow for formation of a seal (Fig. 5.7a). The normal holding force is given by $F_h = \Delta p A_e$ where Δp and A_e denote the pressure difference and the effective contact area respectively. Rough surfaces of the lettuce can affect the effective area in contact, reducing the holding force. Therefore a flat non-contoured area of the leaf should be used. The suction

nozzle surface material affects the leakage flow, with compliant softer materials helping to achieve a seal. The nozzle diameter and materials must be optimized to hold and tear a single leaf opposed to lifting multiple leaves or the entire lettuce.

Leaf Removal point

Picking near the stem reduces the risk of tearing the fragile leaves and achieving partial leaf removal. However, the leaves are more textured (up to $4mm$ variation in height) nearer the stem making it harder to achieve suction. Towards the edges of the leaf the thickness drops to $0.15mm$ with the leaves more fragile, however the radius of curvature is much lower providing a flatter surface minimizing leakage flow.

Leaf Pulling Trajectory

The trajectory in which the leaf is moved after suction is applied affects success. Moving towards the base encourages the snapping of the leaf, as opposed to tearing. The trajectory should be optimized to maximize this snapping force.

5.3 Results and Discussion

5.3.1 Vision and Detection

The lettuce and stem detection algorithm were tested on a set of 180 pictures taken with the camera facing directly downwards above the work space at heights between $70cm$ to $100cm$. 10 different iceberg lettuces were used in various poses with varying light direction, light intensity and background objects (Fig. 5.6). In addition, 30 frames were taken after storing the produce for 3 days, inducing changes in stem color.

The algorithm had 100% detection accuracy, estimating the center with an average accuracy of $20.21px \pm 0.48px$ from the true lettuce center. The stem detection algorithm found 64 of the 79 visible stems in the data-set, with an average distance from the true stem center of $5.76px \pm 0.24px$, reaching a detection accuracy of 81.01%. Given the camera height, on average we detect the lettuce true center within $7.78mm$ and the stem within $1.73mm$, allowing us to achieve the optimum lettuce removal point.

5.3.2 Lettuce Orientation

The lettuce was placed in a pose corresponding to a randomly generated vector with a horizontal force of $1N$ applied at $0.1m^s$ using a UR5 robot arm with a soft pad, rolling the

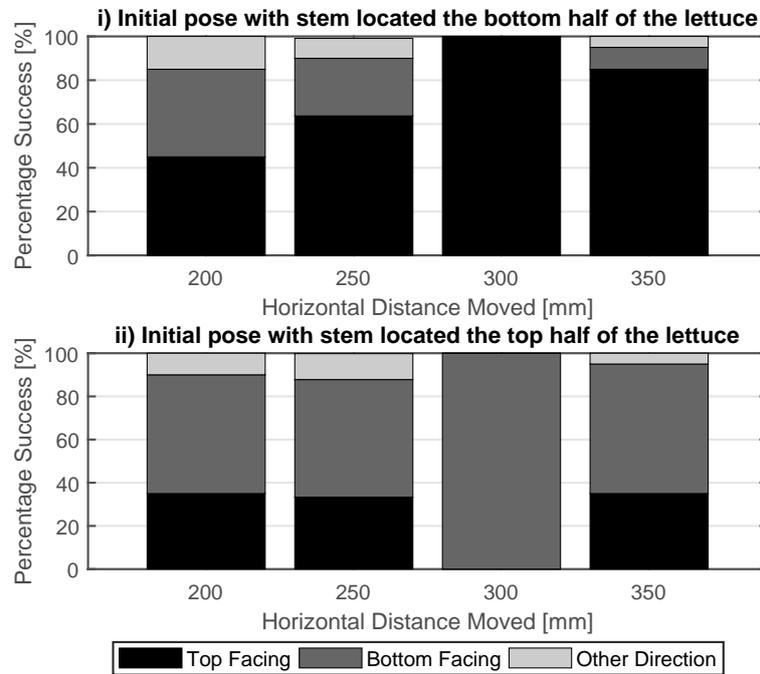


Fig. 5.5 Success rates when rolling a lettuce placed randomly, with the stem in the top half (top graph) or bottom half (bottom graph). 50 experiments performed using 10 lettuce.

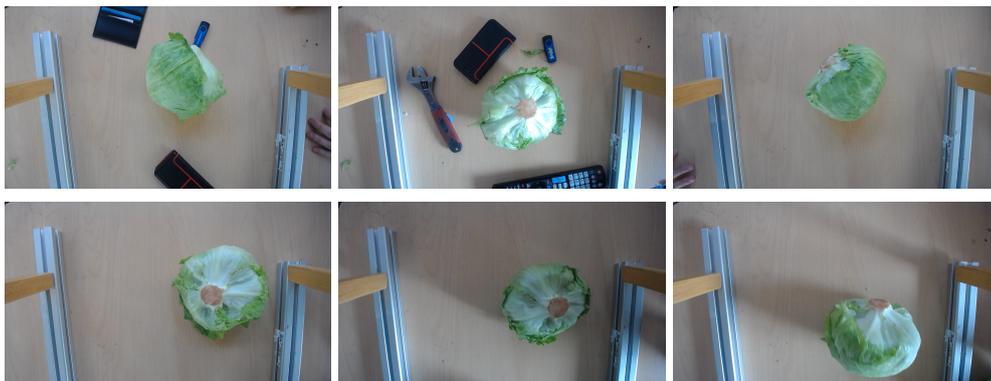


Fig. 5.6 Six example images in the data-set used for testing.



Fig. 5.7 The figure shows (a) the Nozzles tested, and (b) the possible outcomes when lifting the lettuce: partial removal, full leaf removal and the entire lettuce being grasped demonstrated with a number of nozzles and without nozzle.

lettuce (Fig. 5.4). The end position of the lettuce was recorded, with the lettuce considered to be top down (such that the leaves are exposed for gripping), bottom down (stem is exposed) or in another direction. The results for 5 lettuce each rolled 10 times (Fig. 5.5) show that the optimum distance for horizontal rolling is 300mm , where for both tests the lettuce ends in one of two states, top or bottom facing. Therefore, given an initial random orientation, the lettuce should be rolled by 300mm . After this, stem detection should then be used to determine if the stem is exposed (bottom position) and if not, the lettuce should be rolled a further 300mm in the opposite direction (to keep the lettuce within the work-space). If necessary, this process can be repeated until the top surface is exposed.

5.3.3 Nozzle Design

Nozzles of varying diameter and material (plastic or silicone) have been tested with the suction system and nozzle mounted on the end of the UR5 robot arm. A lettuce is placed in the correct ‘top down’ position, with the suction nozzle then lowered onto the middle of a lettuce leaf. The vacuum is then applied and the arm lifted. There are four possible outcomes: failure to grasp, lifting the entire lettuce, partial leaf removal and successful leaf removal (Fig. 5.7.) The results for the different nozzles are shown in Table 5.1. The maximum lift force was measured to be 6.8N , significantly exceeding the typical weight of a lettuce ($\approx 4\text{N}$).

The smaller nozzle (7mm diameter), had insufficient strength to grasp any leaves, only achieving a partial tear. The largest nozzle (37mm) displayed reasonable ability to remove the leaves, both partially and fully, however the increased contact area can result in lifting of the entire lettuce. The addition of a silicone outer ring to the nozzle (27mm and 37mm) increased the effective area, enabling a higher area of leaf tearing and successful leaf removal. The silicone’s elastic properties, in fact, can better fit the non-uniform surface of the lettuce’s outer surface, achieving an appropriate “seal” at the surface level via the soft adaptation of

Table 5.1 Results from grasping lettuce leaves in the middle of the leaf and lifting with a nozzle of given diameter and material, each nozzle tested on 10 leaves.

Nozzle Diameter	Nozzle Material	No Grip	Lift Lettuce (%)	Partial Leaf Tear (%)	Successful Leaf Removal (%)
7mm	Plastic	71	0	29	0
17mm	Plastic	9	0	55	36
27mm	Plastic	0	8	82	10
37mm	Plastic	0	40	38	22
27mm	Silicone	0	6	53	41
37mm	Silicone	0	10	48	42

the end-effector itself. The softness, in turn, would also provide a more gentle “dragging” of the tearing leaf across the lettuce’s body. This is an important morphological process, enabling appropriate leaf manipulation by the robot. The 27mm silicone nozzle showed the best performance.

5.3.4 Leaf Removal Radius

Using the 27mm nozzle from the previous experiment, lettuce leaves were gripped at varying distances along the middle of the leaf and then raised vertically, with the area percentage of leaf removed measured to the nearest 10% (Fig. 5.8).

Due to the extreme variability in lettuces, there is significant variation in the percentage of leaf removed, as shown by the magnitude of the error bars. At the outer edges, the leaf is fragile and tears easily, leading to a limited leaf removal. Near the stem, there is limited leaf removal due to the highly textured leaf. Thus, there exists an optimum picking location at approximately $0.7r$ from the stem. The vision system can detect the lettuce location with an average error of 7.78mm, assuming a typical lettuce to have a diameter of 200mm, the positioning error in achieving this point is less than 5% such that successful removal can be maximized.

5.3.5 Leaf Trajectory

To test the impact of the leaf pulling trajectory, the 27mm nozzle was tested $0.7r$ along the leaf, with the arm moved at varying angles after applying suction and ‘grasping’ the leaf. The inclination of the trajectory relative to the normal vector of the lettuce is kept constant, at 45 degrees, such that the leaf comes clear from the body of the lettuce. The percentage area

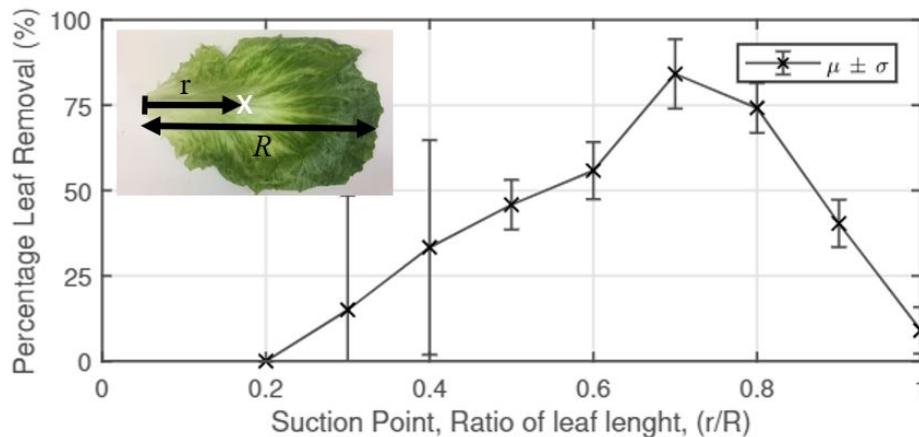


Fig. 5.8 Percentage area of leaf removed, with the leaf removal point along the stem of the leaf. The experiment was repeated for 10 lettuce each with two outer leaves which were removed, i.e. 20 repeats.

of the lettuce leaf removed is shown in Fig. 5.9 for different pulling angles. Pulling at 120 degrees offers the greatest chance of leaf removal as enables the leaf to snap opposed to tear.

5.4 Demonstration

To demonstrate the abilities of the methods and approaches presented, leaves were removed using the techniques discussed. Fig. 5.10 showcases the peeling process and videos shows the different components of the peeling process². A lettuce was placed in a random location and orientation (as discussed previously) and is then rolled to the correct location. The second process, leaves are removed from the optimal position on the lettuce using the 27mm silicon nozzle. The time for this leaf removal process (assuming the lettuce is correctly orientated) was measured to determine the average cycle time. With limited testing (10 lettuces, each with 2/3 outer leaves) there was a success rate of full leaf peel of 50% and partial leaf removal of 30% with an average time of 28.5 seconds. In most cases failure was due to lifting the lettuce followed by leaf tearing.

5.5 Conclusion

This chapter has presented mechanical and vision concepts to enable the automated removal of lettuce leaves; the proposed pipeline is the first autonomous system to perform this task. The mechanical systems have been tested and demonstrated and the required vision

²Video demonstrations can be found at: <https://goo.gl/S68Wnr>

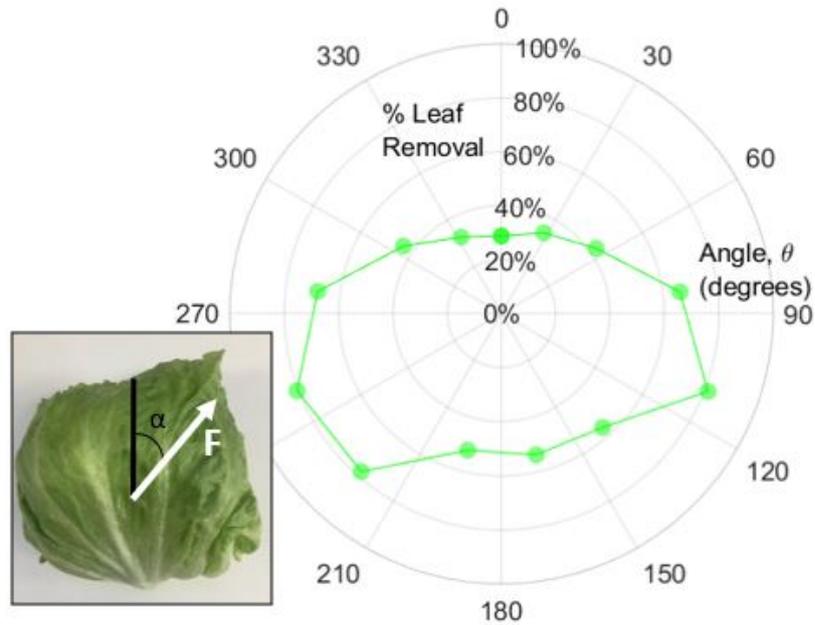


Fig. 5.9 Percentage of lettuce leaf removed with varying trajectory angle α (see inset). Each angle was tested on different outer lettuce leaves.

systems presented. Although the system has been specifically designed for the removal of lettuce leaves, there is significant wider applicability of the system and techniques developed. The model and understanding of the physical parameters affecting successful removal of lettuce leaves can be applied to many other crops such as cauliflower which is far less fragile and presents less visual variations with orientation [8]. The vision, in particular, has wide applicability in other crops where similar information would be required for the post-processing of the produce.

The lettuce and stem detection, tested on a challenging data set, have demonstrated that the proposed solutions are robust to clutter, lighting conditions, camera distance, morphological



Fig. 5.10 Time series of the leaf removal showing the time taken to complete the process.

variations of the produce and changes in its position and orientation in general. The major drawback of the approach lies in the significant reliance on the produce's assumptions. Although the assumptions provide robustness, they have the potential to induce detection failures when they are not met in the real world (e.g. approximate lettuce-to-stem relative size). In this context, the tuned parameters based on the produce assumptions would need to change to reflect different produce types.

Although all components of the system have been tested (rolling, pose estimation and peeling) they are currently distinct processes. Further work to integrate this into one single end-to-end solution is required. The current system is capable of performing the peeling process with full leaf removal 50% of the times with an average time to complete of 27 seconds. This approach was specifically designed to minimize contact with the lettuce, reducing the chances of damage. A two 'arm' approach, however, should also be investigated³.

The work in this chapter highlights how the morphology of the lettuce removal apparatus plays a fundamental role in achieving accurate and damage-free manipulation of the soft lettuce object. This is especially shown in Fig. 5.7 and Table 5.1. The influence of morphology to the soft physical interaction between the robot and the environment is one of the two main underlying themes of this thesis. One of the main limitations, however, is that the task resolution was here achieved via human intuition and design. Future work should instead focus on achieving similar results, while allowing the robot itself to change its morphology appropriately for the resolution of the task at hand. As explained in Chapter 1, to be able to reason about which morphology can structure the stimuli appropriately so to better achieve the task resolution, sensing is necessary. The next chapter will treat this topic in more depth.

³Steps towards achieving this can be found in the Further Work video which shows lettuce peeling performed with a two arm Baxter robot. The pose estimation and peeling process is combined. <https://goo.gl/S68Wnr>

Chapter 6

Soft Morphology to Structure Information

The concept of morphological conditioning, or morphological computation, was first introduced in Section 2.3.1, as well as in 2.4.1, where the role of morphology in filtering and moulding the sensory perception was explained. In the previous Chapter we have shown how morphology plays a fundamental role in manipulation. In this chapter we finally introduce Soft Morphology within *SoMComp*, and focus on the topic of tactile perception. We investigate the possibility of purposefully choosing a morphology to improve the soft interactions such that the physical stimuli arising can aid in object discrimination tasks. We focus on tactile sensing as a modality which is heavily dependent on morphology, as well as physical interactions, and therefore it is a good venue to show the concepts within this thesis.

Reference Publication

This chapter was adapted from an article titled “**Soft Morphological Processing of Tactile Stimuli for Autonomous Category Formation**” [227], published in the 1st IEEE International Conference on Soft Robotics (RoboSoft 2018). The article was written in collaboration with Dr Perla Maiolino and Dr Fumiya Iida and proposes a mathematical framework to make use of the morphology of a soft robot to aid in discrimination tasks. Dr Fumiya Iida and Dr Perla Maiolino contributed with the conceptualization of the topics, design of the experiments and the writing of the article. As first author in the article, my contribution includes conceptualization of the topics, the design and execution of the experiments, the robot control, the formulation of the mathematics for the framework, data analysis and the writing of the article.

6.1 Introduction

Sensor morphology is a fundamental aspect of tactile sensing technology. Design choices induce stimuli to be morphologically processed, changing the sensory perception of the touched objects and affecting inference at a later processing stage. The main contribution of this chapter is to propose a conceptual framework to examine whether any processing or meaningful transformation occurs in robotic tactile sensing due to its morphology and, consequently, how Morphological Computation can drive the robot's internal representation of the world. This work marks a step towards the design of sensors whose morphology can sensibly aid in the information processing of perceptual inputs for a task at hand.

For this purpose we perform tactile discrimination experiments. Haptic sensing differs from other modalities, such as vision, in virtue of its tight coupling with, and need of, physical interactions. The somatosensory system of biological organisms decodes, interprets and categorizes a wide range of tactile stimuli arising from interactions with the environment. This difficult task, if in part achieved at a neural level, is known to be initially performed at sensory receptor's level [2]. As an example, the morphology of the vibrissal system of rats is useful in extracting information relative to object texture, orientation, shape, size and location of surfaces. The system then, preprocesses information from the environment into useful stimuli to be further processed by the brain [10]. In humans, when a scene is explored by touch, the morphology of the skin (in particular of the *Meissner's Corpuscles* together with the *Dermal Papillae*) allows the encoding of edge information [33].

In the last decades, substantial efforts have been made in enhancing the perceptions capabilities of robots by providing them with a sense of touch [45]. Despite the large number of tactile sensors developed, the proposed solutions have been often presented at a prototypical level, where the designs needed be specifically tailored to individual robots and applications. In this context, design principles would be focused on finding trade-offs between aspects such as transduction principles, sensor performances and ease of integration, but only a limited number of research work, mainly in the Soft Robotics community, have focused on the development of tactile sensors with functional morphology [41, 43, 93]. A structured research review about the use of sensor morphology in robotic systems can be found in [106]. Despite the efforts, the role of sensor morphology in encoding and categorizing touch stimuli remains a significant challenge. Moreover, the interpretation of the sensor signals to discriminate between a set of stimuli or to perform object recognition has relied mainly on supervised machine learning techniques [72, 113, 252], burdening solutions with the need of large amount of labelled data.

The chapter is organized as follows: In Section 6.2 we describe the proposed unsupervised method for clustering using the soft filters, in Section 6.4 we describe in detail the tactile

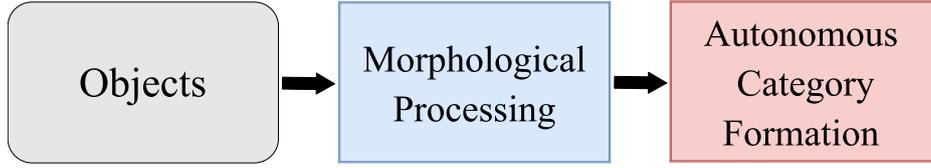


Fig. 6.1 Conceptual map for the Morphological Computation of Sensory Receptors.

sensor technology as well as the experimental set-up used for performing the experiments. In section 6.5 the experimental results are presented. Finally, section 6.6 gives a discussion of the results followed by a conclusion in Section 6.7.

6.2 Autonomous Category Formation

We propose an unsupervised process to automatically cluster a set of objects in two categories. After acquiring tactile images for each object in a set, the autonomous category formation process is mainly divided in two pipelined steps: Principal Component Analysis projection (*PCA*) [257] and K-Means Clustering (*KMC*) [144]. We use the proposed process to observe the influence soft filters with variable thickness have on the categories.

We start the process with tactile sensor readings for each object we wish to cluster. For a set of N different objects, let \mathbf{X} be a $(N \times D)$ matrix where each unique tactile *image* for an object is a D dimensional row in the matrix, where typically $D \gg 2$. We define a tactile image as a one-off tactile sensor reading, where each element in the vector is proportional to the deformation of a tactile element in a predetermined location on the sensor (Fig. 6.3b). As the tactile sensor technology does not affect the processing stages, we leave its description to Section 6.4. We begin by finding the average tactile image by

$$\vec{\mu} = \frac{1}{N} \sum_{i=1}^N \vec{x}_i \quad (6.1)$$

where \vec{x}_i is a row vector in \mathbf{X} . We proceed by computing the scatter matrix of \mathbf{X} as

$$\mathbf{S} = \sum_{i=1}^N (\vec{x}_i - \vec{\mu})(\vec{x}_i - \vec{\mu})^T \quad (6.2)$$

We use Single Value Decomposition to factorize \mathbf{S} into

$$\mathbf{S} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^{-1} \quad (6.3)$$

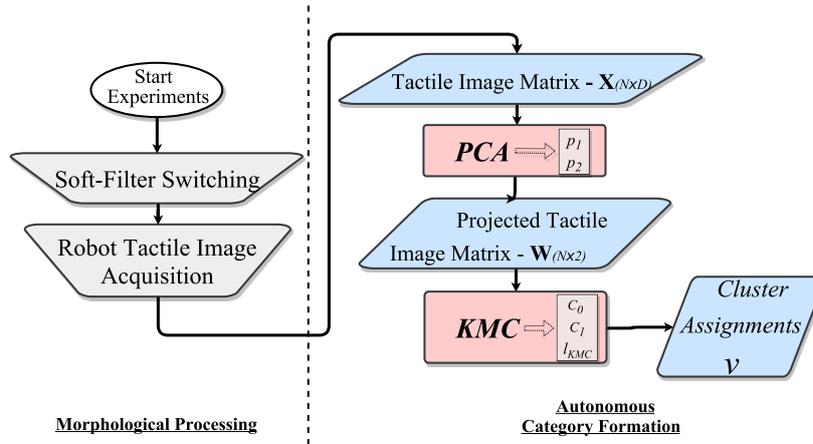


Fig. 6.2 Autonomous category formation steps.

where \mathbf{Q} is matrix such that each column q_j corresponds to an eigenvector of \mathbf{S} , and each element λ_{jj} in the diagonal matrix Λ is its corresponding eigenvalue.

We list the eigenvectors in ascending order of eigenvalue and select the first two in the list. Let \vec{p}_1 and \vec{p}_2 be the selected eigenvectors obtained from PCA . We form a $(D \times 2)$ projection matrix

$$\mathbf{P} = \begin{bmatrix} \vec{p}_1^T & \vec{p}_2^T \end{bmatrix} \quad (6.4)$$

where \vec{p}_1^T and \vec{p}_2^T are column vectors in \mathbf{P} . Finally, we project the D -dimensional row vectors in \mathbf{X} onto a 2-dimensional subspace by:

$$\mathbf{W} = \mathbf{X} \cdot \mathbf{P} \quad (6.5)$$

where \mathbf{W} is a $(N \times 2)$ matrix and each row in it is a 2-dimensional *encoding* of a tactile image. The choice of a 2-dimensional subspace was made to be able to perform clustering robustly. The choice of experimental procedure is such that each tactile interaction between the robot and an object would subject the tactile evidence to noise. Reducing the number of dimensions is beneficial in maintaining only the relevant information for discrimination. In this context, one dimension might be too low to capture the separation of classes across the 7 different tasks (where each object might cluster in the same areas), while dimensions higher than 3 might not induce robust clusters (due to the fluctuations of the tactile sensor response over time across experiments). Moreover, dimensionalities higher than 3 present difficulties for visualization and understanding. We proceed by using KMC ($k=2$ and random centroid

initialization) to split the tactile images in \mathbf{W} into two clusters, thus:

$$\vec{v} = KMC_{k=2}(\mathbf{W}) \quad (6.6)$$

where \vec{v} is an N -dimensional array, $\forall i \in \{1, 2, \dots, N\}$. $\vec{v}_i \in \{0, 1\}$, and $\forall i \exists j. i \neq j \wedge v_i \neq v_j$ (no one cluster can contain all objects). In general $\vec{v}_i = 0$ *iff* the i^{th} tactile image belongs to cluster 0 and $\vec{v}_i = 1$ *iff* the i^{th} tactile image belongs to cluster 1 (Fig. 6.2). The \vec{v} vector then contains the cluster membership of each object in the initial set. To avoid cluster anomalies due to the random centroid initializations we run the K-Means Clustering algorithm three times and discard the clustering attempt if, after convergence, any of the three cluster guesses vectors differs from any other.

As it becomes clearer later, the cluster assignments for each object are largely dependent on the soft filter employed. The change in cluster assignment is the main object of analysis in the following sections.

6.3 Tactile Sensor Technology and Data Acquisition

The tactile sensor technology utilized in the experiments for this chapter has been described in 3.2.3. In the prototype used for the experiments within this chapter the sensor is composed of 6 modules. Each module hosts 10 taxels, as well as the Capacitance to Digital Converter (CDC) chip (namely, the AD7147 from Analog Devices) for converting capacitance values to digital (Fig. 6.3a). A sensor reading from the tactile sensor described is produced at 20Hz, and corresponds to a 60-dimensional array (we exclude the central taxel in Fig. 6.3b), where each element contains the capacitance variation value of the corresponding taxel.

6.4 Methods and Experimental Set-Up

We investigate the influence soft filters with varying thickness have on tactile information encoding. We build three filters using Ecoflex 00-20² from Smooth-on, each respectively 3mm, 6mm and 10mm thick. The material was selected for its mechanical properties, in particular a Shore Hardness of 00-22. We 3D-print a custom-made end-effector with a circular flat surface (*diameter* = 80mm) onto which the soft filters can later be placed and we integrate the referenced capacitive tactile sensor onto its surface to retrieve tactile images of the objects during the experiments (the above set up is described in Fig. 6.5).

²<https://www.smooth-on.com/products/ecoflex-00-20/>

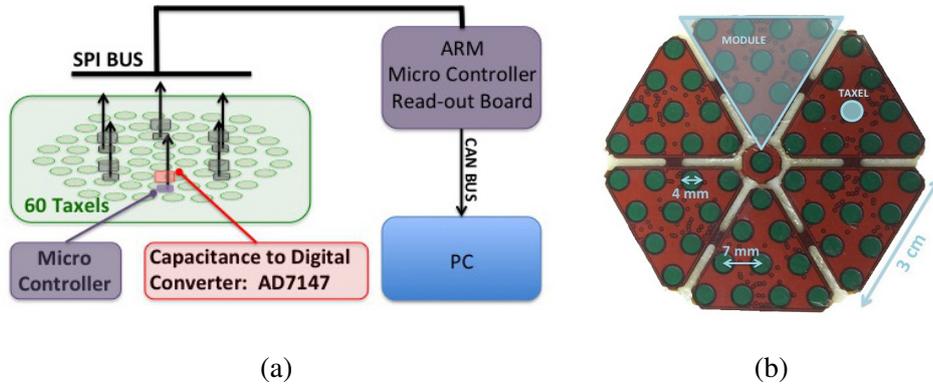


Fig. 6.3 (a) The CySkin technology architecture. The hexagonal patch is connected to a Intelligent Hub Board (IHB) that collect the tactile sensor data and send them to the PC through a CAN bus. (b) The CySkin patch used for the experiments. It is composed by 6 interconnected triangular modules, each hosting 10 taxels.

Task Table	Cluster 1	Cluster 2
<i>Task 1</i>	○	□ ○ □
<i>Task 2</i>	□	○ ○ □
<i>Task 3</i>	○	○ □ □
<i>Task 4</i>	□	○ ○ □
<i>Task 5</i>	○ ○	□ □
<i>Task 6</i>	□ ○	○ □
<i>Task 7</i>	○ □	○ □

Table Legend					
○	<i>Sphere</i>	□	<i>Cube</i>	○	<i>Half-Cylinder</i>
				□	<i>Cuboid</i>

Fig. 6.4 Task Table. Each task is a possible clustering outcome for the object set.

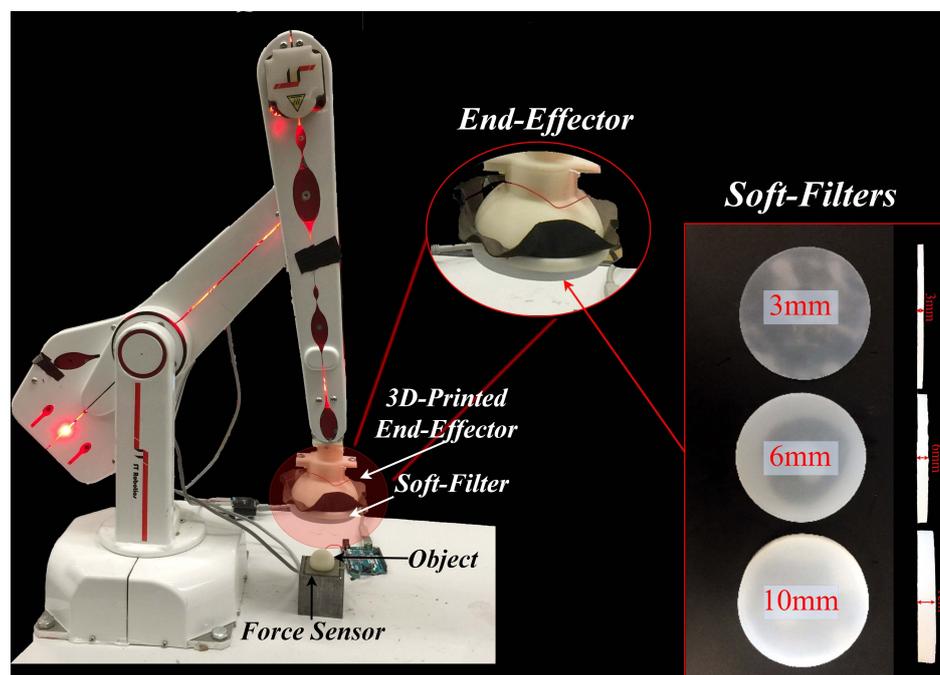


Fig. 6.5 The experimental set-up used for the experiments. The ST robot was used to push the sensorised end-effector against the object. A FlexiForce sensor A502 from TekScan was used for controlling the normal force applied. Three different soft filters were used in the experiments

In this chapter we refer to tactile images as the sensor readings for a specific object. To carry out the experiments we design and 3D-print a minimalistic set of four different objects with distinct features: a Cube (*side = 30mm*), a Cuboid (*side = 30mm, length = 80mm*), a Sphere (*radius = 30mm*) and a Half-Cylinder (*radius = 30mm, length = 80mm*). The objects present mainly two varying properties: long vs short length (Sphere & Cube vs Half-Cylinder & Cuboid) and edged vs non-edged surfaces (Cube & Cuboid vs Sphere & Half-Sphere). We define a task as a unique split of objects into two sets. Given the 4 objects it follows we can derive 7 different tasks (Fig. 6.4). A task here represents one of the possible ways we could wish to perceive similarities among objects. If we were to cluster objects according to Task 5, for example, we would be associating objects based on edges; while optimizing for Task 6 would signify grouping objects by length. Some of the tasks are conceptually less intuitive as no one particular feature can resolve the inclusion of an object in a cluster. As we are interested in the effects of Morphological Computation to the objects' associations, all 7 tasks are considered.

We carry out the experiments by mounting the printed end-effector, coupled with the tactile sensor, onto an ST-Robotics R12/5 robotic arm³. For each set of experiments we secure a different soft filter onto the end-effector flat's surface, and proceed by controlling the arm to descend perpendicularly down on the center of the object (Fig. 6.5).

We place a FlexiForce force sensor A502⁴ at the base of the object in order to apply a controlled perpendicular force when retrieving tactile images. The linear range of the sensor is 0-22N, however, we recalibrate its response in the 0-10N range and choose the maximal calibrated force of 10N, as this falls in the low-pressure regime (characterized as gentle touch [53]) for object exploration. We arrest the arm for the time needed to retrieve 10 consecutive tactile sensor readings and average them to create a tactile image. To further remove experimental bias, we repeat each set of experiments three times and average the computed tactile images, for each object, over the three trials (Fig. 6.6). We finally construct the tactile image matrix \mathbf{X} by setting each of its rows to a computed tactile image.

³<http://www.robotshop.com/uk/st-robotics-r12-5-axis-articulated-robot-arm.html>

⁴<https://www.tekscan.com/products-solutions/force-sensors/a502>

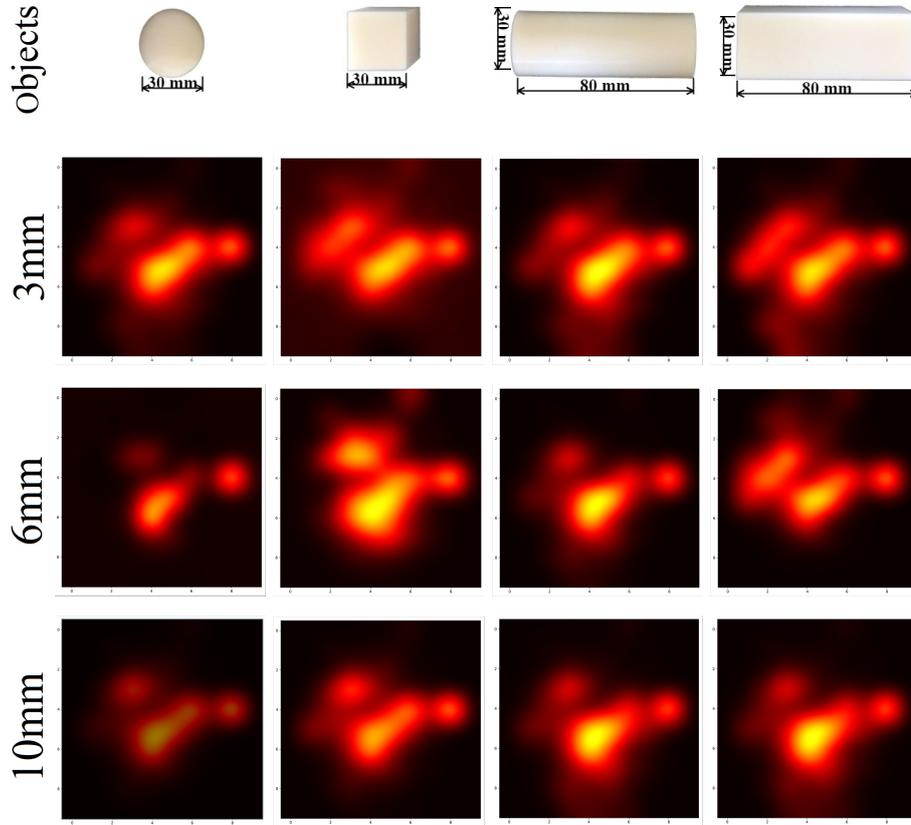


Fig. 6.6 3D-printed Sphere, Cube, Half-Cylinder and Cuboid (view from above) and relative tactile images computed (averaged sensor readings over three trials).

We utilize the process described in Section 6.2 to process the tactile image matrix for each experiment.

The unsupervised part of the process (PCA & KMC) clusters the objects automatically based on the two dimensions of highest variance in the data. We define the cluster matching process \mathbf{CM} as:

$$\vec{v}' = \mathbf{CM}(\vec{v}, \vec{t}_k) \quad (6.7)$$

Given a task \vec{t}_k and a cluster guess vector \vec{v} then, \vec{v}' is a new vector such that

$$\begin{aligned} \forall i \in \{1, 2, \dots, N\}. \\ (\vec{v}_i = 1 \implies \vec{v}'_i = 0) \wedge (\vec{v}_i = 0 \implies \vec{v}'_i = 1) \\ \iff \|\vec{v} - \vec{t}_k\| > \|\vec{v}' - \vec{t}_k\| \end{aligned}$$

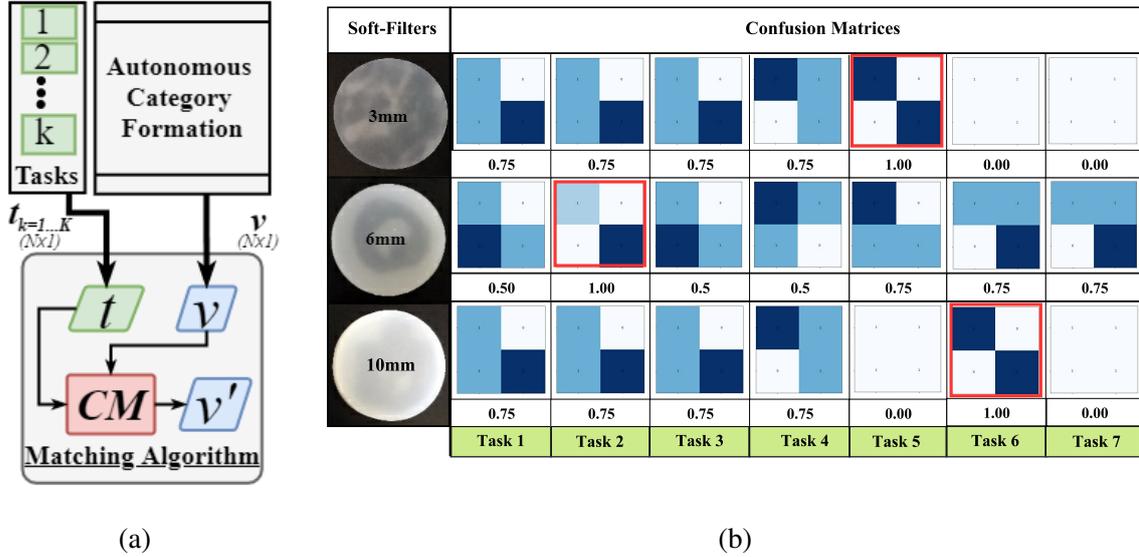


Fig. 6.7 (a) Process pipeline for the Cluster Matching Algorithm. (b) Confusion matrices and accuracy values for 7 different tasks corresponding to the clustering results obtained with three soft filters of 3mm, 6mm and 10mm respectively. The diagonal in each matrix retains the counts for the correct cluster guesses. Each soft filter is optimized for a specific task, highlighted in red.

i.e. we associate a cluster guess to a target cluster maximizing accuracy on a particular task (Fig. 6.7a). A vector $\vec{v} = [0\ 0\ 0\ 1]$ for a task $\vec{t}_k = [1\ 1\ 1\ 0]$, for example, would be re-associated as $\vec{v}' = [1\ 1\ 1\ 0]$. We utilize this to benchmark the performance of the algorithm in the various tasks (the object's inclusion in a cluster does not change after matching).

6.5 Experimental Results

6.5.1 Task Optimization

After the experiments, we observe the accuracy of the clustering with respect to the 7 predefined tasks. Fig. 6.7b illustrates the resulting confusion matrices. For each (2×2) confusion matrix \mathbf{C} , the darkness in square C_{ij} is proportional to the number of times an object class i was matched to a object guess j . The main diagonal then, contains the counts for the correct guesses, while anything outside of it is a mismatch. As clear from the figure each of the soft filters alters the clustering process significantly. The tactile images taken through the 3mm soft filter optimize clustering for Task 5 (*accuracy* = 1); The tactile images taken through the 6mm soft filter optimize clustering for Task 2; and finally, sensing through the 10mm filter clusters optimally according to Task 6.

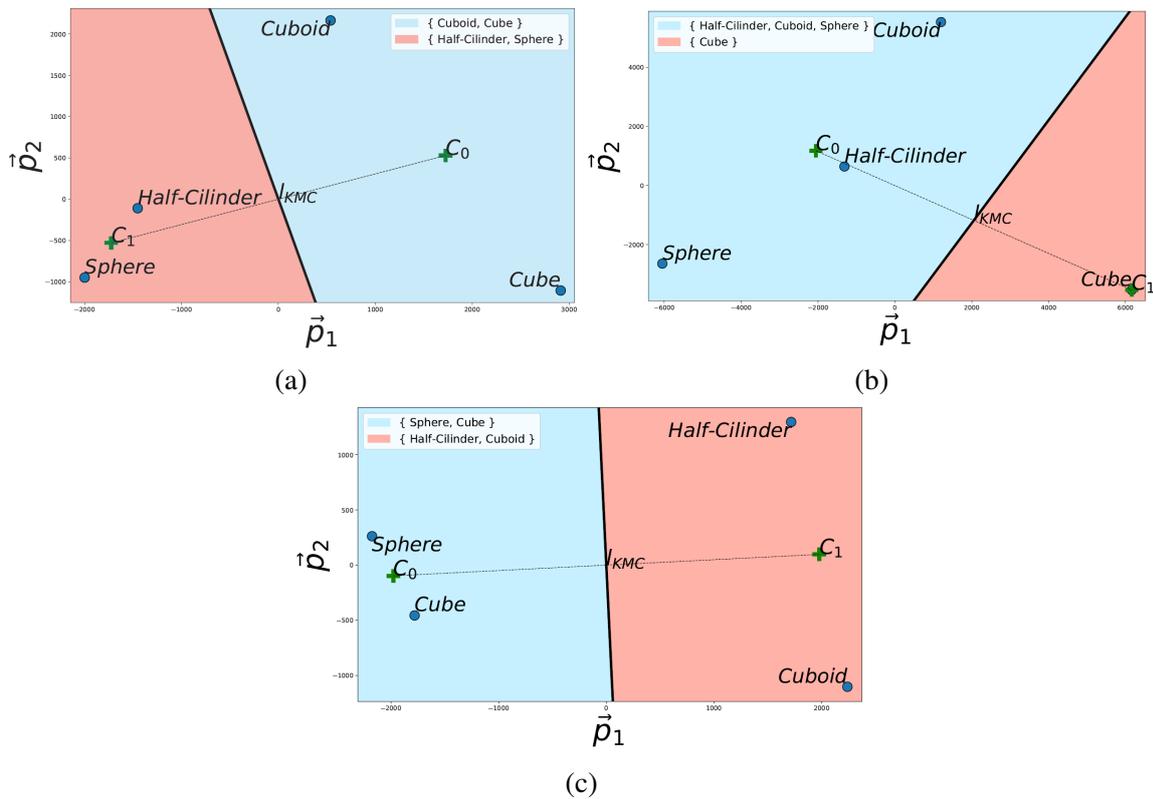


Fig. 6.8 The figure shows the 2-dimensional projection of each object on the two axis of highest variance in the data for the 3mm soft filter (a), 6mm soft filter (b) and 10mm soft filter (c). The line l_{kmc} corresponds to the decision boundary of the two clusters as found by the KMC algorithm (see Section 6.2, equation (6.6)), while C_0 and C_1 represent the cluster centroids. From the figure is it clear how the relative distance between objects changes when changing the soft filter, and the corresponding cluster assignment through the KMC algorithm.

6.5.2 Autonomous Category Formation variations

Fig. 6.8 illustrates the plots for each object in their optimal matched tasks. In the figure, the relative position of the objects to each other changes according to the soft filter used, drawing closer objects with respect to the morphologically processed features. The descriptions retrieved from the 3mm soft filter encode information relative to edges, and therefore draw together in space objects with or without edged surfaces (Cube & Cuboid vs Sphere & Half-Cylinder). As the thickness of the soft filter increases, the tactile sensor response becomes more blurred [238]. With thicker soft filters (10mm) the propagation of forces in the filter changes, and neighbouring taxels to the ones directly under the object are also affected. As edges, in a tactile image, become less and less sharp, another parameter (i.e. length) comes to induce the highest change in sensor readings from object to object. As

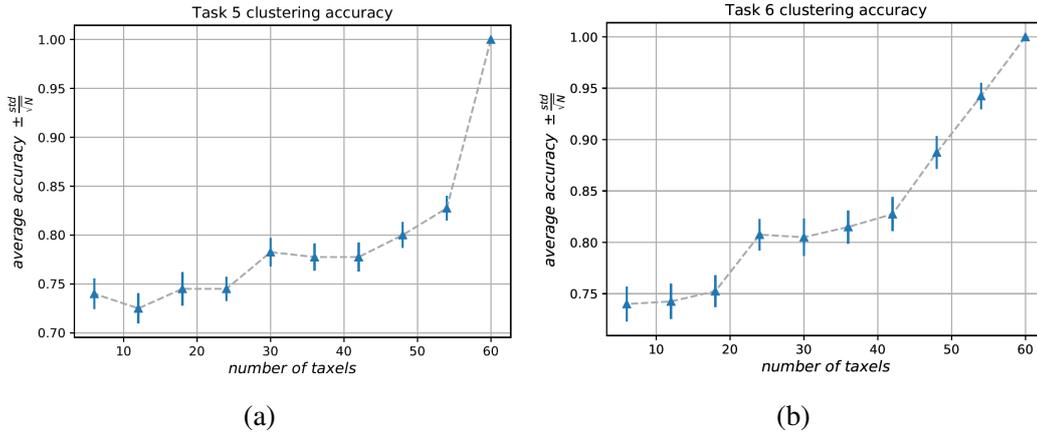


Fig. 6.9 The figure reports the average accuracy and error $e = \pm \frac{\sqrt{std}}{N}$ for the inference algorithm to cluster according to Task 4 (a) and Task 5 (b), when morphologically processing the data with their respective optimal filters ($3mm$ soft filter and $10mm$ soft-filter respectively).

a direct consequence the dimensions of highest variance, appointed by PCA, change from encoding edge information to encoding length deviations of objects, and eventually draw close in space the representation of objects with similar lengths (Cube & Sphere vs Cuboid & Half-Cylinder).

6.5.3 Spatial Resolution Influence

We test the reliability of the findings over tactile spatial resolution by running the Autonomous Category Formation procedure over subsets of s selected taxels. For each subset, we randomly select within the sensor an increasing number of taxels, where $s \in \{6i | i \in 1, \dots, 6\}$, and run the procedure 100 times. In Fig. 6.9, we report the average accuracy levels and errors over the performed runs for the optimal soft filter in Task 5 ($3mm$ filter) and Task 6 ($10mm$ filter). We find that the ability to morphologically process the data is highly dependent on the spatial resolution of the tactile sensor and that results are best when using ≈ 50 or more taxels. The findings highlight the need of a high spatial resolution tactile sensor for the analysis described in this chapter.

6.6 Discussion

After morphologically processing the tactile stimuli, we observe inherently different cluster guesses. Each soft filter alters the sensor response significantly, and induces the object descriptions (based on the two dimensions of highest variance in the data) to be qualitatively

largely different. The experiments provide direct evidence of how changing a single parameter of a soft filter can drastically change the way we perceive objects in the world. In the context of understanding relations amongst objects (for example clustering objects based on different features), the standard approach in the field is to change the inference mechanism to implicitly discern among features. Many of the used algorithms, in fact, need a large amount of data (usually labelled) which allows them to build an internal model of the objects and later do inference on the same. Understanding object properties in an unsupervised manner can be appealing, as there is no need of labelling or explicit modelling throughout the process. The experiments suggest we can drive the unsupervised findings by a careful design of the soft filters for a tactile sensor. As an indirect consequence, we show it is possible to optimize the tactile sensor's soft filter to drive the unsupervised inference algorithm into creating a useful world representation. In the context of manipulation or gripping mechanisms, for example, we may wish to grip an object based on a set of two or more properties. A soft filter can then be carefully designed to be optimal in extracting only the most relevant information for a task while filtering the others. The resulting clusters, then, would be retaining the feature information in terms of object similarities. By simple reinforcement learning, or other more involved strategies, a robot could then learn to grip an object in a cluster, and possibly generalize the gripping mechanism easily on other members of the cluster. In this scenario, no other information, besides cluster membership, would need to be known, and the human input in the process would be minimal.

6.7 Conclusion

We propose a concept to examine the way morphology affects the encoding of tactile sensor stimuli and analyse its effects on category formation. We actualize the concept by developing an unsupervised method for clustering a set of objects into two clusters. After integrating a capacitive tactile sensor onto a custom 3D-printed end-effector, we change the properties of a soft filter to alter the tactile stimuli and observe the change in cluster formation derived from the alteration. Results show that changing one parameter of the soft filter is enough to provide three qualitatively different representations of the objects. When clustering, we find the inference procedure relies on different object properties depending on the Morphological Computation applied. In this context, the *3mm* soft filter optimizes the inference procedure for edge detection while the thicker *10mm* object results optimal for elongation detection. A test on the reliability of the findings over various randomly selected set of taxels shows the results are highly dependent on the tactile spatial resolution of the sensor.

The work in this chapter has shown how it is possible to assess the quality of soft robot morphology via unsupervised methods. The soft interactions arising between the soft filter and the object under touched are heavily influenced by the soft properties of the filter itself. The ability to use information theory to autonomously reason about the influence of soft morphology to the physical stimuli is a powerful tool which can aid severely improve the current robot learning frameworks in Soft Robotics.

Chapter 7

Soft Actuation to Structure Information

As mentioned in previous chapters, morphology is not the only way to influence the sensory stimuli arising from physical interactions. In *SoMComp*, Soft Actuation can pertain a similar role, i.e. to influence the soft interactions, such that appropriate physical stimuli arise. In this chapter, we treat this topic for robotic palpation, where a robot is made to discriminate between areas of a soft phantom containing hard spherical inclusions, and areas without, by touch.

Reference Publication

This chapter was adapted from a journal article titled “**Structuring of Tactile Sensory Information for Category Formation in Robotics Palpation**” [231], published in 2020 in the journal *Autonomous Robots*. The article was written in collaboration with Dr Perla Maiolino, Mr Ed Bray and Dr Fumiya Iida and investigates the importance of appropriate end-effector trajectories when performing robotic medical palpation. Moreover it proposes an unsupervised framework to assess the quality of these trajectories. Dr Fumiya Iida and Dr Perla Maiolino contributed with the conceptualization of the topics, design of the experiments and the writing of the article. Mr Ed Bray contributed with the robot control and data collection. As first author in the article, my contribution includes conceptualization of the topics, the design of the experiments, the formulation of the mathematics for the framework, data analysis and article writing.

7.1 Introduction

7.1.1 Robotic Medical Palpation

In medical palpation diagnosis, given the nature of soft tissues in the human body, haptic perception plays a fundamental role [201].

Palpation is a key diagnostic examination, performed by medical practitioners for the exploration of abnormal masses or lumps within the body [256]. In this examination, practitioners use their hands to explore and feel for abnormalities within the soft tissue of the body, exploiting their physical dynamics and sensing capabilities [17]. This is a technique commonly used for the initial detection and screening of abnormalities within the breast [235], abdomen [28], thyroid [60] or other parts of the body. Abdominal palpation, in particular, is used for the detection of several different conditions [63, 202], ranging from cancer [221] and abdominal aortic aneurysm [132], to appendicitis [83]. Medical practitioners are required to acquire complex sensory-motor skills over many years of training to be able to appropriately perform this type of examination [22, 161]. Appropriate diagnosis can enable the early detection of potentially life-threatening conditions, aiding the success of treatments for the disease.

The strong dependence between the somatosensory system and motor actions in human palpation has been investigated in relation to the development of robotic palpation systems for detection of hard inclusions [87, 122, 124, 244, 277]. In the context of hard inclusion detection, the structure of sensory stimuli generated by physical palpation, helps to understand similarities or differences amongst the palpated objects. Through pertinent physical interactions, sensory stimuli of similar objects will maintain strong invariant similarities in the sensing space, whilst increasing their difference with dissimilar objects. In this context, the invariances allow for the dissociation of stimuli originated from different objects and the association, instead, of stimuli derived from similar objects. This fundamental process, corresponding to the separation and association of sensor stimuli into groups, will be referred to as categorization.

Palpation poses a particularly hard challenge for autonomous robotics systems. For a fully autonomous robotics system to achieve high diagnostic accuracy, it is firstly necessary to have an appropriate sensing and probing apparatus, capable of retrieving enough sensory information. The detection of small masses in complex soft bodies, however, is a challenging sensory task and it additionally requires the selection and learning of the appropriate physical interactions with the soft-tissue of the body, to obtain sensory information which can be used to make diagnosis possible [76, 251]. Human medical practitioners, in fact, must learn appropriate sensory-motor control. Appropriate physical palpation techniques emerge from

the understanding of the consequences of one's actions to the sensory stimuli for a given task or problem, in this case, the specific patient being examined.

Recently, our understanding of haptics and the importance of the human touch, and sensory-motor information has advanced significantly. In particular, the role of sensory-motor co-ordination and haptics has been explored for medical research [38, 253]. Sensor modalities beyond purely visual sensing has recently been emphasized, and the use of tactile sensors for tumour localization has thus been explored in some depth [47, 49, 218, 245, 251]. Finally, it has been investigated how augmenting haptic feedback to surgeons during teleoperated minimal invasive surgery can increase performance [121, 185].

In the past, many robotics palpation systems have been developed, for the detection of hard inclusions in soft tissue. In the early 1980s, the development of the WAPRO-4 showed how a robotics palpation system couple be capable of performing simple breast palpation to identify relatively large inclusions [118]. This has been followed in more recent years with increasingly anthropomorphic palpation systems [48] and abdominal palpation systems [49]. Over a decade prior to this work, Trejos et al investigated the ability to use TSI to augment the abilities of surgeons to perform internal organ palpation under minimally invasive surgery [262]. In [103] and [102], Hui et al investigated the use of SynTouch BioTac tactile sensor, coupled with a Gaussian inference model or Support Vector Machine classifiers, for the detection of soft plastic inclusions within a simulated silicon tissue sample. In [80] Gwilliams et al compared the ability of an artificial tactile sensor to that of a human finger, thus showing key differences between the two, and the ability of the former to outperform the latter in specific scenarios. In [134], Li et al proposed the use of a compliant capacitive tactile sensor array, between the tissue and a probing apparatus, to capture tissue properties during palpation. They show the technology is capable of imaging lumps, if somewhat dependent from their depth within the tissue. Further work focused on the employment of machine learning approaches to palpation [7, 177]

In all previous research, there has been little focus on the physical palpation techniques employed by the robot during the palpation examination. The robotic palpation techniques employed, in fact, were mainly simple vertical displacement, mostly with the tactile sensor constantly normal to the surface of the tissue under palpation. However, it is generally possible to positively influence sensory response through appropriate physical interaction, as advocated by the sensory-motor conditioning framework. In the context of palpation, the efficacy of diagnosis could be significantly improved by influencing the sensory response through appropriate palpation techniques [150, 228].

This Chapter addresses two related problems. First, we wish to investigate how motor actions can aid in the separation and categorization of tactile sensor information. Research

has previously shown that motor actions can introduce structure in sensory information [151, 194, 247], but it is yet to be understood which principles guide the emergence of such structure. Second, as later shown in this chapter, knowing the task to solve may not be enough to understand which physical interaction strategy is appropriate to use, or predict its effects to the tactile information. Here, instead, it is first necessary to understand the properties of the objects in interaction with the agent and the level of abstraction intended for the categorization.

In order to address the above problems this chapter investigates the processing of sensor signals based on dimensionality reduction and clustering. We explore the way active physical interactions with a soft body affect the structure of haptic spatio-temporal information.

7.1.2 Chapter Structure

The chapter is organized as follows: In Section 7.2 we describe the methods used, starting from the experimental set-up in Section 7.2.1, to the acquisition of tactile data through various probing strategies in sections 7.2.2 and 7.2.3. In Section 7.3 we describe the proposed framework. In Section 7.4 we report the results of the experiments followed by a case study in Section 7.5 and the conclusion in Section 7.6.

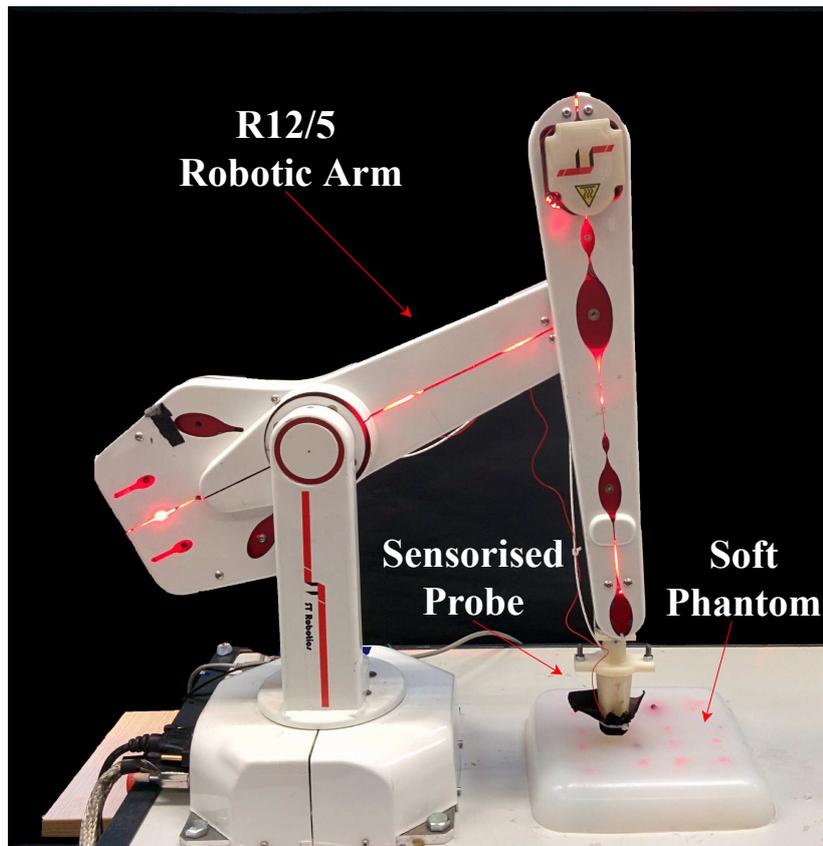
7.2 Methods

We arrange an experimental scenario where a robotic arm, equipped with an end-effector and a tactile sensor, probes the soft tissue of a soft phantom organ, to detect hard inclusions within it. The properties of the phantom organ designed to test the ability of the robotic agent to be detect hard inclusions by their depth and size, as shown to be important in previous systems [87].

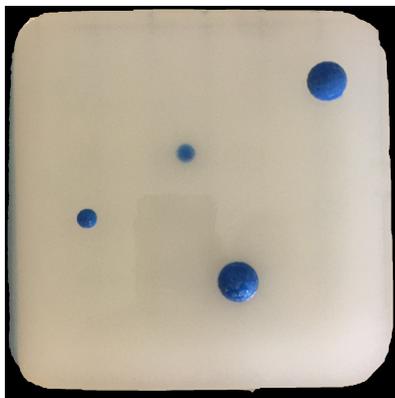
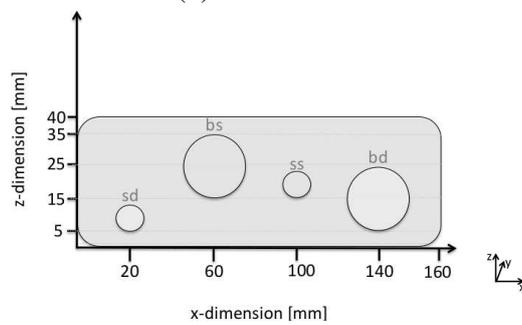
7.2.1 Soft Phantom and Robot Set-Up

We built two $160 \times 160 \times 40 \text{ mm}$ soft phantom organs using Ecoflex 00-10² from Smooth-on. The phantom organs are divided in 16 locations disposed in a coarse grained grid system as shown in Fig. 7.1e. Each location in the phantoms may or may not contain hard inclusions. An inclusion consists of a 3D-printed hard, spherical bead, embedded in the phantoms at a depth of either 5 mm or 15 mm , and having a diameter of 7 mm or 20 mm (Fig. 7.1d). Hereafter we may refer to a 7 mm inclusion placed at a depth of 5 mm as *SS* (Small-Shallow), a 20 mm

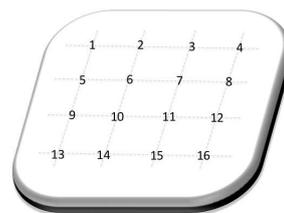
²<https://www.smooth-on.com/products/ecoflex-00-10/>



(a) Robot set-up

(b) *Ph-1*(c) *Ph-2*

(d) Inclusion types



(e) Phantom grids

Fig. 7.1

inclusion placed at 5mm as *BS* (Big-Shallow), a 7mm inclusion placed at 15mm as *SD* (Small-Deep), a 20mm inclusion placed at 15mm as *BD* (Big-Deep) and an area containing no hard inclusions as *NA*.

The experiments were performed on two phantoms: *Ph-1*, containing 12x*NA*, 1x*SD*, 1x*SS*, 1x*BS*, 1x*BD* (Fig. 7.1b); and *Ph-2*, containing 4x*NA*, 3x*SD*, 3x*SS*, 3x*BS*, 3x*BD* (Fig. 7.1c).

We 3D-printed a custom-made end-effector and integrated a capacitive tactile sensor onto its surface to retrieve *tactile images* during the probing experiments (Fig. 7.2a). The printed end-effector, coupled with the tactile sensor, was mounted onto an ST-Robotics R12/5 robotic arm³ (Fig. 7.1a).

7.2.2 Tactile Sensor Technology and Data Acquisition

High spatial resolution is a crucial component of the sensor technology necessary for the analysis in this chapter. The tactile sensor used has previously been described in Section 4.2.1.

In the current prototype, the tactile sensor module is placed on a 3D printed rigid hexagonal prism, to allow for probing experiments (Fig. 7.2a). The sensor schematic architecture is discussed in 3.2.3.

A sensor reading, or tactile image, from the tactile sensor described is produced at 20Hz , and corresponds to a 7-dimensional array, where each element contains the capacitance variation value of the corresponding taxel.

7.2.3 Probing Strategies

We control the r12/5 robotic arm open-loop in Cartesian coordinates. A teach-pendant was used to manually teach the robot the x-y location of the areas to probe. We use the stored end-effector positions in the subsequent control algorithm, where the robot automatically probes each location using the preferred probing strategy. We differentiate between two qualitatively different types of probing strategies, summarized in Fig. 7.2b: vertical and rotatory.

First, the vertical probing strategy is performed with the probe aligned vertically and plunged directly down into the phantom at 0.5mm increments. After each increment, the robot briefly pauses to allow a tactile image to be recorded before continuing with the next movement. This continues until the probe is at a depth d below the surface of the silicon, whereupon it stops recording and returns to a neutral position 10mm above the surface in a single movement.

³<http://www.robotshop.com/uk/st-robotics-r12-5-axis-articulated-robot-arm.html>

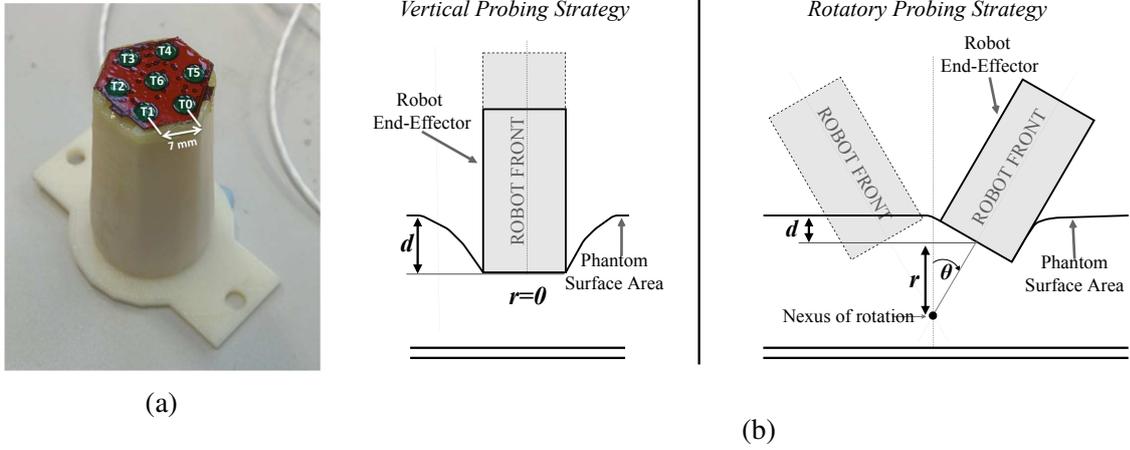


Fig. 7.2 The figure shows the set-up for probing including (a) the sensorized probe coupled with the CySkin patch used for the experiments, and (b) the diagram of the two probing motions employed. The architecture for the sensor is shown in Fig. 3.3a. The vertical probing motion is performed when $r = 0$ and is described by the parameter d . The rotatory motion is performed with $r > 0$, and is fully described by both the d and r parameters.

Second, the rotatory motion is performed with the robot d mm below the surface of the silicone, rotating about a nexus point r mm away in the vertical direction. To reach the initial position of this motion strategy, the robot moves vertically downward from its rest position, until it reaches the position set by d . Hence, a nexus point r distant from the end effector is assumed, and the robot rotates about it in the $+\theta$ direction until it is at an angle of 30° from its initial, vertical, position. Here, the palpation action can begin. The probe rotates in the $-\theta$ direction at 1° increments, recording a tactile image after each step. Once the probe has rotated of 60° it stops recording, and returns to its rest position 10 mm above the surface of the silicone.

In general, a probing strategy can be uniquely identified by a depth d and a radius r , thus:

$$\Theta = \left\{ \begin{array}{c} d \\ r \end{array} \right\}, \quad (7.1)$$

where if $r = 0$, the probing motion will be vertical, while if $r > 0$ the probing takes place via the rotatory strategy (Fig. 7.2b).

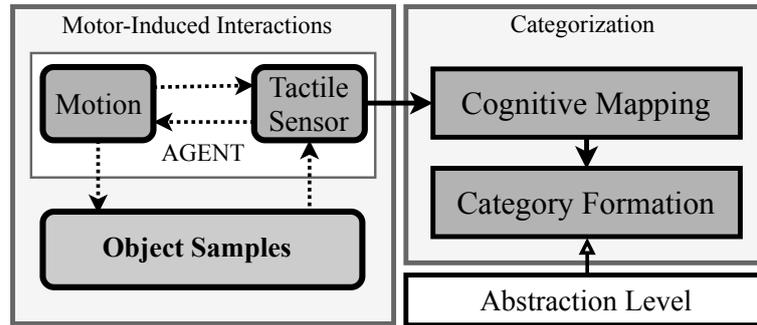


Fig. 7.3 Analytical Framework.

7.3 Analytical Framework

In this chapter we consider the framework in Fig. 7.3. In the framework, an agent retrieves tactile sensor information while interacting with samples of objects, defined by a task. Here, the tactile information is directly influenced by the interactions with the samples. A categorization system allows for the information to be: first, re-encoded into a meaningful, lower-dimensional space (Cognitive Mapping); second, differentiated into useful categories (Category Formation). The abstraction level corresponds to the number of categories that should be observed in the sensor information and has a direct influence on the significance of the formed categories. At its limit, 2 categories might be too coarse to be useful in capturing differences amongst different types of objects, while a number of categories equal to the number of object samples is impractical in identifying any similarities amongst them, and therefore amongst similar objects. The direct influence of the physical interactions to the tactile information, if substantial, should be observable in the category formation process.

7.3.1 Task and Physical Interactions

Within the considered framework, the agent is an embodied system equipped with a tactile sensor, and capable of performing probing actions. The interactions consist of physical probing, through different strategies, of target areas in a soft phantom, as was described in Section 7.2.3. As exemplified in Fig. 7.4, an experiment consists of an agent probing a preselected phantom with a chosen probing strategy. The agent iteratively selects a target area in the phantom to probe, and performs the chosen probing strategy for the experiment (described by Θ) while acquiring and storing tactile information. After probing all intended areas the stored sensor information can undergo categorization.

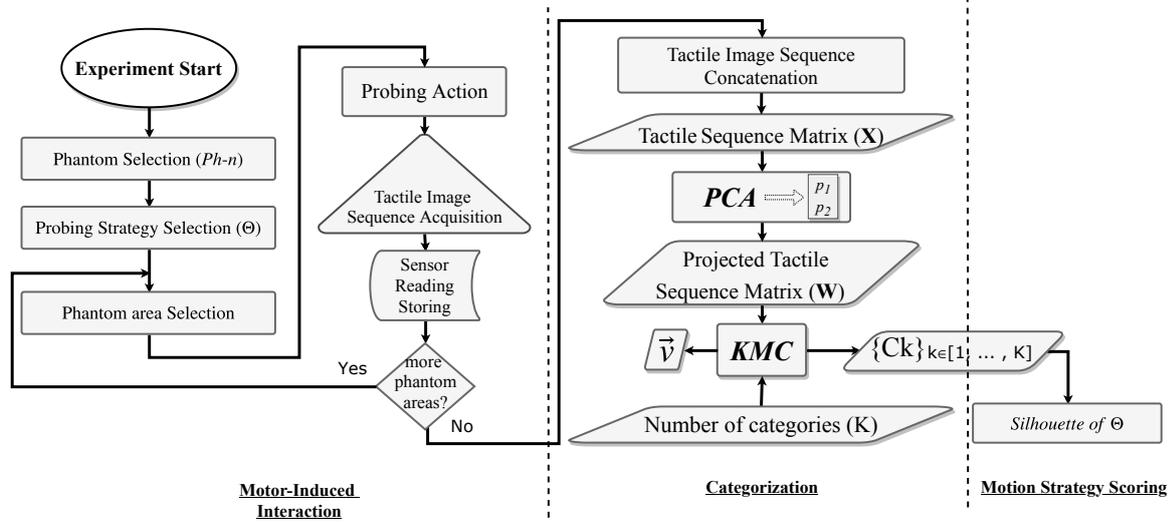


Fig. 7.4 Implementation steps of the Theoretical Framework.

7.3.2 Categorization

Cognitive Mapping

A process is needed to reduce the high dimensionality of the spatiotemporal data acquired through the tactile sensor, while interacting with the environment. We define a tactile image sequence as a series of tactile sensor readings taken at set intervals, and concatenated into a single array. After acquiring tactile image sequences for each probed location, we use Principal Component Analysis projection (*PCA*) [257] to reduce the dimensionality of the acquired data [144].

For a set of N different locations in a phantom, let \mathbf{X} be a $(N \times D)$ matrix where each unique tactile image sequence for a probed location is a D dimensional row in the matrix. The dimension of D , then, will be strictly dependent on the probing strategy and on the interval at which the agent captures each tactile image within the sequence.

After obtaining the tactile image sequences matrix \mathbf{X} , we use equations 6.1, 6.2, and 6.3 to create a scatter matrix \mathbf{S} of \mathbf{X} , and factorize it into matrices \mathbf{Q} and Λ . Similarly to Chapter 6, the \mathbf{Q} matrix is such that each column q_j corresponds to an eigenvector of \mathbf{S} , and each element λ_{jj} in the diagonal matrix Λ is its corresponding eigenvalue. Furthermore, we proceed to form a 2D projection matrix \mathbf{P} with the two eigenvectors corresponding the the two highest eigenvalues in Λ , and use equations 6.4 and 6.5 to form an $(N \times 2)$ matrix \mathbf{W} , where each row in the matrix is a 2-dimensional *encoding* of a tactile image sequence for a probed location.

Category Formation and Abstraction Level

To observe the effects of the probing strategies to the tactile sensor information we wish to have a process to categorize the re-encoded sensor information. We use K-Means Clustering (*KMC*) to find clusters in the data, where each found cluster will represent a potential category of inclusion types. The analysis is analogous to that of Chapter 6, to observe whether the same type of qualitative influence can be observed in the stimuli due to the actions of the robot. The processes, however, differ slightly, as in the context of palpation they need to be generalized for a multi-class classification scenario. The abstraction level is set by the number of clusters we wish to find in the data. We initialize the *KMC* algorithm with random centroids, and split the re-encoded sequences in \mathbf{W} into K clusters by:

$$\vec{v} = KMC_K(\mathbf{W}) \quad (7.2)$$

The resulting \vec{v} is an N -dimensional array, where each element $\vec{v}_i \in \{1, \dots, K\}$, and $\forall i \in \{1, \dots, N\} \exists j \in \{1, \dots, N\} : i \neq j \wedge v_i \neq v_j$ (Fig. 7.4); in other words, none of the resulting clusters can contain all the sample areas in the phantom. Like previously, $\vec{v}_i = k$ only if the i^{th} tactile image sequence belongs to cluster k , thus the \vec{v} vector contains the cluster membership of each probed location in the initial set.

To avoid cluster anomalies due to the random centroid initializations we run the *KMC* algorithm three times and discard the clustering attempt if, after convergence, any of the three cluster guess vectors differs from any other. At the end of the clustering process a list of centroids C is obtained, uniquely dividing the space into K categories (7.4). In this context, the cluster assignments for each probed location is largely dependent on the probing strategy employed.

The category formation is an unsupervised analysis to the data, and it is thus useful if performed on all available data at the time of analysis.

7.3.3 Motion Strategy Scoring

At the end of the clustering process it is necessary to be able to assess the usefulness of the probing in generating meaningful data for classification. For the unsupervised clustering algorithm to be able to find meaningful clusters in the re-encoded tactile data, it is necessary that the data exhibits structure. Therefore we score the probing strategy that generated the data via a metric of structure tightly connected to the type of clustering utilized in this chapter, i.e. the silhouette score [212].

The silhouette score $s(i)$ for cluster i can be computed as:

$$s(i) = \frac{b(i) - a(i)}{\max(a(i), b(i))} \quad (7.3)$$

where $a(i)$ is the mean intra-cluster distance of cluster i , and $b(i)$ is its mean nearest-cluster distance. We will refer to the silhouette score s as the average score for each cluster found by KMC, i.e.:

$$s = \frac{\sum_{i=1}^K s(i)}{K} \quad (7.4)$$

The score will thus be a number $s \in [-1, 1]$, where data exhibiting more structure will score higher s values.

After probing the selected phantom through various probing strategies, the maximum observed silhouette score can identify which probing strategy is capable of generating structured data for hard inclusion detection and classification. The analysis as described thus far can be done without any prior labelling, and can thus be applied to all available data at the time of analysis. After, a supervised method can, for example, be used to perform the classification.

7.3.4 Experimental Procedure

We execute 180 experiments, each of which sees the robot probing all 16 areas of *Ph-1* or *Ph-2* with the preferred Θ parameters. The experiments are carried out for all combinations of $d \in [6.5mm, \dots, 20.5mm]$ at $1mm$ increments and $r \in [0mm, 10mm, 12mm, 14mm, 16mm]$. The bounds were chosen to reach the minimal/maximal experimentally feasible probing depth and rotation with the robotic arm, and the devised soft phantoms.

Given the procedure two datasets are collected, each consisting of time-series tactile data collected from either the first or the second phantom under examination (*Ph-1* or *Ph-2*). Each dataset contains 90 sets of experiments, each of which is composed of 12 data-points, so 1440 data-points are present for each of the datasets. Each data-point is a time series of tactile images with variable length (each tactile image is a 7-dimensional array). For the rotatory motion, each tactile image was taken at 1° increments, while for the vertical motion each image was taken at $0.5mm$ increments.

For each of the experiments, after the probing has ended, the time-concatenated data is used to form the tactile image sequence matrix described (see Section 7.3.2). The matrix can then be used to re-encode the tactile sensor information for each probed location into a lower dimensional space (*Cognitive Mapping*). After clustering, each probed location will be differentiated into one of a predetermined number of categories (*Category Formation*).

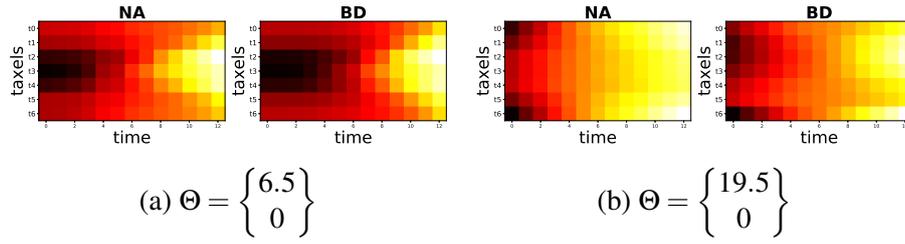


Fig. 7.5 Raw spatiotemporal tactile image sequences, as captured when probing *Ph-2* vertically at varying depths, in an area containing no hand inclusion, and an area containing a 15mm inclusion placed 20mm deep. (a) and (b) correspond to a re-shaped x_j .

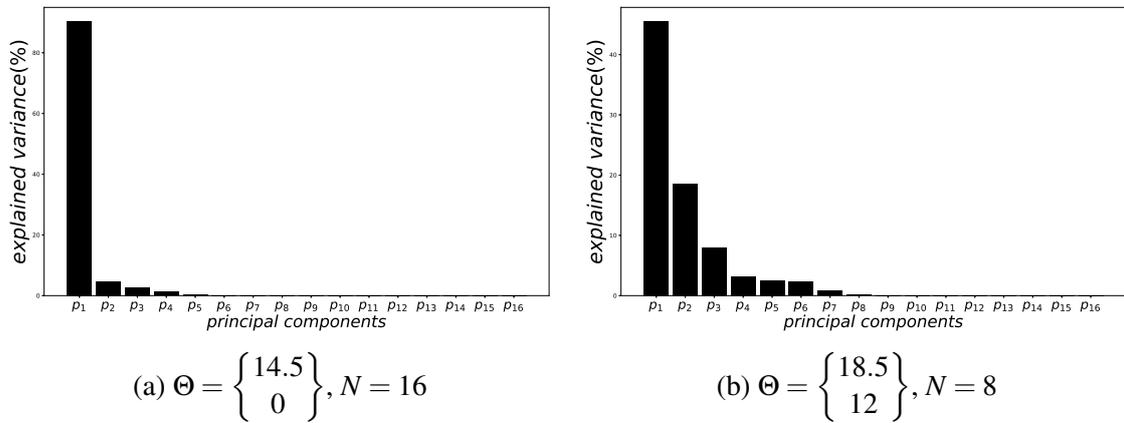


Fig. 7.6 The explained variance of each principal component when projecting the \mathbf{X} matrix belonging to two different experiments where both the number of probed areas in *Ph-2*, to base the *PCA* projection on, and the Θ parameters where changed.

7.4 Results

The following sections will progressively analyse the described framework, starting from the dimensionality reduction process (*PCA*), to the repercussions of physical interactions to categorization (*KMC*).

7.4.1 Sound Dimensionality Reduction

One of the principal components of the proposed framework is the reduction of the high dimensional spatiotemporal tactile information, into re-encoded lower dimensional data. An example of the acquired tactile information is shown in Fig. 7.5. Without knowing which category each tactile sequence vector \vec{x}_i belongs to, it is impossible to assess the quality of dimensionality reduction from \mathbf{X} to \mathbf{W} . However, it is feasible to maximize the information retention in the original tactile sensor data.

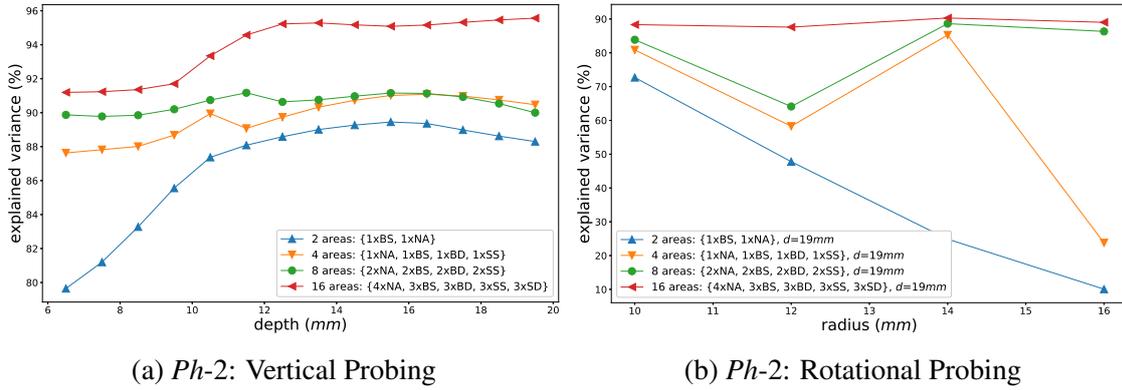


Fig. 7.7 The change in explained variance by the 2D PCA subspace projection, when probing vertically (a) and through the rotatory motion (b), changing the number of samples used to find the principal components (N in \mathbf{X} , see Section 7.3.2).

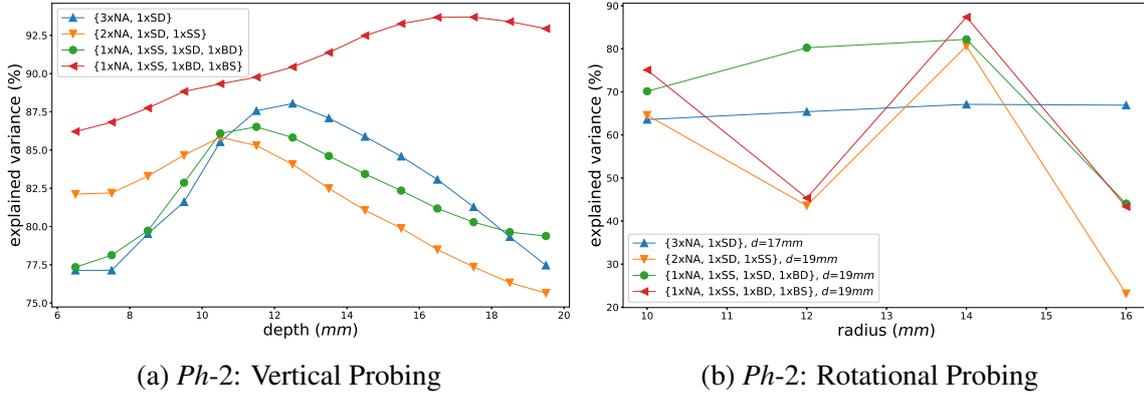


Fig. 7.8 The change in explained variance by the 2D PCA subspace projection, when probing vertically (a) and through the rotatory motion (b), changing the quality of the samples used to find the principal components, while maintaining their number constant.

The explained variance can be thought of as a measure of the information captured by the PCA subspace after projection. As the eigenvalues in Λ (see Eq. 6.3) are proportional to the variance captured by the corresponding PCA principal components, we can compute the explained variance τ_i for the principal component \vec{p}_i as:

$$\tau_i = \frac{\lambda_i}{\sum_{j=1}^N \lambda_j} \quad (7.5)$$

where λ_i is the eigenvalue corresponding to the i^{th} principal component. Here, τ_i is a measure of the proportion of variance in the data, captured along the direction the principal component \vec{p}_i in the original sensor space.

Fig. 7.6 shows the explained variance of each \vec{p}_i , after the robot probed *Ph-2* in two different experiments where both Θ and the number of probed areas used for the projection (N) were varied. As clear from the figure, the number of probed areas and the Θ choice significantly affect the distribution of the sensor data in its original D space. In one case, the sensor data is mainly spread along 7 axis ($\vec{p}_1 - \vec{p}_7$) (Fig. 7.6b), making it unsuitable for dimensionality reduction. In the other, instead, \vec{p}_1 captures the majority of the information in the data (Fig. 7.6a). The figure suggests the suitability of the tactile information to the drastic reduction in dimensionality is dependent both on the properties of the probed areas, and probing strategy employed.

We further explore the way the probing strategy, and the properties of the probed areas in the phantom, affect the amount of information retained after dimensionality reduction. The explained variance achieved prior to categorization is $I = \tau_1 + \tau_2$. Fig. 7.7b shows the explained variance trends when the number of probed areas used for *PCA* projection varies. When the number of probed areas is maximal (16 areas, red plot in Fig. 7.7b), the influence of Θ is negligible. Conversely, with less data to base the *PCA* projection on (2 areas, blue plot in Fig. 7.7b), the choice of Θ can be the sole determinant to induce structure in the data. A second interesting phenomenon can be observed in Fig. 7.7a, when comparing the explained variance obtained after projecting \mathbf{X} based on 4 vs 8 probing areas in the phantom (yellow vs green plots). Here, the agent retains more information, even when basing the projection on less data, if the employed probing is vertical and at a depth of at least 17.5mm.

This result suggests that proper physical interaction can help information retention in the absence of enough data.

Ultimately, we observe the influence of the quality of the data samples to the information retention after *PCA* projection. Fig. 7.8a shows how in presence of very diverse inclusion types (left triangle plot), the effects of the vertical probing strategy Θ to I is negligible. The presence of very diverse data, in fact, is useful for *PCA* to find good projection axis. In absence of good data, or non-diverse inclusion types, instead, appropriate interaction can minimize information loss (peaks in Fig. 7.8a and 7.8b). In the figures, it is possible to see how the least diverse set of samples can yet induce the tactile information to retain most of the information when the phantom is appropriately probed (peak in triangle plot, Fig. 7.8a).

The explained variance analysis shown allows for a preliminary assessment of the robot palpation action employed. For any one palpation we show that good palpation actions correspond to information that can be safely projected onto a two dimensional space, preserving most of the information in but a few principal components.

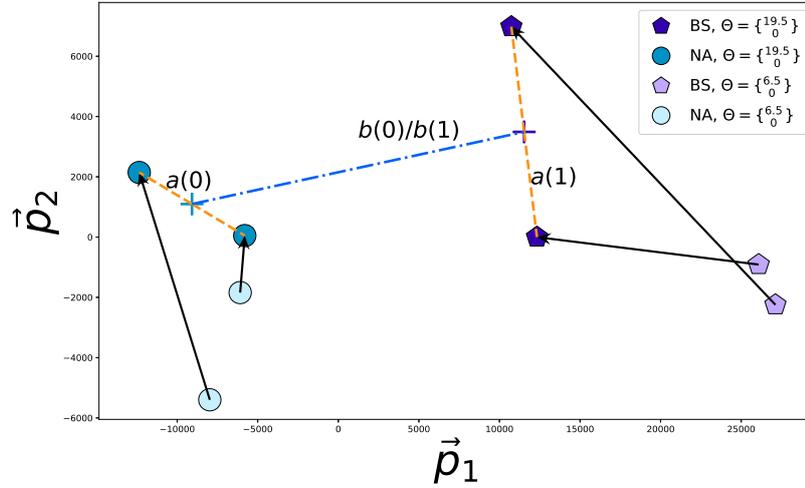


Fig. 7.9 The change in position for the 2D PCA projected *NA* and *BS* samples when probing the phantom vertically at a depth of 6.5mm and 19.5mm . The yellow and blue line show the two parameters on which the silhouette score is based, i.e. intra-cluster distance and nearest cluster distance respectively.

7.4.2 Information Structure and Silhouette Coefficient

Similarly to the previous sections we wish to observe the effects of changing the Θ parameters to the structure of the information after *PCA* projection. The silhouette coefficient, as explained in Section 7.3.3, depends on the mutual mean intra-cluster distance, and mean nearest-cluster distance for each pair of clusters (Fig. 7.9).

Fig. 7.10 and Fig. 7.11 both show how the change in Θ influences the silhouette score. This influence, however, is primarily dependent on N and the diversity of the inclusions probed, as suggested by the change in trends of the plots in each of the figures. Fig. 7.10a shows that little structure emerges when probing *Ph-2* vertically too superficially or too deeply. In both cases, in fact, the sensor response is uniformly too moderate or too steep to have any variation from an area of the phantom to another, thus inducing no variation in the information. Fig. 7.10b, instead, shows how, when in absence of enough data samples (2 areas, blue plot), a correct choice of Θ can be the sole determinant for good or bad structure in the information. In Fig. 7.11a and Fig. 7.11b, interestingly, it is shown how even without much diversity in the inclusion types, good structure can emerge when the phantom is probed appropriately ($\Theta = \begin{Bmatrix} 16.5 \\ 0 \end{Bmatrix}$ or $\Theta = \begin{Bmatrix} 14.5 \\ 16 \end{Bmatrix}$).

At last, we investigate the influence of the number of clusters K to the structure of the information s . The number of clusters sets the level of abstraction that the robot may wish to have to make use of the tactile information, and directly affect the interpretation of the

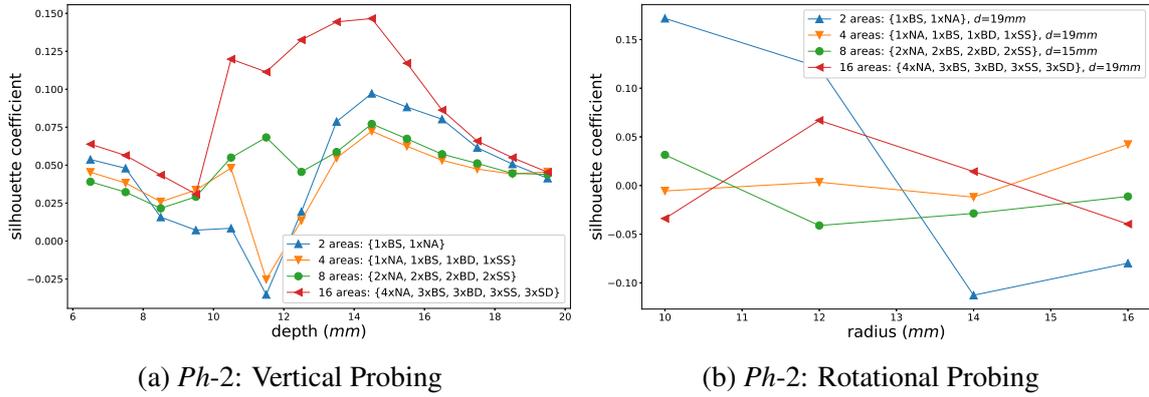


Fig. 7.10 The change in silhouette coefficient by the 2D PCA subspace projection, when probing vertically (a) and through the rotatory motion (b), changing the number of samples used to find the principal components (N in \mathbf{X} , see Section 7.3.2).

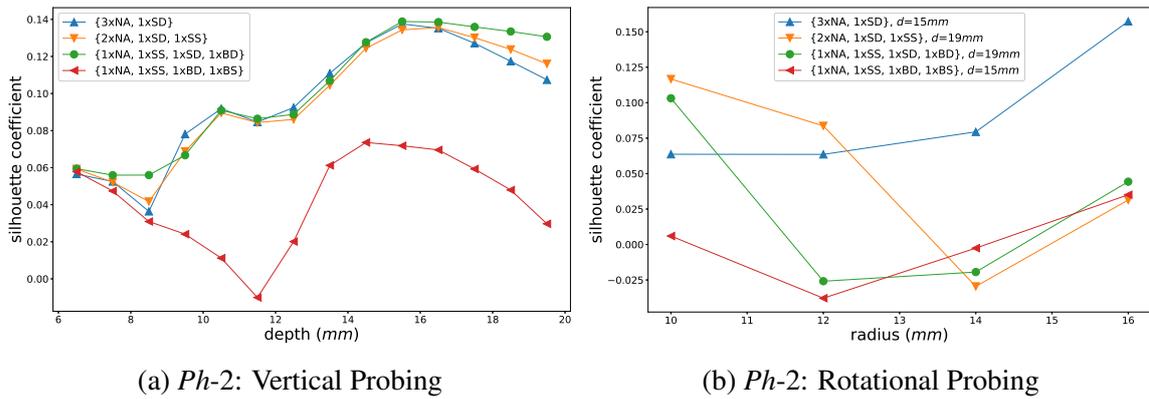


Fig. 7.11 The change in silhouette coefficient by the 2D PCA subspace projection, when probing vertically (a) and through the rotatory motion (b), changing the quality of the samples used to find the principal components, while maintaining their number constant.

emerging clusters. We choose three varying number of clusters: $K = 2$, presence vs. absence of an hard inclusion; $K = 3$, absence vs. small vs. large inclusion; $K = 5$, all inclusion types. Fig. 7.12 shows the trends when probing the soft phantom vertically at varying depths and changing K in the KMC algorithm. The emerging clusters present different structural properties. The different trends in the figure suggest how K directly affect the way the probing strategy influences the structure of the data. Interestingly, probing at a deeper depth increasingly helps to sense inclusions, or detect their size. To dissociate between all different inclusion types, instead, an optimal probing depth is found for $d = 14.5mm$, after which the increasingly high sensor response converges, and renders the clusters less separable, thus decreasing the values of s .

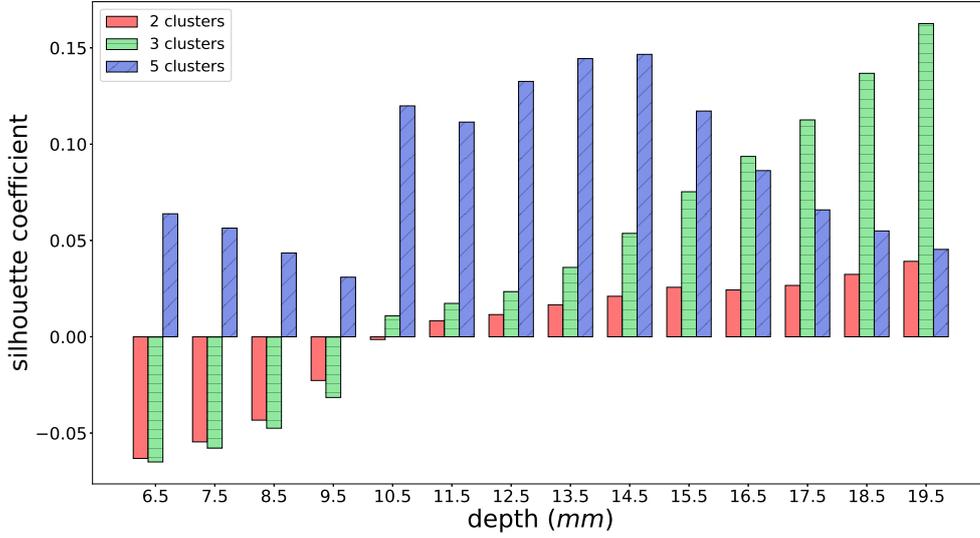


Fig. 7.12 The silhouette score of PCA projected tactile sensor information for every probing area in the soft phantom, when performing the probing action at different depths, and over varying number of clusters.

7.4.3 Motion influence on Cognitive Maps

Predicting the effects of Θ to the low-level encoding of the information in \mathbf{W} is a highly complex process. Understanding such effects, however, would allow an agent to appropriately choose a Θ when solving the probing task.

To understand this relationship we make a plot of the cognitive maps for each set of motion parameters in Θ and observe how the encoding of each probed area changes according to the probing strategy used. Here, to have a better understanding of the motion effects, we perform the experiment on the least cluttered phantom, i.e. *Ph-1* (Fig. 7.1c), which would suffer less from disturbances due to the vicinity of adjacent inclusions. Fig. 7.13a and 7.13b show the plots corresponding to probing the phantom vertically at the minimal and maximal experimental depth. By increasing the depth of probing, two very interesting effects take place: one, nearest cluster distance $b(i)$ between almost all types of inclusions increases, allowing for better dissociation of diverse tactile information; two, the intra cluster distance $a(i)$ between any two probing areas with the same type of hard inclusion decreases, allowing for each possible phantom inclusion type to be better represented.

Extending the analysis to the rotational probing strategy we can similarly observe the effects of changing the parameters in Θ from their minimal to their maximal experimental values. Interestingly, when employing the rotational strategy, the generated tactile information presents a structured layout, by which it is already possible to dissociate one stimulus type from another. In this scenario, then, the effect of the rotational parameter r to the structure of

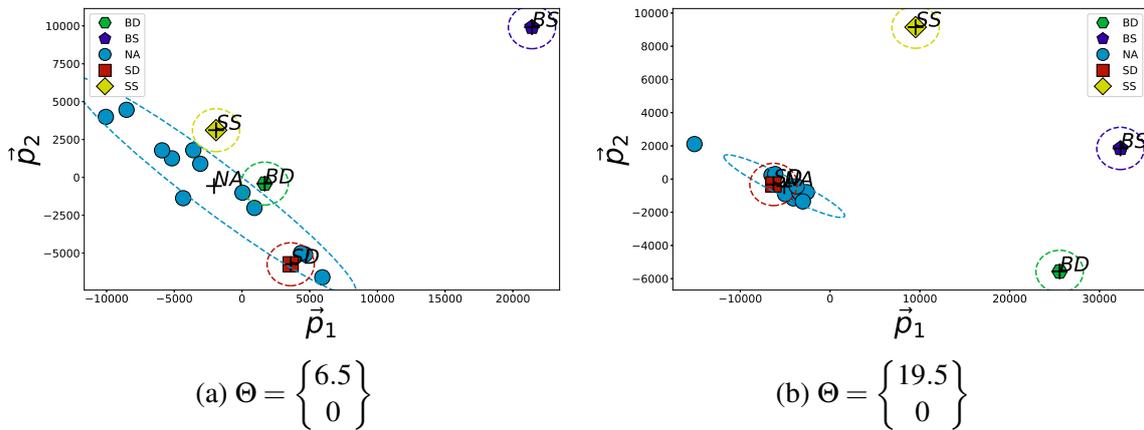


Fig. 7.13 The 2-dimensional projection of the tactile information generated from probing *Ph-2* at varying depths. The ellipses correspond to the distributions of the clusters based on their true inclusion types, at a distance of 2 standard deviations from their cluster centers.

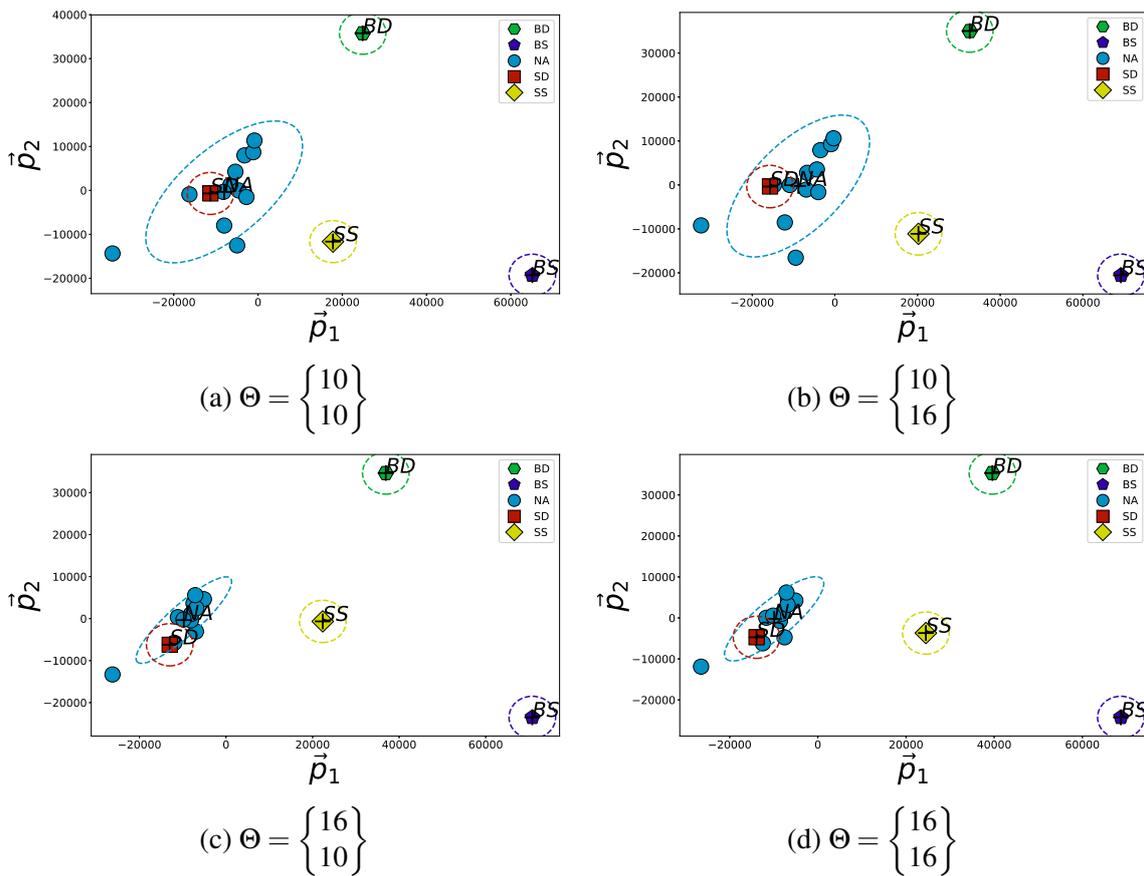


Fig. 7.14 The 2-dimensional projection of the tactile information generated from probing *Ph-2* at varying depths and radii. The ellipses correspond to the distributions of the clusters based on their true inclusion types, at a distance of 2 standard deviations from their respective cluster center.

the data s appears to only mildly act upon the nearest-cluster distance parameter (Fig. 7.14a to Fig. 7.14b). The effect of increasing d , instead, confirms the hypothesis by which the probing depth influence acts upon the intra cluster distance of each stimulus type.

The effect of the depth parameter can be attributed to the strength in response of the sensorised probe. The tactile sensor, in fact, detects pressure levels on its surface. When probing the phantom at the minimum depth, the pressure registered by the sensor is mostly due to the elastic response of the Ecoflex 00-10 soft phantom, almost independently from the presence or absence of inclusions in the probed area. As the depth increases, the elastic response is influenced by the non-elasticity of the hard inclusion, should there be one in the probed area. We hypothesize this influence can be captured by the sensor response in three ways: first, the response should be higher when inclusions are present in the probed area; second, the sensor's increase in detected pressure should arise at slightly different sample intervals depending on where the inclusion is placed in the phantom (deep vs shallow inclusion); third, the area of the response should vary depending on the size and depth of the inclusion.

In this framework, an acceptable probing depth is one which neither saturates the sensor response in each area, nor fails to detect changes in pressure when the probed area contains non-elastic inclusion. The task of dissociating amongst all different types of inclusions is optimized (i.e. maximal silhouette score) for $\Theta = \begin{Bmatrix} 12.5 \\ 0 \end{Bmatrix}$ in $Ph-1$ and $\Theta = \begin{Bmatrix} 14.5 \\ 0 \end{Bmatrix}$ in $Ph-2$. This analysis can be applied to any one dataset, to explore which way the robot action has influenced the haptic data in terms of information structure. The action parameters generating the data with the maximal silhouette score can thus be used to perform palpation on the tissue under examination.

7.4.4 Categorization and Similarity Abstractions

In robotics palpation, proper physical interaction can help in the dissociation of tactile information, such that the emerging clusters can be meaningful with respect to solving a task (e.g. finding hard inclusions in a soft phantom). Besides dissociating amongst different object types, however, another fundamental, yet usually neglected, fragment of information is related to the similarity associations between clusters. The distances between found clusters in the 2D re-encoded tactile information subspace, in fact, grants the agent the possibility to associate types of objects, and order or rank them based on such association.

In the context of probing a soft phantom to find hard inclusions, for example, the agent might need to prioritize possible findings based on the depth of the inclusion, e.g. $[NA, SD/BD, SS/BS]$, we'll refer to this as rank-1. In a different scenario where the size of the hard inclusion should take priority over its depth, the ranking might, for example, change to $[NA,$

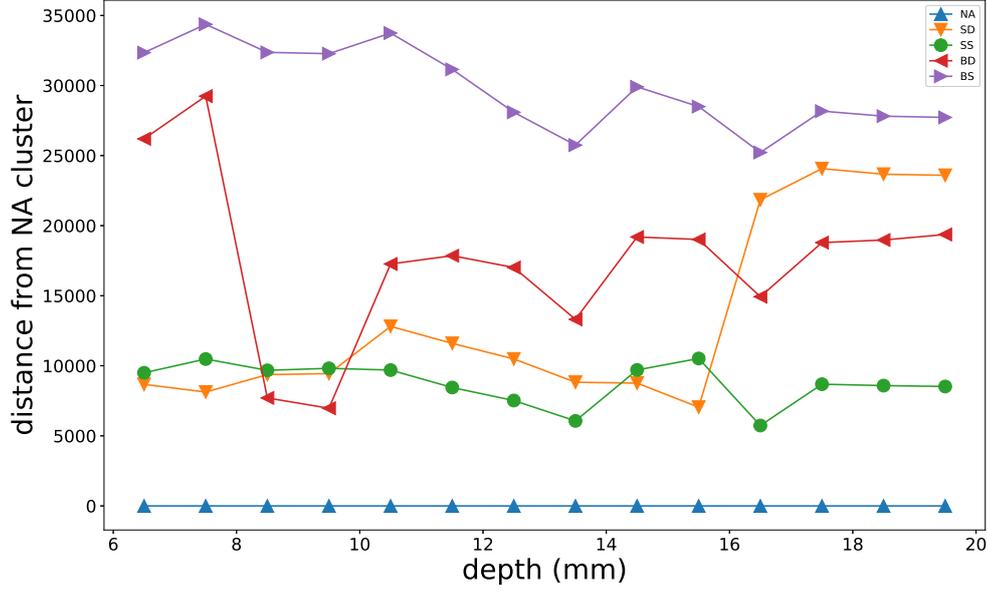


Fig. 7.15 The distance between the cluster-matched *NA* cluster and all matched clusters in the data. The data is captured when probing *Ph-2* and , setting $r = 0$ and varying the d in Θ .

SD/SS, BD/BS], or rank-2. In this scenario, the influence of the physical interactions with the soft phantom may induce the agent to see some inclusion types as more similar to others, depending on which property is deemed more important.

To assess the performance of category formation in each experiment, we first need to match the clusters found by the KMC algorithm to any set of target classes for the phantom under analysis. We devise a cluster matching process based on maximal accuracy.

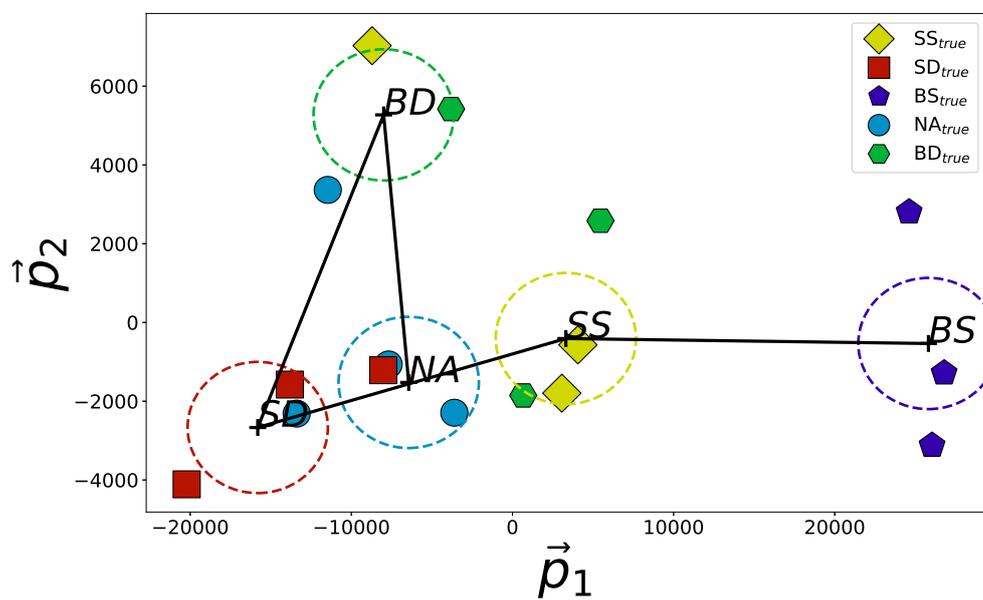
We first define a function Γ such that

$$\Gamma(\vec{v}, \vec{C}) = [x \mid x = \vec{C}_{\vec{v}_i} \text{ for } i \in [1, \dots, N]] \quad (7.6)$$

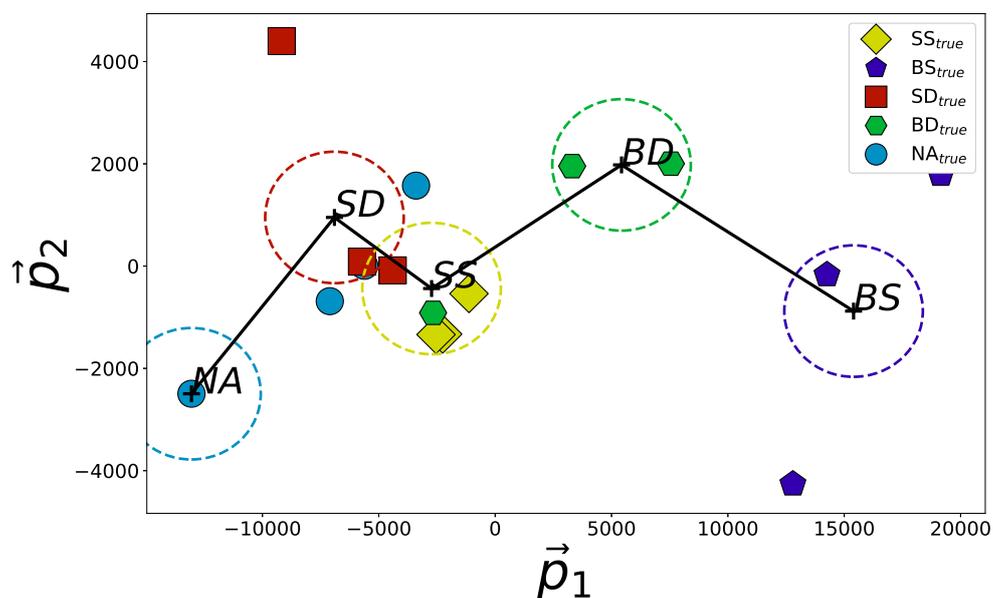
where \vec{v} is the N -dimensional class membership vector, where each element can belong only to one class, and C is a list of centroids which uniquely divide the space into K categories. Moreover, \vec{v}_i is the i^{th} element in \vec{v} , $\vec{v}_i \in \vec{C}$, and $\vec{C}_{\vec{v}_i}$ is the \vec{v}_i^{th} element in C . The function remaps the elements in \vec{v} based on \vec{C} .

Given a target vector \vec{t} we define a function Ψ to re-associate the classes in C such that the distance between the target and the guess vector is minimal, thus:

$$\Psi(\vec{v}, \vec{C}) = \operatorname{argmin}_{\vec{C}'} \|\Gamma(\vec{t}, \vec{C}') - \vec{v}\|$$



(a) $Ph-2, \Theta = \begin{Bmatrix} 9.5 \\ 0 \end{Bmatrix}$



(b) $Ph-2, \Theta = \begin{Bmatrix} 9.5 \\ 0 \end{Bmatrix}$

Fig. 7.16 The emerging cluster similarities when changing the motion parameters and solving for either rank-1 i.e. $[NA, SD/BD, SS/BS]$ (a) or rank-2 i.e. $[NA, SD/SS, BD/BS]$ (b). Each dotted circle is placed on the cluster-matched, KMC found, cluster corresponding to the color coding in the legend.

where $C' \in \mathcal{S}(\vec{C})$, $\mathcal{S}(C)$ is the set of all permutations of C , and $\|\cdot\|$ is the Euclidean norm of a vector. Finally we define the cluster-matching as:

$$\mathbf{CM}(\vec{v}, \vec{t}, \vec{C}) = \Gamma(\vec{v}, \Psi(\vec{v}, \vec{C})) \quad (7.7)$$

We use the cluster-matching process to re-associate the cluster memberships

$$\vec{v}' = \mathbf{CM}(\vec{v}, \vec{t}, \vec{C}). \quad (7.8)$$

Here \vec{v}' is a new vector maximizing accuracy for a particular task given (specified by the target vector \vec{t}). A vector $\vec{v} = [2 \ 2 \ 1 \ 0 \ 0]$ for a task $\vec{t} = [1 \ 1 \ 0 \ 2 \ 2]$, for example, would be re-associated as $\vec{v}' = [1 \ 1 \ 0 \ 2 \ 2]$. We utilize the cluster memberships in \vec{v}' to compute each cluster center and retrieve the mutual distances between clusters. This cluster matching algorithm is a generalization of the one proposed in Chapter 6, to account for multi-class problems.

In this analysis we consider two scenarios where we may want to associate the clusters by depth or size of inclusion, and use the *NA* type as ground zero, we thus consider the distance from the cluster-matched *NA* inclusion type and the remaining types (Fig. 7.15). As clear from Fig. 7.15, by duly interacting with the soft phantom, the distance between each cluster type and the *NA* cluster changes drastically. In this context, then, it is possible to induce a ranked understanding of robot's perceived similarities between different inclusion types by simply acting on the Θ parameters.

We demonstrate the ability to achieve similarity relationships of the kind previously described by finding the parameters for which the agent can rank the system based on rank-1 or rank-2. We perform the experiments in *Ph-2*, and we use the experimental data gathered through the probing of the soft phantom to find the parameters by which we can solve the ranking. We find the robot capable of abstracting similarities relationships according to rank-1 for $\Theta = \begin{Bmatrix} 9.5 \\ 0 \end{Bmatrix}$ (Fig. 7.16a), and according to rank-2 for $\Theta = \begin{Bmatrix} 15.5 \\ 0 \end{Bmatrix}$ (Fig. 7.16b).

7.5 Palpation Test Case

We perform experiments to test the ability of the framework developed to assess and identify the motion control which can best allow an agent to differentiate among different types of inclusions. For this purpose, the robot is set to perform palpation on a phantom containing 4x*NA*, 3x*SD*, 3x*SS*, 3x*BS*, 3x*BD*. The sensorized robotic arm is made to palpate the phantom vertically on each location, as described in Section 7.2.3. At this point, dimensionality

reduction is used to pass from a high dimensional sensor description of each palpated phantom location, to a two dimensional descriptor based on PCA analysis (see Section 7.3.2).

After dimensionality reduction it is possible to utilize Equations (7) through (9) to assess the quality of each motion strategy with respect to the collected data. The motion strategy parameters generating the highest structure in the data can thus be saved.

Here we make use of a standard classification procedure to dissociate amongst the different types of inclusions, and we assess the ability of the framework described in this chapter to assist in determining which motion would have generated the best data for palpation classification. We use a off-the-shelf multi-class Support Vector Machine (SVM) [40] classifier, as implemented in the *scikit-learn* python tool [188].

The dataset utilized for this test scenario consists of 224 data-points, each consisting of sequence of tactile images. The data corresponds to vertical palpations performed at 14 different depths spaced by 0.5mm each, thus for each depth 16 data samples are present. An SVM classifier is trained on a single sample for each type of inclusion (one-shot learning), at each different depth. We thus fit 14 different SVM classifiers, and we show how the unsupervised analysis run on the same data is capable of faithfully predicting the most performing action parameter before any supervised learning is necessary.

Three different type of classification are executed, following the same qualitative analysis in Section 7.4.2. First a classification with two classes, where the SVM classifier is trained to discriminate between locations containing hard inclusions, and locations with no inclusions. Second, three classes, where the classifier is trained to discriminate between large inclusions, small inclusions or no inclusions. Third, 5 classes, where all inclusion types are considered. For each of the three classification types, the classifier is trained on the minimal possible number of inclusions per class, i.e. 1 sample, and the data-set is split into training and test set accordingly. For each probed depth, the 16 data-points are therefore divided into 5 samples for training and 11 for testing. This is done at all 14 different depths. The split was purposefully chosen to observe the classifier performance when lacking large amounts of data.

After training, the SVM classifier separates the two dimensional space according to the two, three or five classes, maximizing the distance to the nearest training data points of any class. Once the classifier has been fit to the training samples, we test the ability of the SVM to classify a new inclusion correctly by testing it on the unseen phantom test locations.

Fig. 7.17 shows the resulting accuracy of the classifier at different probing depths and when classifying the inclusions following the three different sets of classes described. Given the difficulty of the classification task with the limited amount of data, the classifier can only achieve an average classification accuracy of 68.78% when detecting hard inclusions, 36.26%

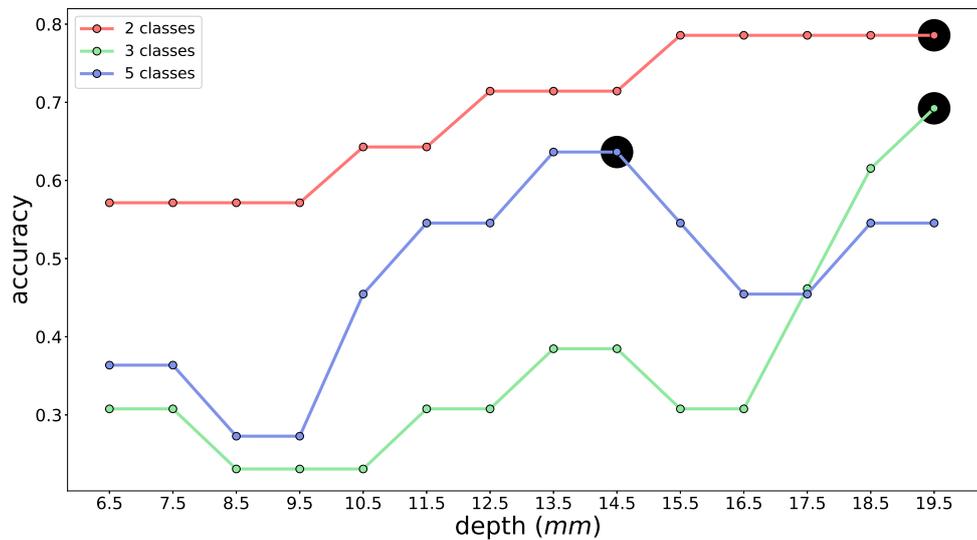


Fig. 7.17 The classification test accuracy of a multi-class SVM trained on a single sample for each inclusion type, when performing a vertical probing action at different depths, and over varying number of clusters. The highlighted black circles correspond to the maximal silhouette score computed through the proposed framework (see Fig. 7.12).

when detecting inclusions based on size and 47.40% when discriminating inclusions based on all their properties. Even in this scenario, the motion strategy detected by the proposed framework can achieve accuracies of respectively 78.57%, 69.23% and 63.63% in the same tasks, improving on the average classification accuracy of up to 10-33%, as shown by the black circles in Fig. 7.17. More significantly, when comparing Fig. 7.12 and Fig. 7.17, it becomes clear how the general performance of the classification can indeed be predicted by the framework proposed, by solely relying on information structure. In fact, additionally to the best performing motion strategy, both the motion parameters resulting in the least accurate classification, as well as the general flow of the accuracy graph in Fig. 7.17 can be almost faithfully predicted based on the scores in Fig. 7.12.

7.6 Conclusion

In this chapter we investigated the effects of various motion strategies to the response of a capacitive tactile sensor, for the task of detecting hard inclusions in a soft body. Actively choosing an interaction strategy, to optimize sensory reception for a specific task at hand, has the potential to be a powerful tool. Such tool could endow robots with the ability to dynamically filter properties of touched objects, actively helping in the completion of a task [21, 182] even before the sensor information arrives to a central processing unit.

The experiments were performed by embedding a capacitive tactile sensor onto a 3D-printed end-effector, and probing two soft phantoms with various hard inclusions through different probing strategies. The sequential sensor data obtained through the probing of each area in the phantom was clustered, and the change in information due to each strategy observed and analyzed.

We found the amount of information retained after PCA projection to be highly dependent both on the probing strategy and the properties of the sample areas in interaction. More interestingly, we found that appropriate probing strategies can help retain information even when lacking a large quantity or good quality of it. Using the explained variance as a measure of information is useful in ensuring large amount of heterogeneity is kept in the data, but it is not capable of ensuring the quality of the information retained. In fact, it could be possible that the projection makes the information relative to highly distinct object, indistinguishable after projection. However, under the assumption of no prior knowledge of target labels, keeping variance in the data is usually a sensible choice. The proposed analysis can therefore help choose those actions which allow sound dimensionality reduction, with the minimum loss of information variance in a low dimensional inference space.

Furthermore, we analysed the impact due to motion on cognitive maps and extracted how the motion influenced the tactile information. This analysis is useful in understanding the effects of motion to the perception of the probed areas, and can be used to appropriately choose an interaction strategy that generates structure. To make full use of such effects, however, it would be ideal to instead be able to predict such change, before interaction takes place. Here, the change in position of each point within a cognitive map could be interpreted as a transformation in the same domain. The transformation function could be learned from initial interaction and used in future tasks to optimize the sensor response for a specific task. The transformation function, however, would not only be dependent on the motion parameters employed, but also on the properties of the sample objects in interaction, like demonstrated in the results.

It is also possible to take categorization one step further and abstract similarities between object types from Cognitive Maps. Here we have shown that the physical interaction can drive the similarity relationship between objects. In an unsupervised scenario, the abstractions can be highly informative and can, for example, be useful to fix an ordering, via mutual distances, on the sensed object types. The object ordering can be purposefully fixed to the agent's advantage. In a real scenario a practitioner might diagnose the gravity of a detected inclusion based on various features. In our fictitious example we show how it is possible for an agent to prioritize over two features by simply changing the palpation strategy.

The unsupervised analysis framework proposed in this work is meant to assess how the robot's palpation technique may influence its ability to diagnose hard inclusions in soft tissue. Without need for explicit labels or knowledge of whether the tissue under palpation has abnormal lumps, the framework can still inform a robotic agent on what type of haptic interaction is likely to be most discriminative. As such the framework is most useful when used as a pre-learning step, before any actual supervised learning takes place. We later show this with a simple supervised test case, where the most discriminative palpation actions are found to be coherent with the unsupervised analysis proposed. In the test case application of the proposed framework a robot is made to palpate a clustered phantom, and an SVM multi-class classifier is trained on the minimal possible number of samples per class. The classifier is shown to perform best when employing the highest scoring motion strategy, as detected by the proposed framework. The chosen strategy is shown to improve the classification accuracy of the classifier of up to 33%. More interestingly, we observe the silhouette analysis based on our method can predict the general relative performance of the classification a priori. The SVM based classifier utilized is effective in showing the usefulness of the analysis, however, it is too simplistic to outperform any other state of the art supervised learning system. The SVM, in fact, performs simple on-shot learning, and tries to classify lumps based on only one example of each type. Although not in the scope of this research, more complex inference methods can also be considered, and are likely to still benefit from analysis shown in this chapter.

As a parametric discretization of the sizes of the lumps was necessary for the analysis in this chapter, the work described palpation on 5 different types of spherical inclusions, thus no evidence was shown for other types or sizes. As the analysis itself was independent of the size or type of lump, we believe it can extend to any-one type of lump parametrization necessary, as shown by the consistency of the results when sub-sampling the lumps in type groups of 2, 3 or 5 types (see Section 7.5). However, a second assumption was due to the location of the inclusion, should there have been one present in the tissue under palpation. Here it is key that the type of lump is approximately the same across experiments, this is true also of its location with respect to the examining probe. Should this not hold true, it is possible the unsupervised clustering method may classify two lumps of the same type into different clusters. Here, additional research is necessary to address the need of haptic search algorithms to locate, rather than discriminate, between lumps in a soft tissue.

Finally, the parameters were optimized with respect to the actual tissue under palpation, and are thus likely to be valid for the phantom devised for these experiments. For any new tissue under palpation, however, the same analysis can be applied, and new optimized parameters retrieved without knowledge of whether an inclusion is or is not present under

the palpated tissue. This can serve as a first filtering procedure before carrying out more expensive, supervised experiments, where the breath of possible robot action may be too large to be employed. Moreover, the analysis can retrieve motion parameters which either maximise information retention in lower dimensional sensor inference space, and/or maximize information structure within the retrieved haptic data.

In the context of this thesis, this chapter has highlighted the important of actions to achieve *Soft Morphological Computation*. In *SoMComp*, this is one of the two principal tools that enable the robot to appropriately influence the soft interactions such that structure emerges in the physical stimuli. This filtering was shown through unsupervised methods, and the ability of the action to influence perception, and thus learning, was shown within the context of palpation. One of the most prominent issues with the current approach is the exponential blow up in the number of experiments necessary to find good parameters for action, and the discretization of the robot's action itself, which is based on human intuition and design. Another limitation instead lies in the parametrization of the action, which limits the complexity of the robot action to achieve the task at hand. The next chapter will treat some of these issues, while leaving others for discussion in chapter 11. Complexity is here a very important topic, as the complexity required in the action to appropriately condition the stimuli might need to be naturally high, so as to off-set the complexity of the soft interactions between the robot and its environment. This topic will be treated more in depth in the next chapters, where large scale physical experimentation is used to achieve appropriate complexity in the robot actions.

Chapter 8

Action Complexity: Soft-Body Palpation

As mentioned in previous chapters, complexity of the action space is a pre-requisite for any meaningful influence of complex soft interactions. This is mainly because the an infinite number of interactions are possible when considering the infinite number of states of a soft robot and/or its soft environment. This chapter investigates the benefits of action complexity, to achieve the appropriate conditioning of the soft interactions. Moreover this chapter starts to treat the topic of complex parameter search. This is an important related topic, as complex actions are usually generated by similarly complex parameter combinations, which may be hard to explore with common methods. In Chapter 6 and Chapter 7 we show how through simple unsupervised machine learning methods it is possible to determine both which

Reference Publication

This chapter was adapted from a journal article titled “**A Bayesian Framework for Multi-Axis Soft-Body Palpation**” [225], currently under revision in the Soft Robotics journal. The article was written in collaboration with Dr Perla Maiolino, Dr Josie Hughes, Mr Liang He, Dr Thrishantha Nanayakkara, and Dr Fumiya Iida and proposes a mathematical framework to use multi-axis palpation trajectories for robotic medical palpation. Dr Fumiya Iida and Dr Perla Maiolino contributed with the conceptualization of the topics, design of the experiments and the writing of the article. Mr Liang He and Dr Thrishantha Nanayakkara contributed with the physical creation of the phantom liver organ for testing purposes. Dr Josie Hughes contributed by developing the phantoms for testing purposes and the writing of the article. As first author in the article, my contribution includes conceptualization of the topics, the design and execution of the experiments, the robot control, the formulation of the mathematics for the framework, data analysis and article writing.

action and morphology to employ to solve a predetermined task. As of yet, the robot must first interact with all objects, with all possible strategies and/or morphologies, before understanding which strategy and morphology improves the information extracted. If the robot were to discriminate between K object types, for example, in a scenario considering N objects of each type, and if the robot motion could be described by one motion parameter which could take any of M discrete values (or a morphology determined by the same parameter), then the number of experiments n_e necessary to retrieve the optimal motion strategy would be $n_e = KNM$. The experiment number increases linearly with each new motion strategy parameter, object number or object type. It is often the case that a motion strategy or morphology is described by more than one parameter with many values. For any more than one parameter $M = J^s$ where s is the number of parameters describing the motion or morphology, and J is the number of discrete values any one parameter can take, thus bringing $n_e = KNJ^s$. The exponential blow up in the search space must therefore be avoided for larger scale problems. In this chapter we assess the usefulness of complex actions, and propose a framework to efficiently search the robot parametric space given sensor evidence, thus avoiding the exponential blow up in the search space derived by the use of purely clustering methods. We also wish to move forward from the robotic medical palpation experiments in Chapter 7, by performing medical palpation experiments in a more complex and realistic soft-tissue environment, as well as using a more dynamic and realistic interaction strategy by the robot.

8.1 Introduction

Like explained in Chapter 7, Palpation is a key examination procedure used by the medical profession for the diagnosis of abnormalities [22, 256]. Practitioners use their hands to explore and feel for abnormalities within the soft tissue of the patient's body, exploiting the physical structure and the sensing capabilities of the human hand [17]. This action is widely used for the initial detection and screening of abnormalities within the breast [235], abdomen [28] or thyroid [60]. Detection of abnormalities through palpation can aid in the diagnosis of conditions including cancer [221], abdominal aortic aneurysm [132], appendicitis [83] and many others [63, 161, 202], making it a powerful diagnostic procedure.

Despite its importance, palpation is still poorly understood, because of the complex nature of the physical interactions involved in this examination. This complexity is due to several factors. Firstly, the human body is constructed from interacting layers of soft tissues, organs and muscular-skeletal systems, each with different mechanical properties. Tissues can have many (or infinite) degrees of freedom and show significant variation in dynamic properties which makes the nature of interactions hard to model, or predict. Secondly, abnormalities are often 'hidden' beneath layers of skin and tissues, making them physically challenging to sense. Finally, every patient is different; practitioner must have the ability to learn and adapt to different interactions across different patients, and immutable solutions are not appropriate.

The importance of medical palpation and the role of robotics within this field has been explained in Section 7.1.1.

Although there has been progress made in understanding palpation, there remains a considerable gap in our knowledge of the physical interactions that occur during palpation. In particular, there has been a limited investigation of the impact of introducing diversity and complexity into the trajectory of the robot hand/probe during palpation. Previous work has only examined robotics palpation systems with simple one-axis vertical displacements [48, 80, 102, 103], or horizontal sliding trajectories [80, 100, 262]. In contrast, medical practitioners use complex examinations techniques, including rotations, twists and percussions[202]. Palpations can be light-touch, deep, or performed at several angles to feel for different organs, including liver, kidneys and the aorta [17]. Thus, to fully understand palpation it is necessary to investigate palpation interactions when multi-axis robot trajectories are introduced. In particular, it is crucial to understand whether more complex palpation actions give rise to richer interaction dynamics and can improve the sensory information that is gained, allowing for better or more confident classification of abnormalities.

As such, following the analysis in Chapter 7, the challenge addressed in this paper is the optimisation of complex palpation trajectories to enable more accurate classification of

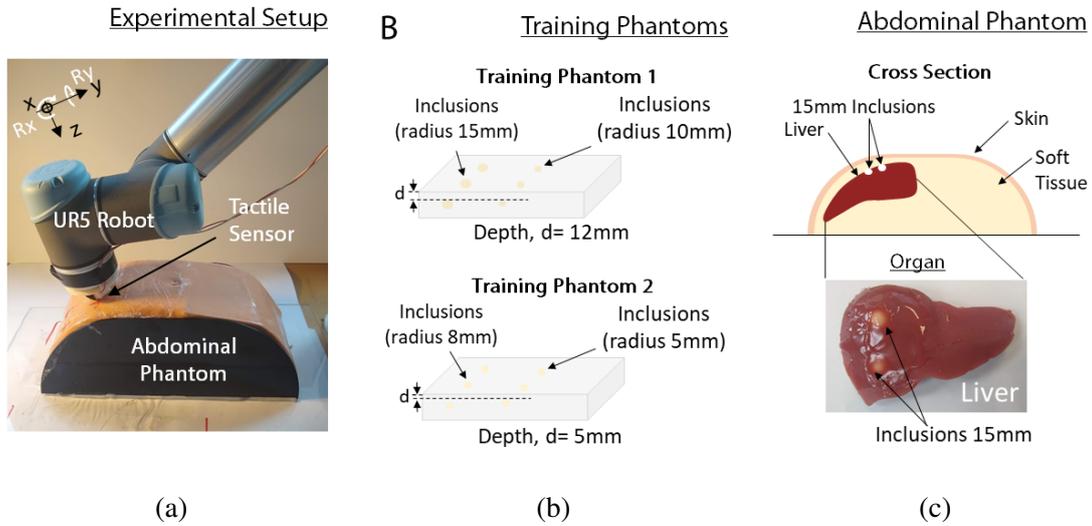


Fig. 8.1 Robotic Medical Palpation, including (a) the Experimental Setup, (b) the Training Phantoms and (c) the Abdominal Phantom developed.

abnormalities. We use a robotic manipulator that is capable of complex rotational and translational motions. To capture the rich palpation dynamics, the manipulator has a sensorized end effector with a capacitive tactile sensor providing sensory information from 7 high-resolution tactile elements (taxels). Using this setup, we perform experiments on anthropomorphic ‘phantoms’, which replicate realistic palpation conditions, including a multi-layered internal soft structure containing hard abnormal inclusions [100].

The optimisation of complex palpation motion trajectories is not a trivial problem. Firstly, there exists a non-linear causal relationship between palpation trajectories and sensory signals, in that the sensory signals often show considerable variations even for small differences in trajectory. This non-linear relationship stems from the complex, continuum body, soft-rigid interactions that occur when the sensorized manipulator interacts with the soft phantom. The second challenge lies with the optimization of the trajectory. Introducing multi-axial, complex trajectories increases the dimensionality of the possible action space combinatorially, so efficient search techniques are required. As such, increasingly informative sensory information can be gained with complex motions, but finding these specific, effective behaviours requires appropriate search processes. Finally, there is the challenge of finding good palpation trajectories that work for different phantoms, or patients.

In this paper, we hypothesize that Bayesian inference is an appropriate technique for identifying optimal complex palpation trajectories. Bayesian inference is investigated for two purposes. Firstly, to analyse the sensory information obtained in palpation, a Bayesian approach can provide an effective analysis of the intrinsically non-linear characteristics of

palpation. This approach can also provide an analysis which is annotated by a confidence metric to help accurate classification, a metric which would be challenging to generate with alternative approaches, such as a linear regression. Secondly, Bayesian inference is also investigated as an effective search strategy (in the form of Bayesian Exploration). Bayesian approaches can leverage on cumulative past experiences to rapidly search motion trajectory parameters, and allow for efficient search of high-dimensional action spaces. This search can enable the robot to select effective trajectories for accurate classification.

8.2 Materials and Methods

8.2.1 Phantom Development

To provide a physical test bed for the abnormality detection experiments several phantoms were created. The phantoms were created by casting Eco-Flex 00-10 Silicone and included a number of spherical 3D printed beads of different sizes and depths. The stiffness of the silicone (6.7 kPa) was chosen to be close to that of human tissue.

The casting of the phantoms included an initial casting of silicone of ‘d’ height at the bottom of a custom cuboidal mould. The initial cast layer would become the top layer of the phantom. After curing, 3D printed inclusions of various sizes are then glued into place on the phantom using a minimal amount of superglue. The beads were placed away from the edges, and with sufficient area around to provide area of the phantom free of any abnormal inclusions (no-inclusion). To complete the process, a second layer of silicone was cast, this time approximately 30mm in height. This depth was chosen so that the dynamics of the palpation experiments would not be dominated by the interaction between the phantom and the surface on which this was placed. After curing, the phantoms could then be released from the moulds. To minimize damage, and extend their usable life, a thin film of cling film was then applied to their surface.

Two training phantoms were developed to test different palpation difficulties in the identification of inclusions in human bodies, and kidneys in particular. The inclusion in *Training Phantom 1* varied from 8mm to 15mm in diameter size and were set at a depth of 10mm, a typical size and depth to tumours which may be identified on the kidney. To demonstrate how the framework developed could also be used to identify smaller, more superficial tumours, a second phantom, *Phantom 2* was created which had inclusions of size 5mm and 10mm, and at a depth of 5mm.

To demonstrate the adaptability of the framework in complex and realistic environments, an *Abdominal Phantom* was created. This was designed to include a silicon kidney organ with

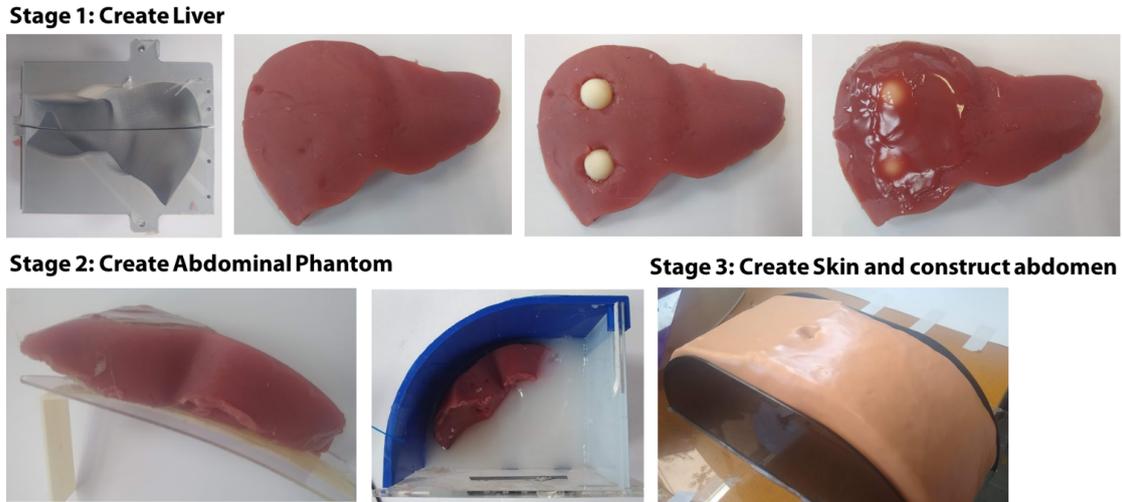


Fig. 8.2 Process for creating the *Abdominal Phantom*, showing the curing of abdomen and adding the inclusions and the overall construction of the abdomen.

inclusions and structure replicating the surrounding tissues and skin. Fig. 8.2 demonstrates the process of creating the phantom. The approach to fabrication was adapted from previous work [111]. The phantom was developed by first fabricating a liver from Eco-flex 00-10 and red dye using a 3D printed mould. Inclusions were added to this liver, and additional silicone used to seal these inclusions. A second mould was then used to cast the abdomen using Ecoflex 00-00, this was done in multiple layers to allow the liver to be positioned in the correct place and with the correct curvature. To create a skin Ecoflex 00-20 was cast into a flat sheet with approximately 1.5mm thickness which was then used to cover the phantom.

8.2.2 Sensor Technology

The choice of the sensory technology is fundamental because its morphology affects the perceptive ability of the robot to feel for abnormal inclusions [231, 232]. The selection of the sensor has been performed on the basis of the characteristics of manual palpation that have been analysed in [123]. A capacitive tactile sensor array was used for the palpation experiments [154].

The sensor technology is CySkin, adapted from the work in previous chapters. The module utilized corresponds to the one previously used in Chapter 4. (see Section 4.2.1).

A sensor reading, or tactile image, from the tactile sensor corresponds to a 7-dimensional array, where each element contains the capacitance variation value of the corresponding taxel.

8.2.3 Robot Control Experimental Set-up

For the purpose of this chapter, a palpation experiment consists of N_s seconds of contact between the sensorised robot end-effector with a target phantom. To achieve this, we manually teach the robot the location of the phantoms areas to palpate, and set the palpation starting position with the end-effector alighted normally to the surface of the phantom organ (Fig. 8.3a). The end-effector is thus driven downward until a touch event is detected by the capacitive tactile sensor at its extremity, whereby the palpation experiment begins.

The parameterization of the robot control action was designed to generate point-based exploration techniques, while maintaining freedom of wrist rotations and probing depth. The end-effector of the robot was controlled in real-time in Cartesian coordinates, acting upon the depth (z tool axis), R_x (rotation around the x tool axis), and R_y (rotation around the y tool axis) axis simultaneously. Distinct sinusoidal displacements profiles are generated for every axis, each of which is controlled by two separate parameters, thus a total of 6 parameters were used to control the robot for each palpation procedure, i.e. A_{rx} , A_{ry} , A_z , ω_{rx} , ω_{ry} and ω_z . For an arbitrary axis ‘ ax ’, a sinusoidal displacement $s_{ax}(t)$ over the course of the palpation experiment is defined as:

$$s_{ax}(t) = A_{ax} \cos(\omega_{ax}t) \quad (8.1)$$

where ‘ t ’ is the time, in seconds, elapsed since the start of the palpation experiment, and A_{ax} and ω_{ax} the axis-dependent parameters (Fig. 8.3 b,c). To achieve the displacement profiles, a UR5 robotic arm was speed controlled at 60Hz via:

$$v_{ax}(t) = \frac{\partial s_{ax}(t)}{\partial t} = A_{ax} \omega_{ax} * \cos(\omega_{ax}t) \quad (8.2)$$

where $v_{ax}(t)$ is the axis velocity in ‘ ax ’ at time ‘ t ’. Two values for each of the six parameters were explored, for a total of 64 different available palpation techniques. The Z parameters were set to drive the end-effector into the flesh of the silicon at a variable height between $2mm$ and $10mm$, thus $A_z \in [0.002, 0.01]$ and $\omega_z \in [0.5, 2]$. The parameters for the R_x and R_y rotations were set to achieve variable rotations between ± 10 degrees, thus $\omega_{rx}, \omega_{ry} \in [1, 3]$ and $A_{rx}, A_{ry} \in [-\frac{\pi}{18}, \frac{\pi}{18}]$. The rotations and depths were chosen as the maximal achievable end-effector rotations and probing depths within all phantoms and inclusions under examination, without achieving the saturation of the tactile sensor response.

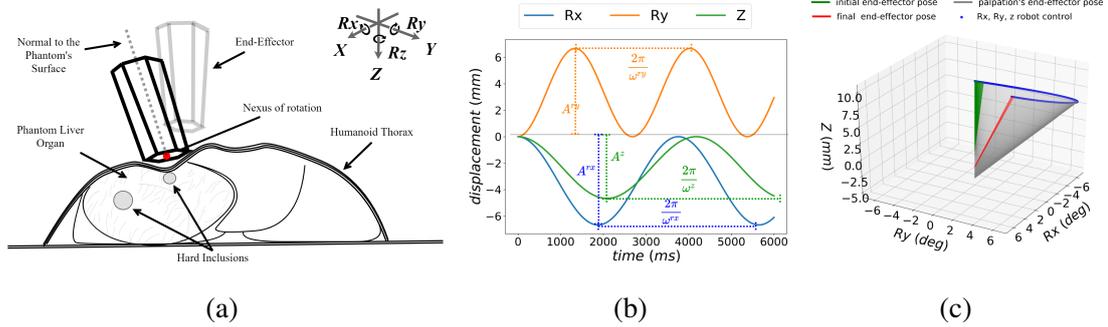


Fig. 8.3 Execution and parameterization of the robot control, and palpation trajectory. a) Diagram of the end-effector robot control during the palpation, (b) Example robot control strategy over time, given the parameters: A_{Rx} , A_{Ry} , A_z , ω_{Rx} , ω_{Ry} and ω_z , (c) Corresponding generated robot trajectory over the 3s palpation.

8.2.4 Bayesian Framework

The approach taken in this work is large scale physical experimentation, with the robot iteratively performing many palpation experiments. The framework we propose for these large-scale experiments has three key phases: palpation training, palpation inference and evaluation (8.4).

During the first phase, ‘palpation training’, the robot generates sensory data by repeatedly palpating the different inclusions with different palpation trajectories. In the second phase, ‘palpation inference’, classification processes are performed on new samples. In the final phase, ‘evaluation’, the accuracy is determined by comparing the robot classifications to the ground truth.

Training Phase

The palpation training phase involves the iterative execution of multiple robot palpation experiments. Each robot experiment involves a palpation trajectory, or action A_m where $A_m \in \mathbb{R}^6$, being performed on a specific class of inclusion (C_k) in a phantom. We define a palpation iteration as when the robot performs palpation experiments on all types of inclusions C_k once, under a specific trajectory A_m .

Data sampling

Let \mathbf{X} be an $N \times D$ dimensional vector, where each unique temporal tactile image for a probed location is a D dimensional row in the matrix. A temporal tactile image is a sequence of tactile images sampled at constant time intervals. For the purpose of the experiments

presented in this chapter the sampling period is of 500ms. By limiting the palpation to three seconds, we gain 35 pressure points over time, limiting the dimensionality of the data ($D = 35$). The value of N increases with the number of palpation experiments. In each reported experiment the value of N is initially 0 and for each ‘palpation iteration’ $N = N + K$ where K is the number of discriminative classes, or types of inclusions in the phantom to palpate (Examples of rows of \mathbf{X} can be visualized as heatmaps in Fig. 8.5 a-d).

Dimensionality Reduction

For dimensionality reduction we utilize the equations in 6.1, 6.2, 6.3, 6.4 and 6.5 to pass from \mathbf{X} to an $(N \times 1)$ matrix \mathbf{W} , and each row in the matrix is a 1-dimensional encoding of a tactile image sequence for a probed location. We will refer to w_i as a generic projected tactile sensor measurement of a probed location after palpation. The choice of a 1-dimensional subspace was to enable palpation with the minimum possible information. In this context, the robot’s action can induce the tactile information to present diversity across the 1-dimensional subspace. The methods presented in this Chapter, however, can extend to higher dimensional spaces.

Bayesian inference and PDFs update

For each palpation experiment on a specific phantom, the class of inclusion under palpation (C_k) and the type of palpation action (A_m) generate an observable sensor measurement w_i . The likelihood that a specific inclusion class $C_k \in C$ has generated the haptic observation w_i can be computed as:

$$P(C_k|w_i, A_m) = \frac{P(w_i|C_k, A_m)P(C_k)}{P(w_i, A_m)} \quad (8.3)$$

where A_m is a specific palpation action and $P(C_k)$ is the prior probability of inclusion C_k . The prior probabilities for each class can be extracted from the domain knowledge of the task to perform. For breast cancer detection, for example, the prior probability of being diagnosed with cancer can vary between 0.0017 and 0.0102 depending on risk factors [3], while for the experiments within this manuscript there is an equal probability of tissues areas with and without abnormal inclusions. To find the probability of observing w_i when performing a palpation A_m on a class C_k , instead, the central limit theorem suggests that we

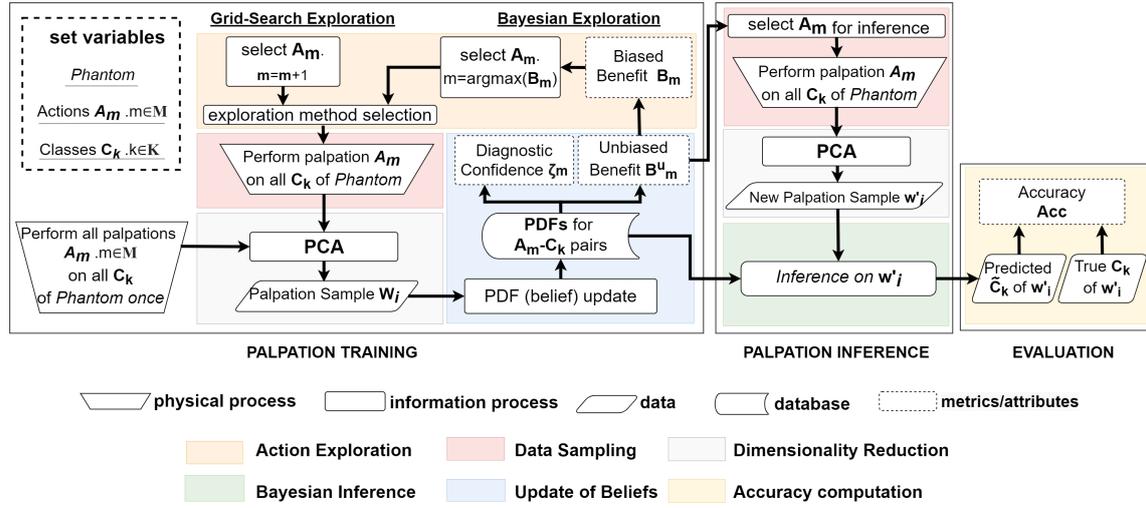


Fig. 8.4 Flowchart of experimental procedure. During the palpation training phase, the robot performs palpations A_m on different types of inclusions C_k to form PDFs. After an initial set of palpations to generate PDFs, the robot performs additional experiments to improve its classification capabilities based on the biased B_m score. In the palpation inference phase, the PDFs are used to perform inference on new samples. Moreover, an unbiased benefit and a confidence level for each palpation trajectory A_m can be estimated. In the evaluation phase the performance of the robot can be evaluated if the ground truth classification of the palpated area is known.

can approximate the findings with the probability density function $p(w_i|A_m)$ defined by a mean $\bar{\mu}_{k,m}$ and a standard deviation $\Sigma_{k,m}$ as:

$$P(w_i|C_k, A_m)P(C_k) \propto p(w_i|C_k, A_m) = \frac{1}{\sqrt{(2\pi)^2 |\Sigma_{k,m}|}} e^{(-\frac{1}{2}(\bar{w}_i - \bar{\mu}_{k,m})^T \Sigma_{k,m}^{-1} (\bar{w}_i - \bar{\mu}_{k,m}))} \quad (8.4)$$

The evidence in the projected matrix \mathbf{W} can thus be used to update the $\bar{\mu}_{k,m}$ and $\Sigma_{k,m}$ of each class of inclusion under every exploratory palpation action, improving the estimates with evidence from every new palpation iteration.

From this training data we can generate two key metrics which help assess the quality of a palpation trajectory. The first metric is the ‘Unbiased Benefit Estimator’ $\hat{\mathbf{B}}_m$, which provides a measure of how useful the sensor data from a given trajectory is in performing classification of abnormalities. In general, a good palpation trajectory leads to significantly different sensor responses for different classes of inclusion. This can be quantitatively measured by considering the overlap between the PDFs for the different classes of inclusions. A possible

measure of this overlap is the Bhattacharyya coefficient [203]. The Bhattacharyya coefficient between two probability density functions p and q is defined as:

$$BCoeff = \int \sqrt{p(x)q(x)}dx \quad (8.5)$$

We use the Bhattacharyya coefficient to compute a confusion probability matrix $\Psi_{ks,m}$ for each possible exploratory palpation $A_m, m \in M$. Each element in $\Psi_{ks,m}$ is a mutual confusion between any two classes C_k and C_s , which under the assumption of normal distributions can be computed as:

$$\Psi_{ks,m} = \sqrt{\frac{2\vec{\sigma}_{k,m}^2 \vec{\sigma}_{s,m}^2}{\vec{\sigma}_{k,m}^2 + \vec{\sigma}_{s,m}^2}} e^{-\frac{(\bar{\mu}_{k,m} - \bar{\mu}_{s,m})^2}{4\vec{\sigma}_{k,m}^2 + 4\vec{\sigma}_{s,m}^2}} \quad (8.6)$$

where $\vec{\sigma}_{k,m,h}$ is the diagonal vector of $\Sigma_{k,m,h}$. The Ψ probability confusion matrix can be used to find the benefit of making an exploratory action A_m . Based on this coefficient, we can define a new measure of how well an action separates the PDFs for different inclusions, and therefore how ‘beneficial’ it is for making a classification of abnormalities. This measure, called unbiased benefit estimator ($\hat{\mathbf{B}}_m$), can be computed for a specific action A_m as:

$$\hat{\mathbf{B}}_m = \sum_k^K \frac{P(C_k)^2}{\sum_s^K \Psi_{ks,m} P(C_s)} \quad (8.7)$$

and its value will be higher for palpation actions with class probability density functions with least overlap. Ψ_m is the confusion matrix for action A_m , but $\Psi_{ks,m}$ is a number, corresponding specifically to the confusion under action A_m of class C_k and C_s . In this context, $\hat{\mathbf{B}}_m$ is a very important metrics as it can be considered as the benefit of making an exploratory action A_m . One of the unique advantages of calculating PDFs based on Eq. 8.4 is that we can also obtain a probabilistically measure of confidence of the tactile sensor data for a specific palpation trajectory A_m . Using the PDFs, we can compute a measure of confidence based on $\hat{\mathbf{B}}_m$ which indicates the amount of confusion the robot has for each palpation trajectory. This measure of confidence, ζ_m , is defined as:

$$\zeta_m = 1 - \frac{1}{2 + e^{-\hat{\mathbf{B}}_m}} \quad (8.8)$$

This metric has been devised as to increases monotonically when the discriminatory confusion reduces, providing a value between 0 and 1. Higher values of ζ_m signify a higher confidence of classification for all classes of inclusion under a specific trajectory. The confidence is a fundamental metric used in this experimental framework as actions with a higher confidence

levels signify that the PDFs for different classes of inclusion can be better separated, and hence more accurate classification can be made. These two metrics ($\widehat{\mathbf{B}}_m$ and ζ_m) can be used to assist in selecting and choosing the appropriate palpation trajectories for discrimination of abnormal inclusions in different scenarios.

Exploratory Action Identification

As the number of parameters which describe a palpation trajectory increases, there is an exponential increase in the number of possible actions that must be searched.

We propose using Bayesian Exploration, an approach that leverages on past experiences to improve the search of the palpation trajectory parameters. Bayesian Exploration was first used in [64] for tactile discrimination of textures, and was later argued to be meaningful from a neuroscientific perspective [145]. Bayesian Exploration has since been shown to aid in several other tactile discrimination tasks [232, 274]. Fig. 8.4 illustrates how Bayesian Exploration can be implemented to find optimal palpation actions. Bayesian Exploration provides a method of iteratively selecting, and exploring, the most “promising” action trajectories in a probabilistic, evidence based, manner [64].

This search requires a metric to guide the selection of actions. For this we can use the $\widehat{\mathbf{B}}_m$ score, which captures how favourable an action is by identifying how it minimizes the discriminatory confusion between classes of inclusion. However, for exploration, it is important to also bias this measure by factor which encourages exploration of new, unexplored, trajectories. Thus, we create a biased benefit score based on $\widehat{\mathbf{B}}_m$ as follows:

$$\mathbf{B}_m = 1 - (1 - \widehat{\mathbf{B}}_m)^{\frac{1}{n_m}} \quad (8.9)$$

Where n_m is the number of times action m was performed iteratively during the palpation experiments. The biased benefits are discounted by the number of times the action has already been performed during action exploration, to discourage excessive exploitation and eventually encourage the explorative update of belief states under less exploited actions.

At the start, the robot is made to palpate each class of inclusion under every action once, to gather initial experimental evidence. For each new iteration the palpating robot is then made to palpate each class of inclusion under a specific action once more. The action to use is chosen based on the highest scoring action under the \mathbf{B}_m , and the belief states are updated accordingly.

Bayesian Inference Phase

In the second phase of the framework, Bayesian inference, the robot performs the classification of abnormal inclusions, identifying the class of unseen sensor data obtained through additional robotic palpations. This classification is made via Bayesian Inference, using the PDFs generated in the palpation training phase through several palpation iterations. To perform inference on a new tactile sample w'_i , we evaluate the sample at $p(w'_i|C_k, A_m)$, under every C_k for a chosen action A_m . The C_k of the PDF yielding the highest value will be inferred as a class for w'_i .

$$\tilde{C}_k = \operatorname{argmax}\{p(w'_i|C_k, A_m) : k \in C\} \quad (8.10)$$

where \tilde{C}_k is the class estimated for C_k . This inference process is used throughout the results section to test the abilities of different palpation trajectories, and will be referred to as ‘Bayesian inference classification’.

Evaluation Phase

In the third and final phase of the framework, evaluation, we evaluate the performance of the classification made during the palpation training phase. This is achieved by comparing the ‘true’ class of inclusion C_k for the unseen sample, against its inferred class during phase two. Over several iterations, we can count the number of correctly classified abnormal inclusions as True Positives (TP), and the number of correctly classified inclusion-free areas of the phantom as True Negatives (TN). For a total of N_C classifications, or palpation inferences, the accuracy can be formally computed as:

$$Acc = \frac{TP + TN}{N_C} \quad (8.11)$$

8.3 Results

During the process described in the previous sections, an ‘a posteriori’ probability density $p(w_i|C_k, A_m)$ of sensor evidence w_i , given an inclusion C_k and an action A_m is formed. Under each palpation technique, the overlap between the probability densities belonging to different types of inclusions is indicative of the amount of discriminative confusion the robot has, ‘a posteriori’, under that palpation technique.

As a consequence of the probabilistic treatment of the sensor response, it is possible to rank the robot’s palpation techniques, favouring those actions which minimize the overlap of different distributions, and thus the discriminative confusion of the robotic agent when

Phantom	Class of Inclusions	Depth of Inclusions	Number of Inclusions	Complexity Demonstration	Exploration Experiments	Validation Experiments	Total Number
Training Phantom 1	10mm	10mm	3	1920	3840		
	15mm	10mm	3	(10 iterations per trajectory)	(20 iterations per trajectory)	/	5760
	no incl.	N/A	4				
Training Phantom 2	5mm	5mm	3		3840		
	7mm	5mm	3	/	(20 iterations per trajectory)	/	3840
	Healthy	N/A	4				
Abdominal Test Phantom	15mm	10mm	2		800		
	no incl.	N/A	2	/	(20 iterations per trajectory)		800

Table 8.1 Experimental breakdown of robotic palpations and palpated phantoms. The values in parenthesis represent the number of samples gathered by the robot for any trajectory (A_m) and class of inclusion (C_k) pair

	Complexity Demonstration	Exploration Experiments	Validation Experiments
Number of trajectories attempted	64	64	20
Parameter combinations	$A_{rx} \in [0, \frac{\pi}{18}]$ $A_{ry} \in [0, -\frac{\pi}{18}]$ $A_z \in [0, 0.01]$ $\omega_{rx} \in [0, 1]$ $\omega_{ry} \in [0, 1]$ $\omega_z \in [0, 0.5]$	$A_{rx} \in [-\frac{\pi}{18}, \frac{\pi}{18}]$ $A_{ry} \in [-\frac{\pi}{18}, \frac{\pi}{18}]$ $A_z \in [0.002, 0.01]$ $\omega_{rx} \in [1, 3]$ $\omega_{ry} \in [1, 3]$ $\omega_z \in [0.5, 2]$	20 Highest score trajectories from " Exploration Experiments "

Table 8.2 Experimental breakdown of robotic palpation trajectories and parameters over experiments.

making a classification. One such measure is the Bhattacharyya coefficient, which we use to rank each of the possible robot actions based on the amount of discriminative confusion observed after performing the palpation procedure.

Through Ψ_m a biased and unbiased benefit estimations are defined. The unbiased benefit estimation will have a higher value for actions with class probability density functions with least overlap. The biased benefit, instead, discounts the unbiased estimation by the number of times the action has already been performed during action exploration, to discourage excessive exploitation and eventually encourage the explorative update of belief states under less exploited actions.

8.3.1 Exploring Action Complexity in Robot Medical Palpation

The first set of experiments investigates the influence of the palpation trajectory on the robot's ability to distinguish different classes of inclusion. We investigate how the classification

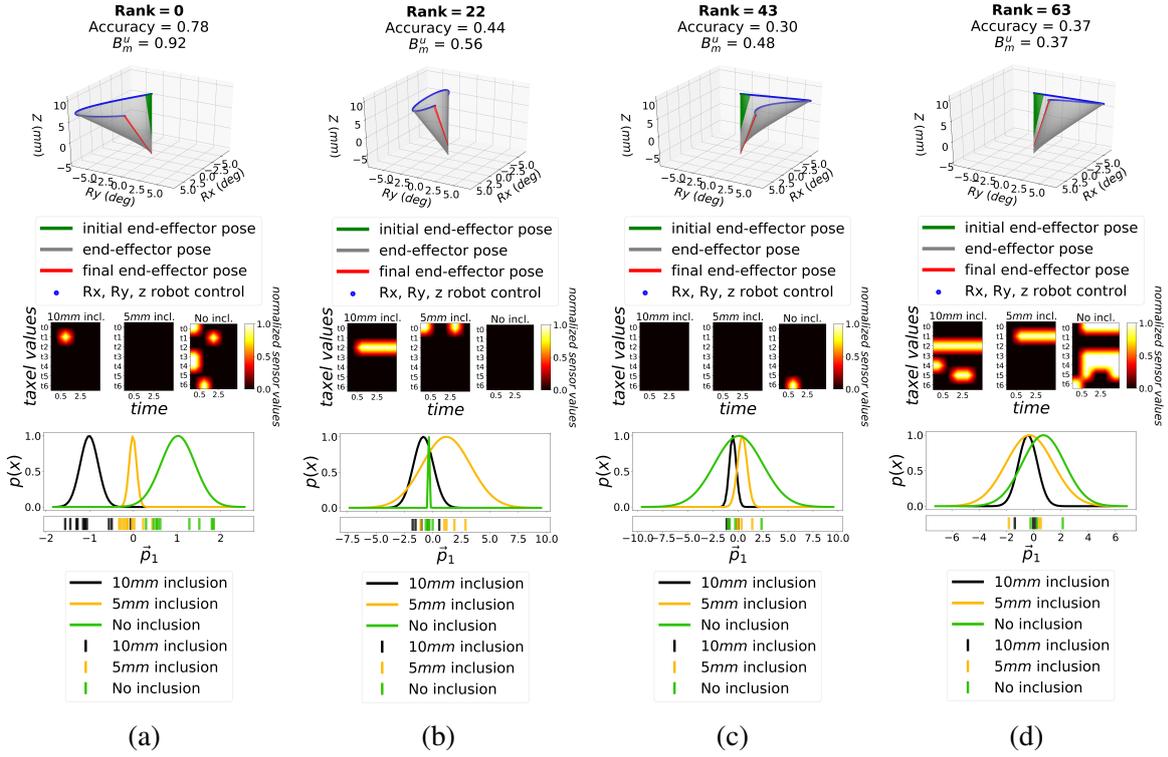


Fig. 8.5 Influence of the palpation trajectory to the PDFs in Training Phantom 2. The 64 robot actions are ranked by $\hat{\mathbf{B}}_m$, and Rank 0-22-43-63 (best to worst) are shown respectively in figures (a), (b), (c) and (d). The top plots show the robot trajectory generated during palpation, where R_x , R_y are rotations about the x-y axis and Z is the probing depth. The middle plots show the normalized spatio-temporal tactile images generated during palpation. The brightness corresponds to the normalised taxel values at specific time intervals (proportional to pressure). The bottom plots show the PDFs generated from 60% of the data detailed in the “Exploration Experiments” of Table 8.2. Here, p_1 is the first principal component onto which the original sensor data were projected.

capabilities of different palpation trajectories depend on the separation of the PDFs in sensor space for different classes of inclusion.

In these experiments we examine 64 different palpation trajectories, and analyse how they influence the separation of PDFs. The 64 palpation trajectories are generated through the combination of 6 parameters which describe a trajectory. We conducted these experiment on *Training Phantom 1*, performing all the palpation actions on all the different inclusion types. For each type of inclusion and palpation trajectory, we perform the palpation 20 times, as shown in the “Exploration Experiments” column of Table 8.2 and Table 8.1.

Fig. 8.5 shows the PDFs for four exemplar palpation trajectories, ordered with respect to the $\hat{\mathbf{B}}_m$ scores. The PDFs are created via Eq. 8.4, and the unbiased benefit score $\hat{\mathbf{B}}_m$ is

computed for each palpation action using Eq. 8.7. The different motion parameters result in different palpation trajectories with very diverse PDFs. Moreover, as shown in Fig. 8.5, the raw tactile sensor data for each class is influenced by the palpation strategy itself. Ideally, PDFs for different inclusion classes should have minimal overlap for discrimination purposes. The results shows how it is possible to have motion parameters, and hence trajectories, that give rise to PDFs which are fully separated across the PCA principal component p_1 for all classes of inclusion (Fig. 8.5a). Conversely, some motion parameters lead to PDFs which are heavily overlapped (Fig. 8.5c & d).

The figure also shows that the degree of these overlaps can be represented by using the $\widehat{\mathbf{B}}_m$ scores. The trajectories with less overlap (Fig. 8.5a & b) results in higher $\widehat{\mathbf{B}}_m$ scores whereas those with more overlap (Fig. 8.5c & d) have a far lower score. As such, the score represents the discriminative performance of the palpation trajectory used by the robot and can be used to compute a ranking for the different trajectories. It is also important to note that $\widehat{\mathbf{B}}_m$, and thus the resulting ranking order, also indicates the degrees to which each PDF is separated from the others, in addition to measuring the amount of overlap. For example, although the PDFs of the 43rd and 63rd ranked actions are similarly overlapped, the former is ranked higher because of the lower overlap between the PDFs of the 10mm and 5mm inclusion classes.

In the next experiment, in contrast to the previous one, we compare the separation of the PDFs for the same trajectories but across different phantoms, i.e. *Phantom 1*, *Phantom 2*, and the *Abdominal Phantom*. We perform this experiment to assess whether a palpation trajectory optimized for one phantom can perform well on other phantoms.

To achieve this, we identified the best trajectories for *Phantom 1*, *Phantom 2* and the *Abdominal Phantom*, by finding the trajectory with the highest $\widehat{\mathbf{B}}_m$ score for each phantom. These three top-ranking trajectories, together with the resulting PDFs, are compared in Fig. 8.6a. The experimental data used is detailed in the ‘‘Exploration Experiments’’ of Table 8.2 and Table 8.1. The first observation is that the optimum trajectories are significantly different for the different phantoms. The best palpation trajectory for *Phantom 1* is a counter clockwise rotation in an almost horizontal plane, whereas that for *Phantom 2* is a clockwise trajectory with a similar amplitude. The optimum trajectory for *Abdominal Phantom* is significantly different, with a clockwise rotation occurring with smaller amplitude, and a higher palpation depth. The second observation that can be made considers the PDF overlaps. Fig. 8.6a shows that the highest ranked actions do not show high separation of the PDFs on the other phantoms. The best trajectory for *Phantom 1*, for example, does not perform well in *Phantom 2*, with the action resulting in high overlaps of PDFs belonging to different classes of inclusions. The *Abdominal Phantom* is a relatively easier task, in comparison

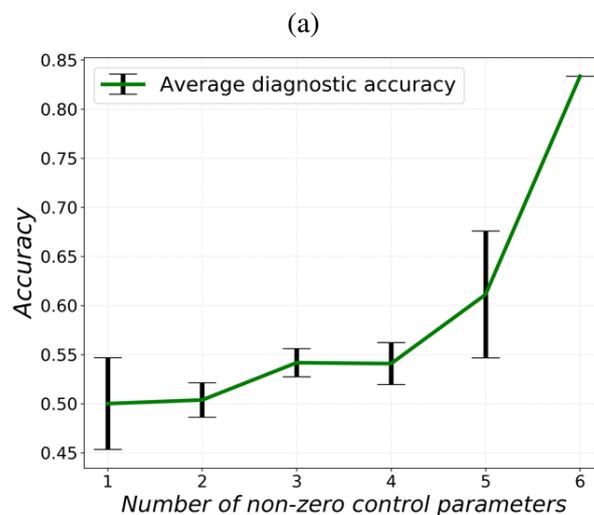
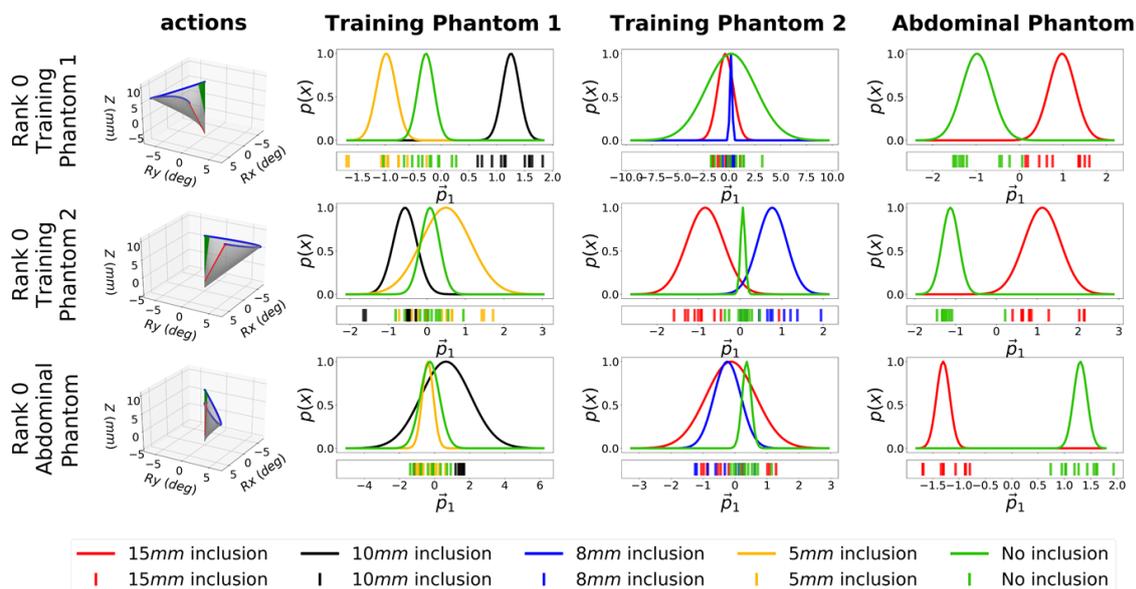


Fig. 8.6 The figure shows the complexity of robotics palpation. The diagonal plots in (a) show the PDFs of the best performing palpation trajectories for each phantom, while the off-diagonal plots show PDFs of the same trajectories in all other phantoms. (b) shows the accuracy of a Bayesian Inference classifier trained on sensor data generated via palpation trajectories with a varying number of parameters.

to *Phantom 1* and 2, with all of three palpation strategies achieving high separation of the PDFs. However, the trajectory ranking higher for *Abdominal Phantom* still achieves higher separation of PDFs in the same phantom, while it does not perform well in *Phantom 1* and *Phantom 2*.

In the next set of experiments, we examine the necessity of complex trajectories for more accurate palpation. This is achieved by comparing palpation trajectories which are described by different number of control parameters. As each axis of motion is controlled by a specific pair of parameters (i.e. $A_{rx}-\omega_{rx}$, $A_{ry}-\omega_{ry}$ and $A_z-\omega_z$ to control R_x , R_y and Z respectively), reducing the number of parameters decreases the number of DoF, and hence the complexity of the palpation trajectory.

To systematically vary and reduce the complexity, every possible combination of the 6 parameters is set to zero in turn. As such, each palpation strategy can have between 1 to 6 non-zero parameters. The 64 palpation trajectories defined by these parameters are performed 10 times on all types of inclusions in Training *Phantom 1*, and the corresponding tactile data is stored for evaluation. This data corresponds to tactile information from 2160 palpations; the parameter values explored and the number of experiments are detailed in the ‘‘Complexity Demonstration’’ columns of Table 8.2 and Table 8.1.

To evaluate the performance of each set of motion parameters, Bayesian Inference classification is performed on the computed PDFs as previously described. The classification inference is performed on each palpation trajectory separately, with 60% of the sampled palpations used for training, and the remaining 40% used for testing. Fig. 8.6b shows the average performance of the classifier for across all palpation trajectories with different numbers of active parameters. As illustrated in Fig. 8.6b, trajectories with one or two describing parameters achieve accuracy rates of 50% on average, thus little above chance (33%). With the full employment of the 6 descriptive parameters, the generated trajectories can achieve accuracies above 60%. As shown in Fig. 8.6b, when the dimensionality of the actions, and hence number of motion parameters, is increased, there is up to 35% improvement in the average classification accuracy of the robot. This justifies and demonstrates the need for complex trajectories when performing palpation.

From this first set of experiments we can make several conclusions, which allow us to better understand the role of the palpation trajectories to perform better soft tissue palpation. Firstly, the correspondence between the palpation trajectory and discriminatory performance is complex and non-linear. Slight changes in the palpation trajectory can significantly affect the discriminatory abilities of the robot. Secondly, the optimum trajectories vary from phantom to phantom. There is not one ‘optimum’ motion for all phantoms. Thirdly, introducing more complex palpation trajectories allows for better action profiles to emerge,

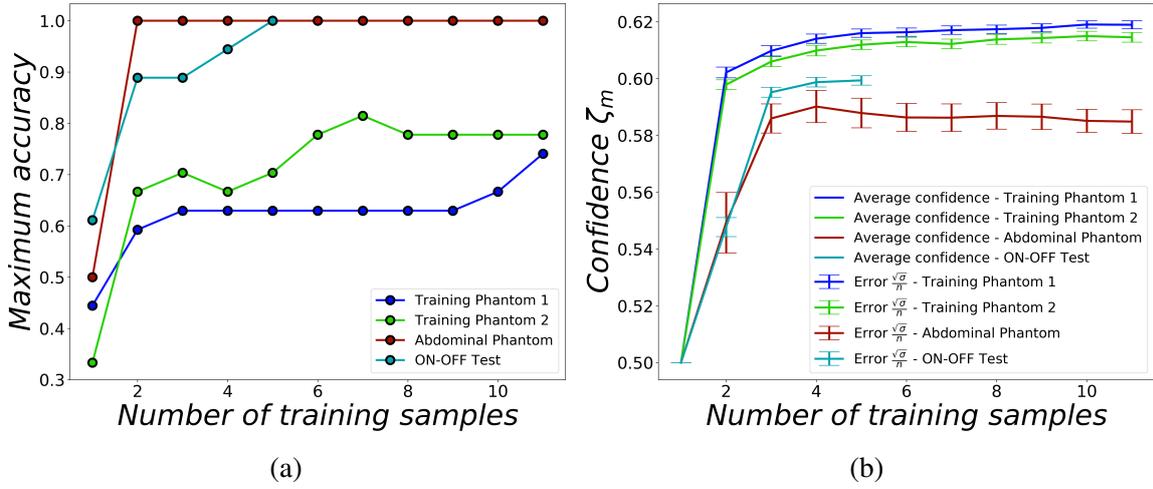


Fig. 8.7 The figure shows the performance of a Bayesian inference classifier within the framework developed. (a) shows the relationship between the maximum classifier accuracy and the number of samples gathered for each palpation trajectory-class pair. (b) shows the relationship between the developed confidence level and the number of samples gathered for each palpation trajectory-class pair. The vertical bars in the plot illustrate the errors of the confidence at that point

demonstrating that increasingly complex actions increase the ability to make more accurate classification of abnormalities.

8.3.2 Bayesian Approaches for Confident Abnormality Detection

The next set of experiments examines the levels of confidence (ζ_m) and the experimental accuracy (Acc) when computing the PDFs based on a different number of training samples.

In these experiments the same dataset from the previous experiments is used, where palpation training was performed on each class of inclusion, using each action 20 times (see “Exploration Experiment” columns of Table 8.2 and Table 8.1.

For every trajectory, we consider the data originated from palpating each class of inclusion 12 times. This data is used systematically, i.e. the first computation of the PDFs is based on one sample of tactile sensor data per class-action pair, while the last computation of the PDFs is based on 12 samples per class-action pair. Every time the PDFs are computed, we also compute the benefit and the confidence as previously described. These steps are performed on every trajectory available, and the resulting PDFs are used to compute the accuracy, as described in the palpation inference and evaluation phase. The robot is tested on 40% of held out sensor data, which corresponds to 8 samples for each class-action pair.

Fig. 8.7a shows the highest accuracy of all palpation trajectories, as a function of the number of training samples used to compute the PDFs. As the number of training samples increases, the evidence used to build the PDFs increases, leading to the classifier performing more accurate classification.

In Fig. 8.7b, the confidence metric is also plotted as a function of the number of samples used for training. As expected, we can see that the confidence metric increases with the number of training samples. The confidence, however, saturates at different values for each phantom. These values indicate how ‘reliable’ is the classification of the robot under a specific trajectory. This measure will first and foremost depend on the overlap of the PDFs, which will in turn depend on the similarity of the tactile sensor data for different classes of inclusions. In Fig. 8.7b, the robot achieves highest confidence for *Phantom 1*, followed by *Phantom 2* and the *Abdominal Phantom* in order. *Phantom 1* contains relatively large inclusions (15mm and 10mm), while *Phantom 2* contains smaller, and more similar, inclusions (7mm and 5mm).

From this we can hypothesise that the sensor data originating from *Phantom 1* is generally more diverse across the different classes of inclusions, than it is in *Phantom 2*. The *Abdominal Phantom* reaches a lower level of confidence with respect to the two training phantoms. We hypothesise that this is due to the higher levels of structural complexity of this phantom, with a curved surface, separate skin layer and liver organ. Despite this lower confidence in classification, as shown in Fig. 8.7a, the robot’s accuracy remains high.

8.3.3 Online, Rapid Bayesian Exploration: Pruning the Search Space

As shown in the first set of experiments, good palpation trajectories lead to separable PDFs for different classes of inclusion, enabling confident and accurate classification. However, we already identified that firstly, there is a significant benefit to using trajectories which are described by a higher number of parameters (more DoF), and secondly, the optimal palpation trajectories vary for each phantom. Therefore, the next set of experiments focuses on efficiently finding these high dimensional optimal palpation trajectories for different phantoms.

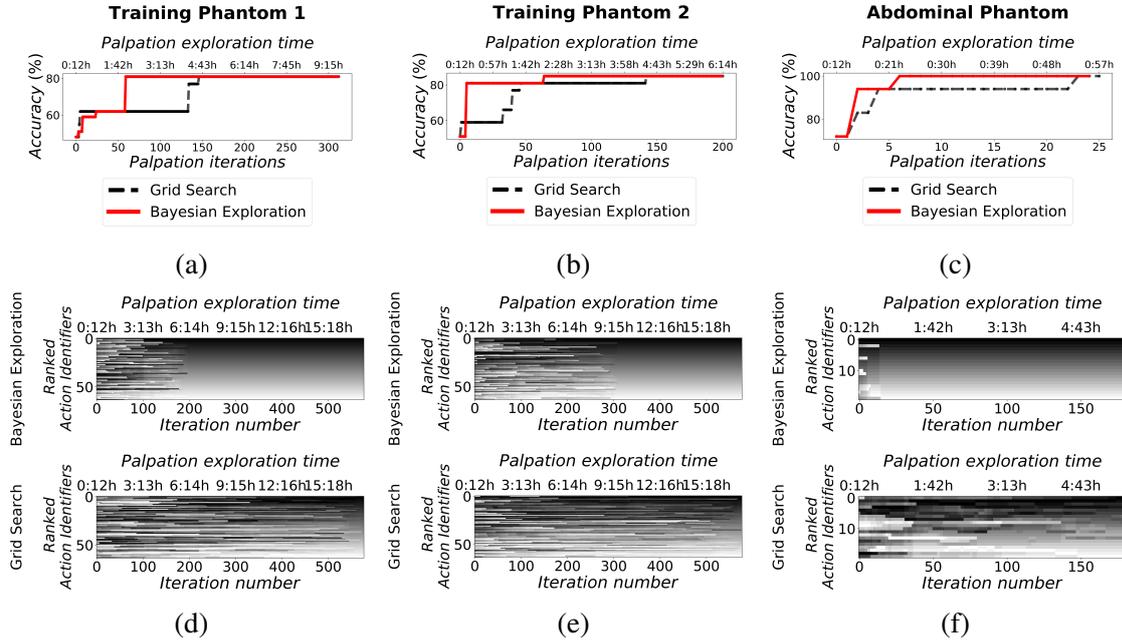


Fig. 8.8 Training comparison between Bayesian online exploration and Grid Search. Figures (a), (b) and (c) show the maximal accuracy achieved at each training iteration by both Grid Search and Bayesian Exploration, on a left-out validation set of 8 palpation samples for every class of inclusion. Figures (d), (e) and (f) show the trajectories on the y-axis, ordered based on their final $\hat{\mathbf{B}}_m$ rank during training both for Bayesian Exploration and Grid Search. The color intensity for each action indicates its rank at each iteration. Bayesian Exploration achieves a final ranking faster, bringing the robot’s trajectories to a ranked solution in at least half the time than when training through a systematic action space search.

In the next set of experiments we compare the Bayesian Exploration approach with a benchmark, a grid-based search method. For this comparison we perform palpation on all phantoms. The data utilized for these experiments is detailed in the “Exploration Experiment” columns of Table 8.2 and Table 8.1. When using this Bayesian Exploration framework, the robot first palpates each class of inclusion once with every possible action. This is unavoidable, as it allows a PDF to be generated for each inclusion-action pair. This initial gathering of evidence allows the Bayesian Exploration process to then start (Fig. 8.4).

In this exploration process, all palpation trajectories are ranked using the biased B_m score. The action which has the highest B_m score is then used to palpate each class of inclusion once, and the PDFs are updated with the new tactile information. This corresponds to one iteration of the Bayesian Exploration framework. The B_m score is then computed again and used to select the next palpation trajectory to test, with the steps then iteratively repeated. To evaluate each iteration of the exploration process, we take the top scoring action at that time, as defined by the unbiased benefit score, and use this action perform Bayesian Inference. The

inference is performed on 40% of unseen data from ‘Exploration Experiments’, and provides the robot with the ‘best accuracy’ for every iteration of the exploration process. Importantly, the top scoring action is selected by the unbiased benefit score, as we want to find the top performing action which is purely based on the ability to separate the PDFs in sensor space. As a benchmark, the results from the grid-search method is also presented. During grid-search, contrarily to Bayesian Exploration, the action is selected based on a breadth-first parametric search, with the rest of the experiments performed in the same manner.

It is important to note that although under Bayesian Exploration it is first necessary to initially palpate every inclusion via each action available, our hypothesis is that by using Bayesian Exploration the robot will ultimately need less training (and therefore palpations) to reach a good perceptual understanding of each inclusion under touch. The initial evidence, in fact, should drive further collection of tactile information prioritizing only the actions which lead to less confusionary states.

We compare the performance of these methods by considering the number of ‘palpation iterations’ necessary to train the robot. As previously described, a ‘palpation iteration’ involves the palpation of all classes of inclusions C_k under a specific action A_m . The action A_m is here iteratively selected through Bayesian Exploration or Grid-Search. As shown in Fig. 8.8 a-c, Bayesian Exploration achieved its highest performance after around only 60 iterations in both training phantoms. On the *Abdominal Phantom* this took approximately 7 iterations. Conversely, grid based systematic search performed poorly, finding equally good palpation strategies after 150 palpation iterations on the training phantoms, and 23 iterations on the *Abdominal Phantom*.

In Fig. 8.8 d-f, the intensity of the color shows the final ranking of the actions based on the figure shows how this ranking is ‘unstable’ for grid search, i.e. the ranking keeps changing throughout the experiments, before reaching the final rank. Bayesian Exploration, however, induces a stable ranking much sooner, where the final ranking of trajectories is found much earlier on in the experiments.

By applying Bayesian exploration, and leveraging the ranking provided by the score, the actions which best separate the PDFs across different classes of inclusions are preferentially explored. By using this exploration technique, the robot can efficiently search a high dimensional parameter space. This complex high dimensional action space has previously been demonstrated to be necessary for accurate classification of abnormal inclusions in soft tissues. From these results we can observe that by using Bayesian exploration, the time taken to find the optimal strategy is halved in comparison to a systematic grid search.

Finally, after performing Bayesian Exploration, we can report the final accuracy of the entire framework across all palpated phantoms. As previously explained, this is computed as

ACCURACY % highest (average)	15mm vs NA	10mm vs NA	8mm vs NA	5mm vs NA	15mm vs 8mm vs NA	10mm vs 5mm vs NA
Training Phantom 1	0.944 (0.462)		0.778 (0.394)		0.740 (0.412)	
Training Phantom 2		0.889 (0.438)		0.833 (0.387)		0.778 (0.423)
Abdominal phantom	1.0 (0.85)					

Table 8.3 Highest and average classification accuracies achieved by the palpation system when training the Bayesian classifier on 14 samples of each class of inclusion and testing on 6 unseen samples.

the accuracy achieved on 40% of unseen palpation samples from each phantom. Table 8.3 reports the final highest test accuracy observed after training. Since the hypothesis in this paper hinges on the postulate that appropriate palpation trajectories can aid in abnormality detection via palpation, we also report the average accuracy across all attempted palpation trajectories in Table 8.3.

Notably the system is capable of achieving over 80% accuracy when discriminating between 5mm inclusions and no inclusions. On the *Abdominal Phantom*, the robot achieves 100% accuracy when discriminating between 15mm inclusions and no inclusions. Moreover, the highest performing motion strategies outperform the average performance of any one action by approximately a factor of two in almost all scenarios, confirming and emphasizing the need for appropriate palpation trajectories during abnormality detection.

8.4 Discussion

Medical palpation is an impactful preliminary diagnosis tool which is performed widely by primary care physicians. However, it is considered extremely challenging for a robot to perform due to the complexity of the interactions. The interactions between the palpation device and the human body are non-linear; the complexity of the action space and the interactions is significant, and potentially infinite; and the solutions are different for every ‘patient’. Thus, to gain a more insightful understanding of this problem we need to go beyond typical robotic approaches, including modelling and optimization. In this work, we perform large scale physical experiments to understand how increasingly complex palpation trajectories affect the sensory information that is gained, and thus the ability to make accurate classification of abnormal inclusions. The framework presented in this work (Fig. 8.4) allows for the fast exploration of a high dimensional action space, which arises from the palpation of soft bodies. The framework identifies palpation strategies that allow for a confident classification of the presence, or absence, of abnormal inclusions. The identified palpation

strategies have been shown to enable the confident detection of abnormal inclusions as small as 5mm in diameter (Table 8.3).

Through this experimental approach to palpation we have furthered our understanding of palpation. We have identified that increasing the complexity of the palpation trajectory improves the robot's classification capabilities. In addition, we have shown that slight changes in the trajectory, or the patient, significantly affect the performance. This demonstrates that the optimum palpation trajectory must be found or identified for each patient through physical experimentation; this mirrors the method in which human practitioners find the best palpation motion for each patient. To make intelligent decisions in this non-linear highly complex space, we have demonstrated how a probabilistic Bayesian approach allows for accurate and efficient search and decision making. This gives us insights into how to process and make decisions based on the sensory evidence. In the light of these experiments, one remaining question lies with the parameterization of the palpation trajectories. The explored parameters, in fact, are still based on human design and intuition, and as such, they are limited. In future scenarios the parameterization and trajectories would ideally emerge from the haptic interaction with the soft tissue itself.

The coupling between motion and sensory data goes beyond simple optimisation of the sensory data through motion optimization. The motion also affects the dimensionality reduction of the sensory data obtained providing additional benefits and improving the ability of the system to build a strong belief knowledge base. This approach starts to explore beyond the optimisation of sensory data, but rather the augmentation of the robot's understanding of its sensory perception, and the role it plays in exploring and building beliefs about the world. This is an exciting research direction and should be investigated further.

Going forwards, this knowledge is important in several ways. In the long-term, we can use the methods to develop 'robot doctors' who can perform accurate and confident diagnosis. The framework development provides a starting-point for the experiment procedure for such a robot. However, to achieve this it is necessary to find appropriate ways to perform knowledge transfer across patients or phantoms. In the short term, we can use this understanding of the process to help with the training and teaching of doctors. We can also apply the methods and approaches to other similar problems, where the Bayesian treatment and large-scale physical experiments would further our understanding of the problem at hand.

In this work we focused on problems which are typically 'challenging' for human doctors and push our sense of touch to the limit. Developing a robotic system which could outperform humans in this diagnosis and enable identification of inclusions at an earlier stage, would be highly impactful. Robotic technologies have the advantages of high sensitivity, repeatability, precision, and perfect memory and recall. Using our increased understanding of palpation, it

may be possible to further explore and optimize motion strategies, and intelligently learn and reason to push the boundaries on what can be detected.

The role of Bayesian Exploration, in this context, goes beyond simple exploratory action decision-making, but hinges on the perceptual understanding, by the robot practitioner, of the experienced tissue abnormalities, and the exploitation of said experience to improve exploration and decision-making. Going forward, we believe it is necessary to find appropriate ways to perform knowledge transfer, which restore some of the robot's beliefs about abnormalities from training to more complex testing phantoms, and that go beyond the transfer of learned palpation techniques. Exploring techniques for adapting and updating the beliefs for a new patient or phantom without requiring any additional training would increase the potential applicability of the work.

In this chapter we show how complexity in the action space can be beneficial for the conditioning of the soft interactions arising from the robot and its environment. This is an important part of the sensory-motor conditioning in *SoMComp*. This chapter has also shown how it is possible to bypass some of the limitations imposed in the previous chapters. More specifically, how to efficiently search the action space for high dimensional action parametrization, and how to achieve complex trajectories through the combination of only 6 parameters of motion. The discretization of the parameters to perform search, however, is still an issue. The next chapter introduces a method to take information conditioning through actions one step forward, and how to employ complexity in the actions to achieve human-like interactions with a piano instrument.

Chapter 9

Action Complexity: Expressive Piano Playing

In the last two chapters the topics of tactile perception in the context of sensory-motor coordination were treated, the action was purposefully devised to augment the perceptive capability of the robot, and perform accurate tactile diagnosis. The sense of touch is extremely suited to show how morphology and action can influence sensory perception, since it is characterized by the physical interactions which give rise to the stimuli themselves. However, the *sensing* treated in *SoMComp* is not limited to tactile sensing, instead, it refers to any sensing modality which captures information which the robot can directly influence. In this Chapter we choose to show results over a different sensing modality, such as hearing. We

Reference Publication

This chapter was adapted from a Journal article titled “**Gaussian Process Inference Modelling of Dynamic Robot Control for Expressive Piano Playing**” [226], published in the journal PlosOne in 2020. The article was written in collaboration with Dr Fumiya Iida and Ms Cheryn Ng, and proposes a mathematical framework to use jointly optimize both the morphology and the action of a robot to improve on discrimination tasks. Dr Fumiya Iida contributed with the conceptualization of the topics, design of the experiments and the writing of the article. Ms Cheryn Ng is the co-first author of the paper, and contributed with the robot control, experiment execution, data analysis and paper writing. As co-first author in the article, my contribution includes conceptualization of the topics, the design and execution of the experiments, the robot control, the formulation of the mathematics for the framework, data analysis and the writing of the article.

move away from the classification scenarios shown until now, towards a regression-learning problem instead, where a robot learns how to perform key-strokes on piano keys to approximate the playing of notes according to different playing styles. Like previously, large-scale physical experimentation is used to show how a robot can find complex actions which appropriately condition the soft interactions between the environment and the robot.

9.1 Introduction

Since the dawn of robotics there has been an interest in making machines perform artistic and creative tasks in a human-like manner [139]. Music instrument playing, in particular, is an important challenge, because the skills necessary to play music from physical instruments are often beyond state-of-the-art robotics technologies, while the attempt to mimic musicians would give us insights into the artistic nature of humans [242].

Piano, among others, is a complex instrument with rich and complex acoustics, which is difficult to master even for humans after many years of training. The production of rich expressive sounds requires appropriate key-press trajectories with a suitable mechanical apparatus. A key-press event, as performed through a finger by a human musician, can therefore not be seen from the point of view of the finger, or the instrument, in isolation. Rather the action of the finger and the instrument is coupled, where the dynamics of the piano are linked to the bio-mechanics and neuromuscular dynamics of the pianist, and their coupling produces rich and complex acoustic energy radiating from the soundboard [74, 94].

Previous attempts to reproduce piano playing by robots mainly focused on two aspects: the mechanical actuation of the fingers and the algorithms for finger motion planning across keys. A large variety of actuation mechanisms was proposed by using DC motors [117], servomotors [260, 281], pneumatic cylinders [109, 141], and tubular solenoids [61]. These actuation mechanisms were then integrated to various control and planning architectures, such as hard-coded motion paths [281], optimal path planning algorithms [61, 117, 136, 260], and more advanced algorithms including collision avoidance [117, 136].

Although these robotics studies demonstrated impressive accuracy and speed for complex music playing, very little attention has been paid to the understanding of delicate embodied interactions of players and instruments for expressive sound generation. So far [94, 138] have analysed the importance of dynamic interactions for expressive playing, but it is still largely unknown how music expressions can be systematically analysed and understood. Generally speaking, expressive piano playing is a manifold problem involving the dynamics of the instrument, note arrangement in music instructions (sheet musics), and player's action, and we are not able to independently investigate each of these components in isolation as they are mutually related to each other [74].

The problem addressed in this paper is therefore the development of a method to systematically analyse the relationship between these three components, by employing a state-of-the-art digital piano, robot arm platform, and a statistical computational tool based on Gaussian Process (GP) inference. For a systematic analysis and comprehensive understanding of the landscape of this framework, we employed a minimalistic approach where we consider 10 basic playing styles, expressed by a single note, with a finger performing key-presses on

a piano instrument. As exemplified in later sections, even with this simplified setup, the systematic understanding of expressive piano playing is nontrivial.

For this challenging problem, this paper argues that the relationship between the motor control of a player and the corresponding expressive auditory output on the piano is intrinsically nonlinear, thus specific treatments are necessary when designing and analysing motor control of piano players. The expressive piano playing is known to be analysable by the MIDI format of music sound representations, in which expressive sounds are related to the velocities of piano key-pressing and interval times between them. Based on this framework, we will extend the analysis to robot control to show the non-linearity of the relationship between expressions and a player's motor control. The identification of this nonlinear nature of piano playing is particularly important in order to understand players' (bio)mechanical dynamics, control, and learning processes. In this context, the mechanical dynamics (impedance) of players' fingers, arms and hands are important. Additionally, linear regression methods may not be flexible enough to cope with the nonlinear dynamics of this system, and other nonlinear control optimisation (learning) processes become instead necessary.

In the past, humans have been shown to learn and make decisions with processes akin to Bayesian inference, above all in tasks involving sensory-motor control [125]. It is in this context that we propose a fully probabilistic GP-based framework to capture the relationship between the piano music and key-press events that generated it. Advantages of this approach include a mathematically meaningful measure of uncertainty in key-press trajectory prediction. While these implications are valid for both human and robot players, it would be particularly interesting for designers of piano playing robots, because a hard-coded linear mapping of motor control would not be sufficient for human-like playing of piano but an integrative view of morphology and sensory-motor control become more valuable in the context of dexterous manipulation tasks [227, 229, 232].

This paper is structured as follows. Section 9.2 reports the methods in this Paper, including the GP-based Learning framework in Section 9.2.1, and the robotic experimental set up in Section 9.2.2. In Section 9.3 we report the results of this work. Finally, in Section 9.4 a discussion and a conclusion are provided.

9.2 Methods

9.2.1 Learning Framework

The framework developed for this work aims to capture the relationship between piano key-press events and the corresponding piano sound outputs, thus optimizing the robot's

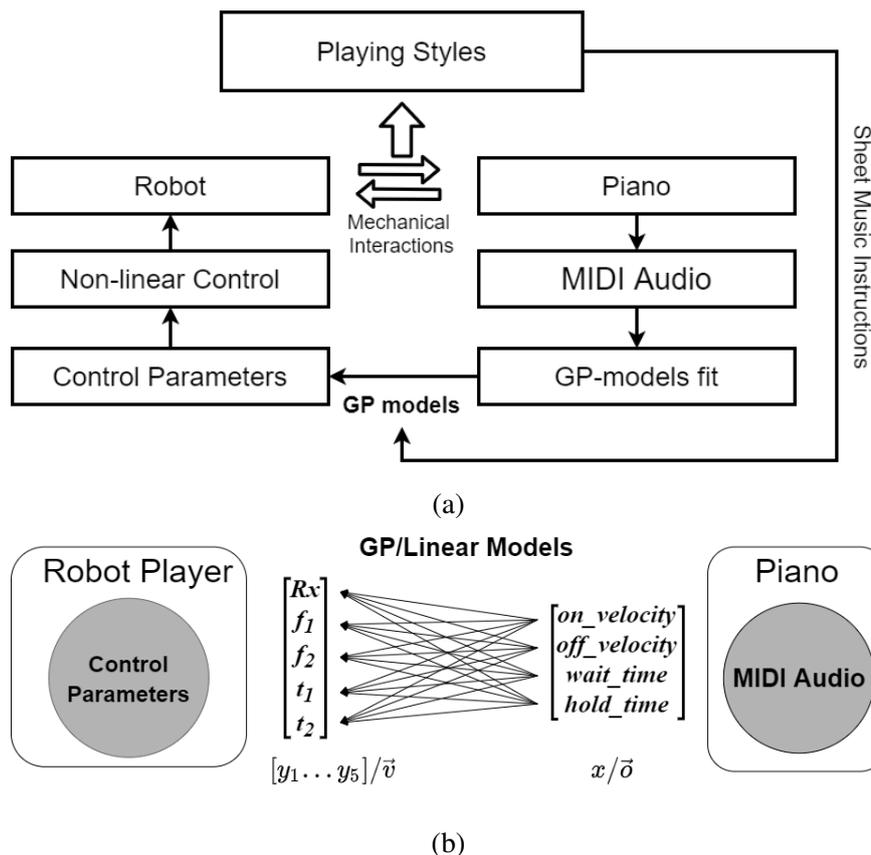


Fig. 9.1 Adaptive piano playing diagram, including (a) the GP-based framework developed and (b) the model relationship between robot key-press and piano sound outputs.

key-press trajectories for different styles, through a single demonstration. Much like a human player, the robot can perform key-presses on a piano, observe the resulting music output, and then explore its own action space and the consequences of its actions through sound feedback (Fig. 9.1). The music styles chosen for the experiments are commonly used in piano playing to evoke different musical expressions. Two types of fundamental musical parameters governing musical events are explored, articulation and dynamics, for which a musical event is typically a single note or phrase of notes. Music articulations shape the attack, decay and length of an event, while dynamics determine the loudness of an event relative to the entire passage. Articulation methods *tenuto*, *staccatissimo* and

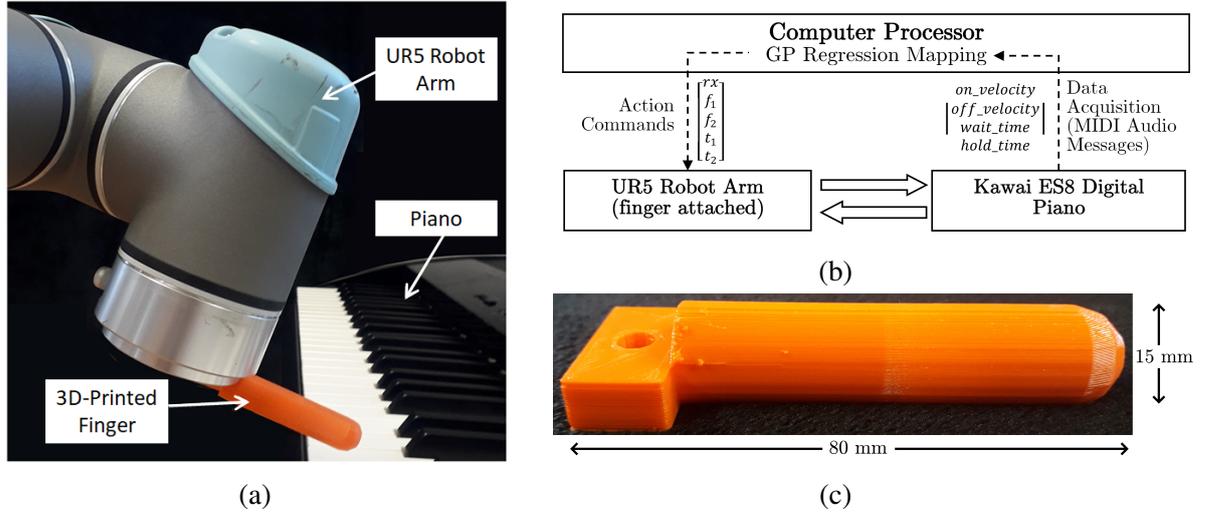


Fig. 9.2 The set-up for the experiments. In the figure, (a) shows robotics set up, including a schematic of the robot connection to a processing unit and the musical instrument, while (b) shows the 3D finger used for piano playing.

staccato were chosen for the experiment for their relevance to the piano instrument and suitability for monophonic (single-note) playing. The styles of *fortissimolff* (very loudly) to *pianissississimolpppp* (extremely softly) were instead chosen for their dynamic range. These playing styles introduce a wide range of music which require very diverse types of key-press action to be performed on the instrument.

Linear Models

The first set of models are Linear models, capturing the relationship between the detected key-press piano sounds and their corresponding robot key-press actions. Fig. 9.1a and Fig. 9.1b show a qualitative diagram of the framework in the context of these experiments. For each key-press let \vec{v} be the d_v dimensional vector of control parameters utilized for the robot control, and \vec{o} the d_o dimensional vector of the corresponding sound outputs. An in-detail explanation of the parameterization of the robot control and the sound output is irrelevant to the learning framework, and will be provided in later sections.

For notation's sake we will impose $x = \vec{o}$ and $y = \vec{v}_i$. In this context, \vec{v}_i is a one dimensional vector, corresponding to one dimension of the control action parameters (Fig. 9.1b). The following equations will be repeated for every control action dimension, thus $i \in [1, \dots, d_v]$, where $d_v = 5$ for the duration of the experiments.

For the Linear models we impose $x' = [x \ 1]$ and :

$$y = \vec{w}x' \quad (9.1)$$

where \vec{w} is a $d_o + 1$ dimensional vector of weights capturing the relationship between the sound outputs and the control parameter \vec{v}_i under consideration. The values for \vec{w} are approximated by a Least Square fit.

Gaussian Process Framework

In the past, humans have been shown to learn and make decisions with processes akin to Bayesian inference, above all in tasks involving sensory-motor control [125]. In contrast to the linear models, a GP-based framework is used to capture the relationship between the sound produced by the piano and the robot control of the key-press generating it [271]. Given key sound observations x , generated by noisy robot-controlled key-presses y , the relationship of sound output to motor control can be captured by:

$$y = f(x) + \varepsilon \quad \text{where } \varepsilon \sim \mathcal{N}(0, \sigma_y^2) \quad (9.2)$$

i.e. the noisy key-press by the robot control are assumed to have a Gaussian process prior and be drawn from:

$$y = f(x) \sim GP(m(x), k(x, x') + \delta_{pq}\sigma_y^2) \quad (9.3)$$

where the mean is $m(x) = 0$, and the covariance of any two noisy observations y_p and y_q is:

$$\text{cov}[y_p, y_q] = k(x_p, x_q) + \delta_{pq}\sigma_y^2 \quad (9.4)$$

where x_p and x_q are the inputs to the corresponding observations, and $\delta_{pq} = I(p = q)$. The relationship between y and x in Eq. 9.3 is thus dictated by how any two musical outputs co-vary in terms of their generating key-press trajectory. The covariance of any two points is governed by Eq. 9.4, and thus the choice of the kernel is here important. We build on a linear kernel, and account for non-linearities in the relationship of x and y by a Radial Basis Function Kernel, thus:

$$k(x_p, x_q) = x_p x_q \sigma_f^2 e^{-\frac{\|x_p - x_q\|^2}{2l^2}} \quad (9.5)$$

where σ_f^2 and l are hyperparameters which decide the magnitude of influence of adjacency when evaluating the function at any one point. From [165], we can write the mean μ_* and variance Σ_* for any new test audio input X_* , prior inputs X and generating observed control key-press y as:

$$\mu_* = k(X, X_*)^T K_y^{-1} y \quad (9.6)$$

$$\Sigma_* = k(X_*, X_*) - k(X, X_*)^T K_y^{-1} k(X, X_*) \quad (9.7)$$

where $K_y = k(X, X) + \sigma_y^2 I_N$, to account for the noisiness of the observations.

Finally, it is desirable not to manually pick the hyperparameter σ_f^2 and l of the covariance function. We therefore perform model selection by initializing the σ_f^2 and l to 1 and iteratively minimizing the negative marginal log likelihood $-\log p(y|X)$ over 100 training steps as implemented in [157].

The equations described can capture the relationship between any sound output parameter \vec{v}_i and control input \vec{o} . In this paper $d_v = 5$ and $d_o = 4$, thus five 4-dimensional Gaussian Processes are built to automatically capture the relationship between the sound output and robot key-press control.

9.2.2 Experimental Set-Up

For the experiments we use a UR5 robotic arm, equipped with a custom end-effector (Fig. 9.2a and Fig. 9.2c). The music instrument is a Kawai Es8 Digital Piano, which provide the possibility to retrieve event-based, MIDI audio messages when a key is pressed. An audio message is generated when one of two events is detected: a key press or a key release. For every pair of detected MIDI message, four variables are going to be relevant for the purpose of the experiments in this paper, namely: the velocity of the key-press, the time the key was held down, the velocity of the key release and the wait time before performing the next key-press; the four variables will be referred to *on_velocity*, *hold_time*, *off_velocity* and *wait_time* for the remainder of the experiments.

Finger Design

The finger was designed to be the simplest end-effector to allow the UR5 robotic arm to perform single key-presses on a standard piano. The finger is a $80\text{mm} \times 15\text{mm}$ cylindrical attachment, with a flat origin and a rounded finish, to perform key-presses at various stroke angles without compromising the area of contact (Fig. 9.2c). The finger was 3D-printed using FilaFlex, a Thermoplastic Polyether-Polyurethane elastomer (TPE) filament of shore hardness $82A^2$, and thus it presents some room from flexing and bending.

²<https://recreus.com/en/12-filaflex-original-82a>

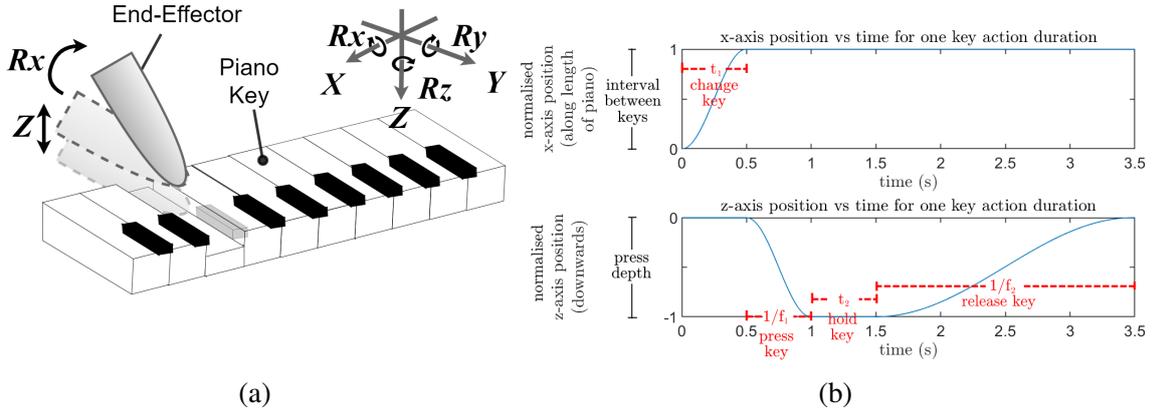


Fig. 9.3 The robot control, including (a) an illustration of a robotic key-stroke, and (b) the hybrid sinusoidal control designed for the experiments.

9.2.3 Robot Control

The robot was controlled in Cartesian coordinates at $\approx 125Hz$, acting upon the Z and Rx tool axis simultaneously, to generate the desired contact between the end-effector and piano key for a key-press. A hybrid sinusoidal displacement profile was generated for the Z axis, parameterized in both amplitude and frequency, while an angle of rotation about the finger tip was chosen for the Rx axis, Thus a total of 5 parameters were used to control the robot for each key-press experiment ($d_v = 5$), i.e. Rx, f_1, f_2, t_1, t_2 (Fig. 9.3b).

For the Z axis of motion, a sinusoidal displacement over the course of the key-press is defined as:

$$s_z(t) = \begin{cases} 0 & \text{if } 0 < t \leq T_a \\ \frac{1}{2}A_z[\cos(\frac{2\pi f_1}{2}(t - T_a)) - 1] & \text{if } T_a < t \leq T_b \\ -A_z & \text{if } T_b < t \leq T_c \\ -\frac{1}{2}A_z[\cos(\frac{2\pi f_2}{2}(t - T_c)) + 1] & \text{if } T_c < t \leq T_d \end{cases} \quad (9.8)$$

where

$$T_a = t_1, T_b = t_1 + \frac{1}{f_1}, \\ T_c = t_1 + \frac{1}{f_1} + t_2, T_d = t_1 + \frac{1}{f_1} + t_2 + \frac{1}{f_2}$$

Here, ‘t’ is the time, in seconds, elapsed since the start of the touch experiment and A_z defines the amplitude of the generated sinusoidal displacement for the key press, and it is here set to $32mm$ throughout the experiments. Additionally, a parameter Rx sets the angle of rotation of

the end-effector, between 0° and 90° throughout the experiments (Fig. 9.3a). The X axis of motion controls the robot's ability to shift key along the piano as it plays different notes. The control on the X axis is achieved through:

$$s_i(t) = \begin{cases} \frac{1}{2}A_i[\cos(\frac{2\pi}{2t_1}t) - 1] & \text{if } 0 < t \leq T_a \\ A_i & \text{if } T_a < t \leq T_d \end{cases} \quad (9.9)$$

$$i \in \{Rx, x\}$$

where $A_x = k_d n_k$, k_d represents the key-width and n_k is the number of keys between the previous and current key. The standard modern piano keyboards are designed with white keys 23.5mm wide, thus we set $k_d = 23.6$, taking into account the gap between keys.

The choice of parametrization is here important. The sinusoidal displacement at the beginning and at the end of the key-press were chosen to provide a smooth robot trajectory of key-press and key-release, which could be affected by simple amplitude and frequency parameters. The t_1 and t_2 parameters were instead introduced to allow the robot to directly influence the times of key-press and key-change. Our choice of parameters was also to induce a strong correlations between the control parameters and piano MIDI outputs, as later shown.

The following sections will explore how different control parameters can approximate different playing styles, and how these may be learned online through sound feedback.

9.3 Results

In the following sections we wish to understand the delicate embodied interactions of players and instruments for expressive sound generation. We first show how expressive piano playing is a manifold problem, involving the dynamics of the musical instrument, note arrangement, and player's action. Here, we show that the relationship between motor control and piano is intrinsically non-linear. We will further show the viability of the GP-based framework developed in capturing the non-linear dynamic relationship of this system, and its advantages with respect to simpler linear regression methods. Finally, the optimized controllers for 10 different playing styles are compared with the performance of an expert human player.

9.3.1 Robot key-press control to sound feedback

In the first set of experiments we investigate the relationship between the robot control parameters and the generated sound outputs following the robot key-press control. This analysis is based on observations on a large-scale set of over 3125 key-press experiments performed

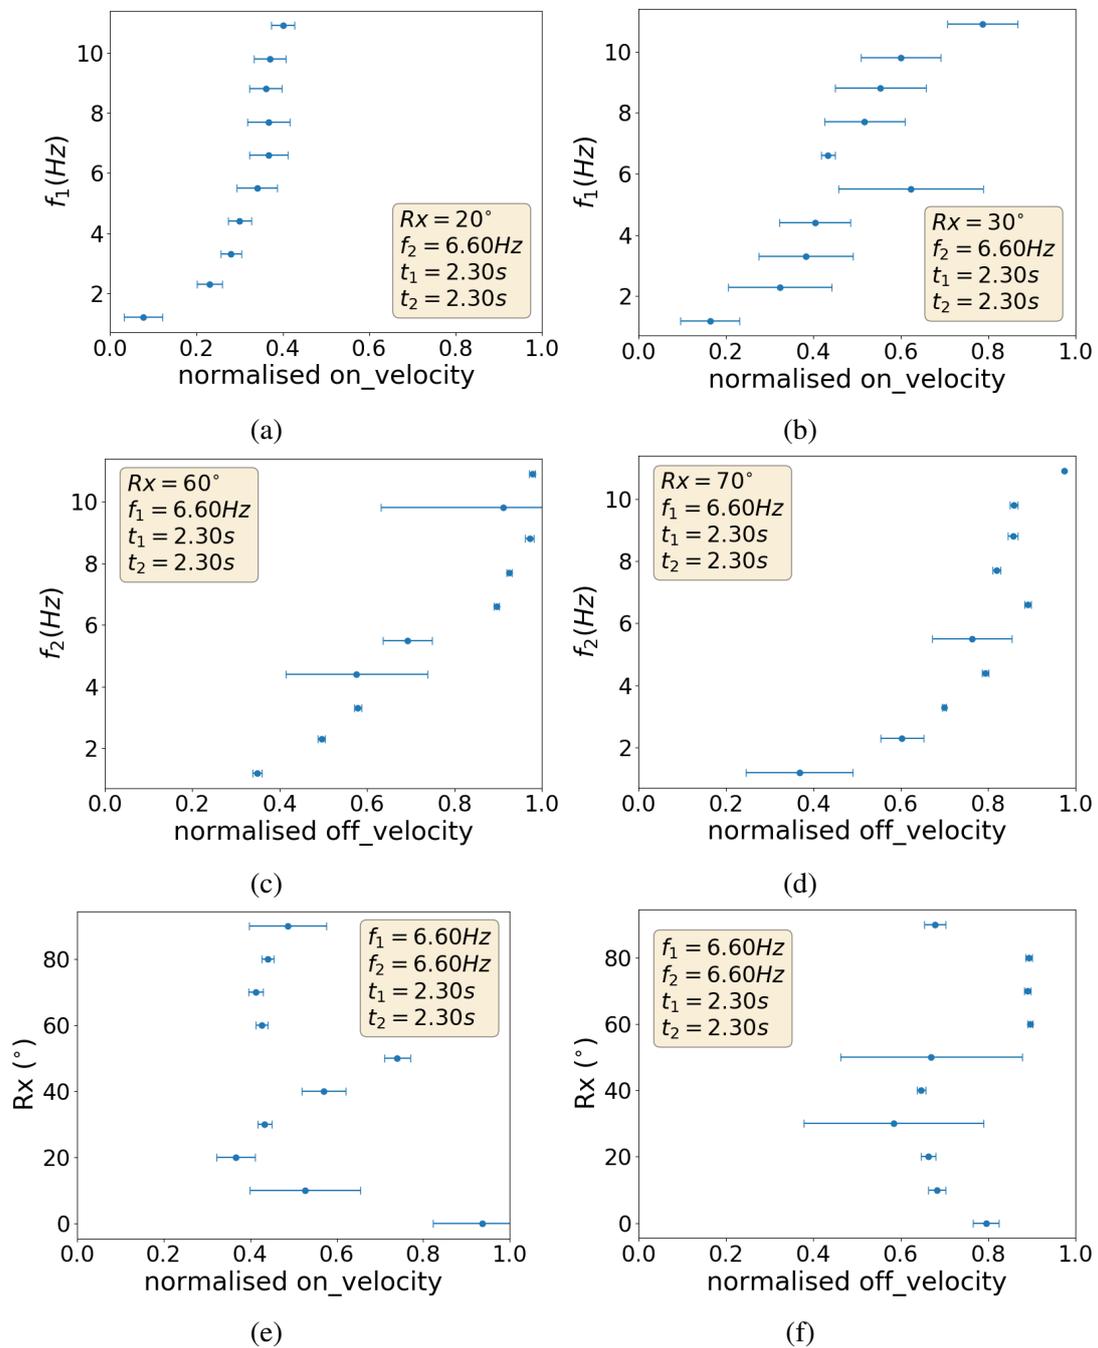


Fig. 9.4 The sample raw data corresponding to the MIDI sounds registered by the piano and the control parameters generating the key-presses execution, averaged over 10 trials.

with the set-up described in the previous section. Fig. 9.4 shows example relationships between the MIDI parameters of *on_velocity* and *off_velocity* with the control parameters of f_1 , f_2 , and Rx . From Fig. 9.4a, it is possible to see how the *on_velocity* increases as the robot control's f_1 parameter increases, while all other control parameters are kept constant. However, the normalized value of *on_velocity* saturates at ≈ 0.4 regardless of an increase in f_1 from 6.6 Hz. Beyond a frequency threshold of 6.6 Hz, any higher imposed frequencies in the key-press control appear indistinguishable by the piano key's velocity-sensitive trigger sensor. The piano key trigger has not reached its velocity sensing saturation, as we observe from Fig. 9.4e that at other finger rotation angles (Rx), the normalized *on_velocity* values are able to reach up to ≈ 0.8 compared to the saturation at ≈ 0.4 previously observed at $f_1 = 6.60$ Hz. It is likely that this is in part due to the elasticity of the finger, which is capable of flexing and bending to some degree, combined with the sinusoidal parameterization of key-press trajectory. Both the finger's make and the choice of action parameterization for the robot key-press control, in fact, induce slight changes in both the stiffness and contact point of the finger with respect to the stroked key during a key-press, with higher angles inducing higher degrees of stiffness in the end-effector.

We observe a similar trend between f_2 and *off_velocity* in Fig. 9.4c and Fig. 9.4d. The normalized *off_velocity* detected by the piano increases as the f_2 parameter is increased in the robot control, while all other control parameters are kept constant. However, the value of *off_velocity* saturates at values of f_2 frequencies of $\approx 6.6\text{Hz}$. Beyond a frequency threshold of 6.6Hz , higher release velocities are indistinguishable by the trigger.

The non-linear relationships between Rx and audio parameters *on_velocity* and *off_velocity* respectively are harder to capture as there are additional factors at play. As the finger's material is non-stiff material, and given its elongated structural composition, key-presses at different angles may vary the finger's stiffness, as the generated forces derived from the key-press may be more or less normal to its longest side.

The non-linear relationships observed are also not representative of those observed at other constant variable values. This is illustrated by comparing Fig. 9.4a and Fig. 9.4b for which all constant robot control variables are the same except for $Rx = 20^\circ$ for the data in Fig. 9.4a and $Rx = 30^\circ$ for that in Fig. 9.4b: there is a less significant plateau observation of *on_velocity* in the latter as the value continues to increase gradually to 0.8 at $f_1 = 10.9$ Hz. This is likely because at the higher rotation angle, the depth at which the piano key is electronically triggered corresponds to a different point along the gradient of the sinusoidal curve in Fig. 9.3b, causing a velocity difference that is distinguishable by the trigger as f_1 increases. Other factors may also contribute to this difference, such as the different

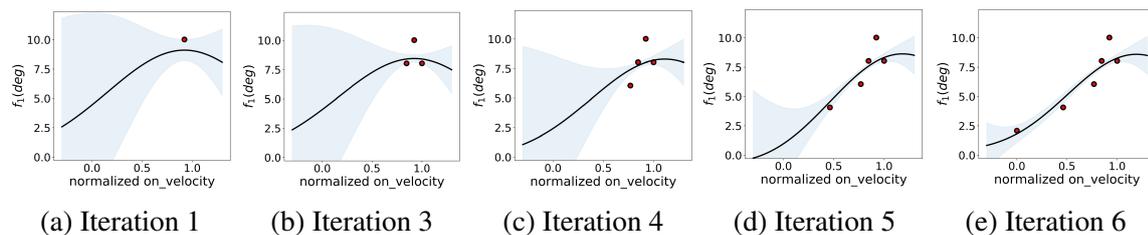


Fig. 9.5 The GP-based exploration fit over different iteration steps when running the framework with simple f_1 control on the $on_velocity$ parameter.

mechanical properties of the finger at different rotation angles and the differing point of contact of the finger on the key.

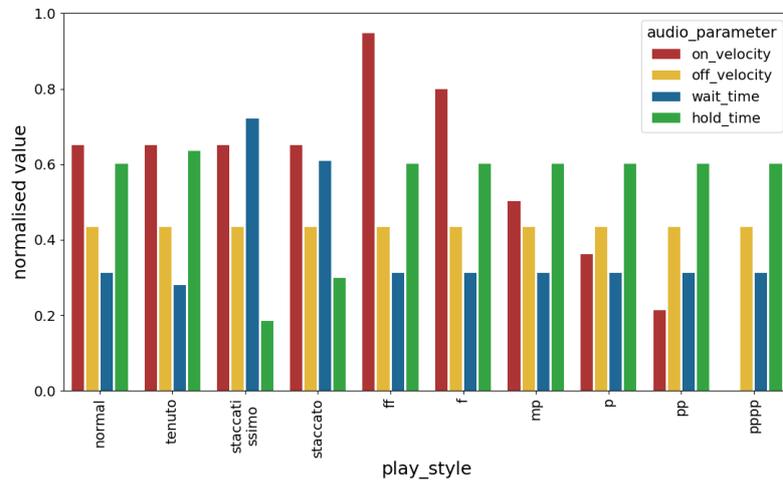
Similarly when comparing Fig. 9.4c and Fig. 9.4d, which plot the data obtained from setting $Rx = 60^\circ$ and 70° respectively, we observe a more significant plateau between $f_2 = 7.7$ Hz and $f_2 = 9.8$ Hz in the latter figure, where the velocity change due to f_2 is indistinguishable by the key trigger.

The analysis of the raw data shows the complexity, multi-dimensionality and non-linearity of problem at hand, where the physical interaction of the robot's finger and the piano instrument is quantified experimentally.

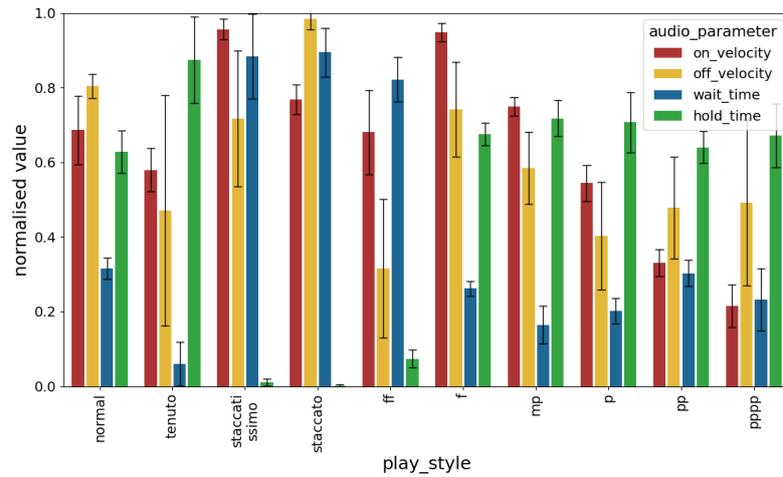
9.3.2 Gaussian Process based Framework analysis

As shown in the previous section, the relationship between the control parameters and resulting note musical outputs is both non-linear and multivariate dependent. Gaussian Processes can capture both the non-linear nature of the relationship between the inputs and outputs, and the dependence across parameters.

In the second set of experiments it is shown how the GP-based framework developed can approximate a parametric fit during training. We initially thus ignore the complexity of multivariate fits and run the framework by optimizing a single control parameter with respect to one MIDI output. We chose a control parameter and MIDI output which should show some degree of correlation, e.g. f_1 and $on_velocity$, and run the algorithm to train the robot over 12 key-press, or iterations. Fig. 9.5 shows the algorithm at 5 different stages within the 12 iterations. As shown in the figure, for each parametric value attempted by the robot, the uncertainty of the fit at that point collapses, and is later related to the variance of the fit at that point. By iteration 6, the robot has found a fit over almost the whole controllable parameter space. At this point, the GP model trained on the same samples can be used to inference the control necessary to reproduce a wanted MIDI output.



(a) Digital



(b) Human

Fig. 9.6 The 10 different playing styles addressed in this work. (a) The playing styles generated by MuseScore digital score writer, and (b) the play styles as played by the human player. The variance between the MIDI parameters shows the fundamental differences between the various styles.

In the next set of analysis we now consider all parameters, i.e. five 4-dimensional GP models are fit, to capture the relationship between the 4-dimensional MIDI audio piano outputs and each of the 5 control parameters. At each iteration, we use Equation 9.6 on the five 4-dimensional GP models, to choose controls to approximate each of the playing styles shown in Fig. 9.6a. We use MuseScore³, a digital score-writer computer program that supports MIDI output, to generate each playing style, and the resulting sound output. The actual sound outputs generated by the inferred control can then be compared to the MIDI outputs to reproduce, and an error can be computed by:

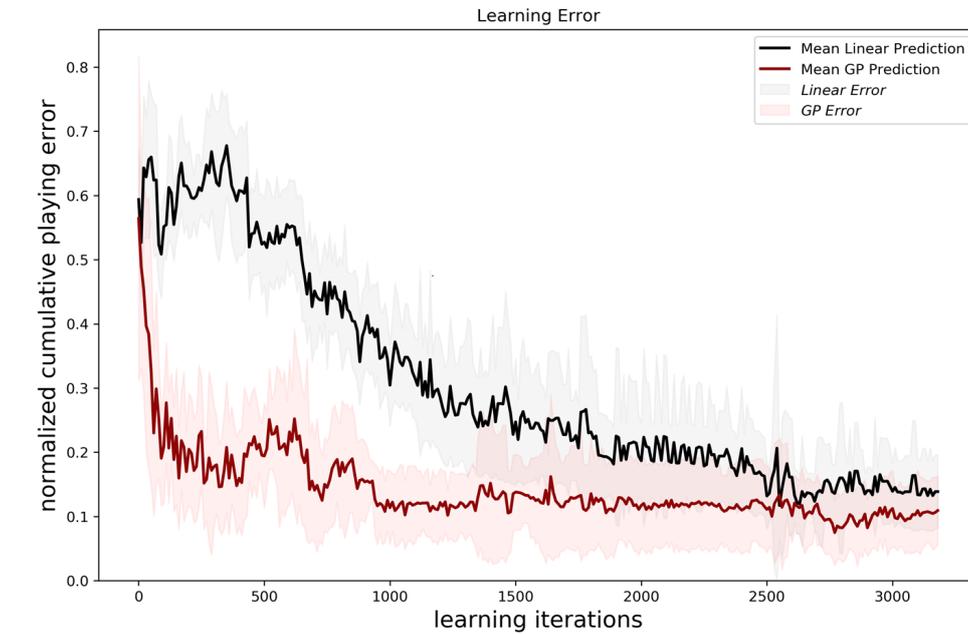
$$error_s = \sqrt{(\vec{o}_{inv,*} - \vec{o}_{s,midi})^T (\vec{o}_{inv,*} - \vec{o}_{s,midi})} \quad (9.10)$$

where $\vec{o}_{inv,*}$ is the MIDI output generated by applying the inferred control $\mu_{inv,*}$ of the inverse model, and $\vec{o}_{s,midi}$ is the reference MIDI output for the playing style under consideration.

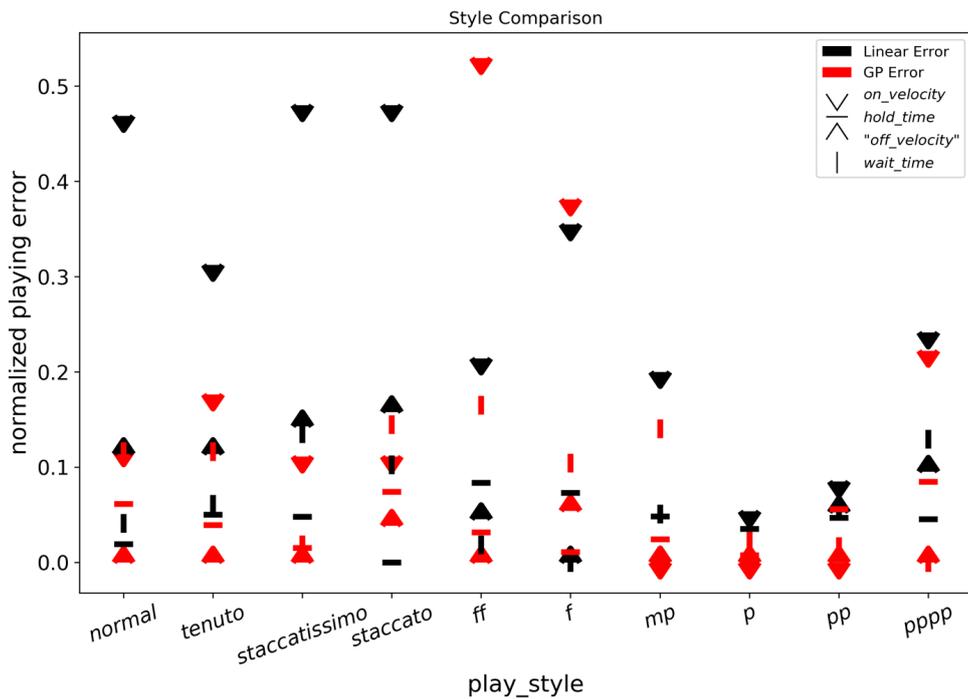
We compare the robot note error over the 10 different playing styles when learning through the GP-based framework developed, against a linear fit of the $x - y$ parameters in Section 9.2. In both cases the robot searches the space of each control parameter in a breadth-first grid-search fashion, with a parametric discretization of each parameter into 5 equally spaced values. The robot searches each parameter combinatorially, so a total 3125 key-press are performed to incrementally train the linear and GP models. Fig. 9.7a shows the sound errors of testing key-presses on the piano, after testing the fits every 30 training key-press iterations. For each testing epoch, the robot is made to test each playing style 3 times for the linear model, and three for GP-based framework, bringing the total number of experiments to 9375 piano key-press for both training and testing, with a split of 50% and 50% respectively. From Fig. 9.7a it is clear how the GP-based framework developed is capable of outperforming the simpler linear model, bringing the lowest error to 0.0747 MIDI units as opposed to 0.117. More interestingly, the GP-based framework reaches convergence after approximately 1000 iterations, a factor of three times smaller than the time necessary to approximate the playing styles by the linear models. The limitations of the linear model, in this setting, follow the relationships observed in Fig. 9.4, where the control parameters showed a non-linear effect on the MIDI outputs. This non-linearity captures the reality of key-press experiments with an elastic finger, where the dynamics of the collisions of the finger and the piano key, with respect to the generated output, can at best be approximated, but not truly captured, with simple linear relationships.

Fig. 9.7b sheds some light into the limitation of robotic piano playing with a set up analogous to our own. The robot, in fact, is incapable of matching the key-press velocity necessary to approximate each playing styles, when learning off sound feedback through

³<https://musescore.org/en>



(a)



(b)

Fig. 9.7 The comparison between the Linear and GP fits to approximate 10 different piano playing styles. (a) The testing Error over 3125 training key-press, and (b) the average testing error by play-style, for the best validating epoch during training for the linear and GP models respectively.

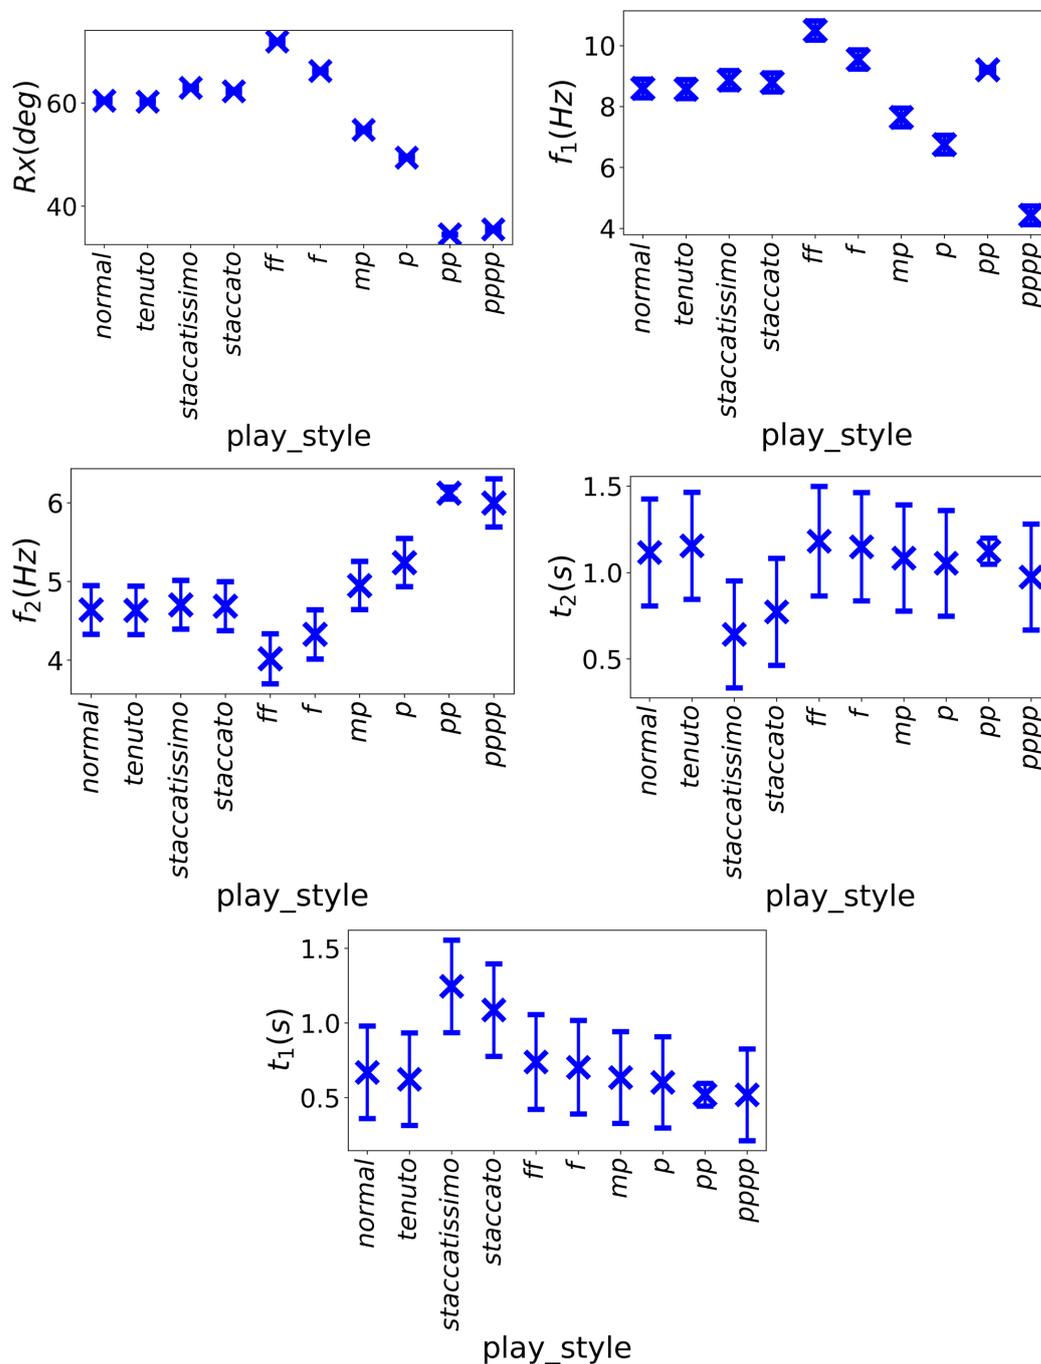


Fig. 9.8 The final predictions, after 3125 learning epochs, of the required control parameters for 10 different playing styles. The Confidence intervals are shown as normalized predictive variance, and show the predictive uncertainty of the GP model.

simple linear models. The GP-based framework outperforms the linear models by a larger margin in 8 out of 10 playing styles. The framework results better suited in capturing duration relationships between target temporal patterns and control key-stroke dynamics, additionally to highly accurate control on slow-speed downward key-strokes, effectively reaching lower MIDI errors for the styles of *normal*, *tenuto*, *staccatissimo*, *staccato*, *mp*, *p*, *pp* and *pppp*. The Linear model results capable of better capturing the relationship between high speed robot control for downward key-press actions, regulated by f_1 , and target louder sound outputs, finally achieving better performance in the styles of *ff* and *f*. The difference in performance for the styles of *f* and *ff* can be explained by the intrinsic tendency of the linear model to overestimate the levels of downward velocities required to achieve high *on_velocity* outputs. The downward velocity levels, regulated by f_1 , can in fact be observed to saturate at certain levels in Fig. 9.4a and Fig. 9.4b, levels which depend mainly on the finger angle to the piano key. The linear model will not be able to capture the *on_velocity* plateau, overestimating downward key-stroke velocities, but effectively achieving louder outputs for the *f* and *ff* styles.

Fig. 9.8 shows the control parameter values attempted by the robot as generated from the GP model prediction. It is clearly shown that for playing styles *normal*, *tenuto*, *staccatissimo* and *staccato*, control parameters Rx , f_1 and f_2 have very similar values, at approximately 60° , $8.5Hz$ and $4.5Hz$ respectively. The control parameters t_1 and t_2 vastly vary across these playing styles, showing variations within $0.5s$, and indicating a large contribution of these in the playing style's unique characteristics. On the other hand, for playing styles *ff*, *f*, *mp*, *p*, *pp*, *pppp*, the change in dynamics, which clearly defines these playing styles' unique characteristics, is largely contributed by control parameters Rx , f_1 and f_2 , with observed changes of the magnitude of 15° , $6Hz$ and $3Hz$ respectively. Moreover, an invaluable advantage of the GP framework proposed is the uncertainty estimation. Analogously to the prediction computation through Eq. 9.6, we use Eq. 9.7 to compute the uncertainty, or variance, of a control when attempting to generate a target sound output. In Fig. 9.8, the uncertainties are shown in terms of normalized variance for each prediction, to better visualize the plot trends. From the figure it is clear how the robot results more confident in generating both rotations and downward velocities. Temporal parameters (controlled by t_1 and t_2) and upward velocities (controlled by f_2) result somewhat harder to grasp over the different playing styles. The high uncertainty over t_1 , t_2 is indicative of one of two factors: one, that the grid-search parametric exploration of t_1 and t_2 did not attempt any combination which was close to perform any of the playing styles accurately, and thus no ample evidence is present for the inferred control parameters; two, the robot wait time control through t_1 and t_2 shows high degrees of variability in terms of the actual wait time outcome. Given

further results shown in Section 9.3.3, the first case is more likely. The f_2 parameter is also not capable of achieving full control of piano key-release velocities. This is likely due to the dynamics of the piano key-release action, which limits the speed at which each key springs back to its original position after a key-press. For higher key-lift velocities performed by the robot, and controlled by f_2 , then, the detected velocity of a key release by the piano will eventually saturate.

9.3.3 Human vs Robot Piano Playing

Finally, we investigate the ability of the robot to perform the 10 different playing styles in Fig. 9.6a as compared to an expert human pianist. We use the controllers optimised by the GP-based framework developed. The human performer is a veteran pianist with 15 years of history in piano playing. Upon listening to the sound output the pianist is made to reproduce the note on the piano. We collect 10 different key-press samples at 40 beats per minute (BPM) or 1 key-press every 1.5s, performed 4 times by the human pianist, for each playing style, so as to have an idea of the playing variation within each style. The resulting normalised MIDI output from the pianist's playing is shown in Fig. 9.6b. We let the robot perform according to the $\mu_{inv,*}$ (Eq. 9.6) extracted for each playing style after learning through 3125 iterations, log the resulting $\vec{\sigma}_{inv,*}$ from the robot playing and $\vec{\sigma}_{human,*}$ from the human player, and compute errors with respect to the computer generated MIDI for each playing styles.

Fig. 9.9 compares the human and robot's normalized error performance for each playing style. Fig. 9.9a shows the error of the human and robot's performances in terms of *on_velocity*. The performances of the robot and the human are highly similar for most playing styles, on average within 0.05 normalized error units from each other for all playing styles. The GP-optimized robot's errors are 0.05 and 0.03 units higher than the human player at *ff* and *f* as the robot is unable to achieve equally high downward velocity required for the loud playing. On the other hand, the robot's errors are lower at *mp*, *p*, *pp* and *pppp*, with error differences ranging from 0.05 to 0.02 normalized units. The precise motion control at low velocities achieved by the robot is, in fact, capable of precisely approximate soft key-presses, which the robot optimizes with respect to the reference MIDI target style. Fig. 9.9b shows the differences between the human player and the MuseScore generated playing styles in terms of *off_velocity*. Due to an innately more dynamic and highly varied key release motion by the human player, the human tends to show a diverse range of release velocities, with errors of up to 0.4 normalized units. The robot, on the other hand, has low variance and error across all play styles due to its consistent speed control for key releases. In Fig. 9.9c we compare the human and the robot performance over the *wait_time* parameter, i.e. the time necessary to wait between key-presses for each playing style. Surprisingly, the robot

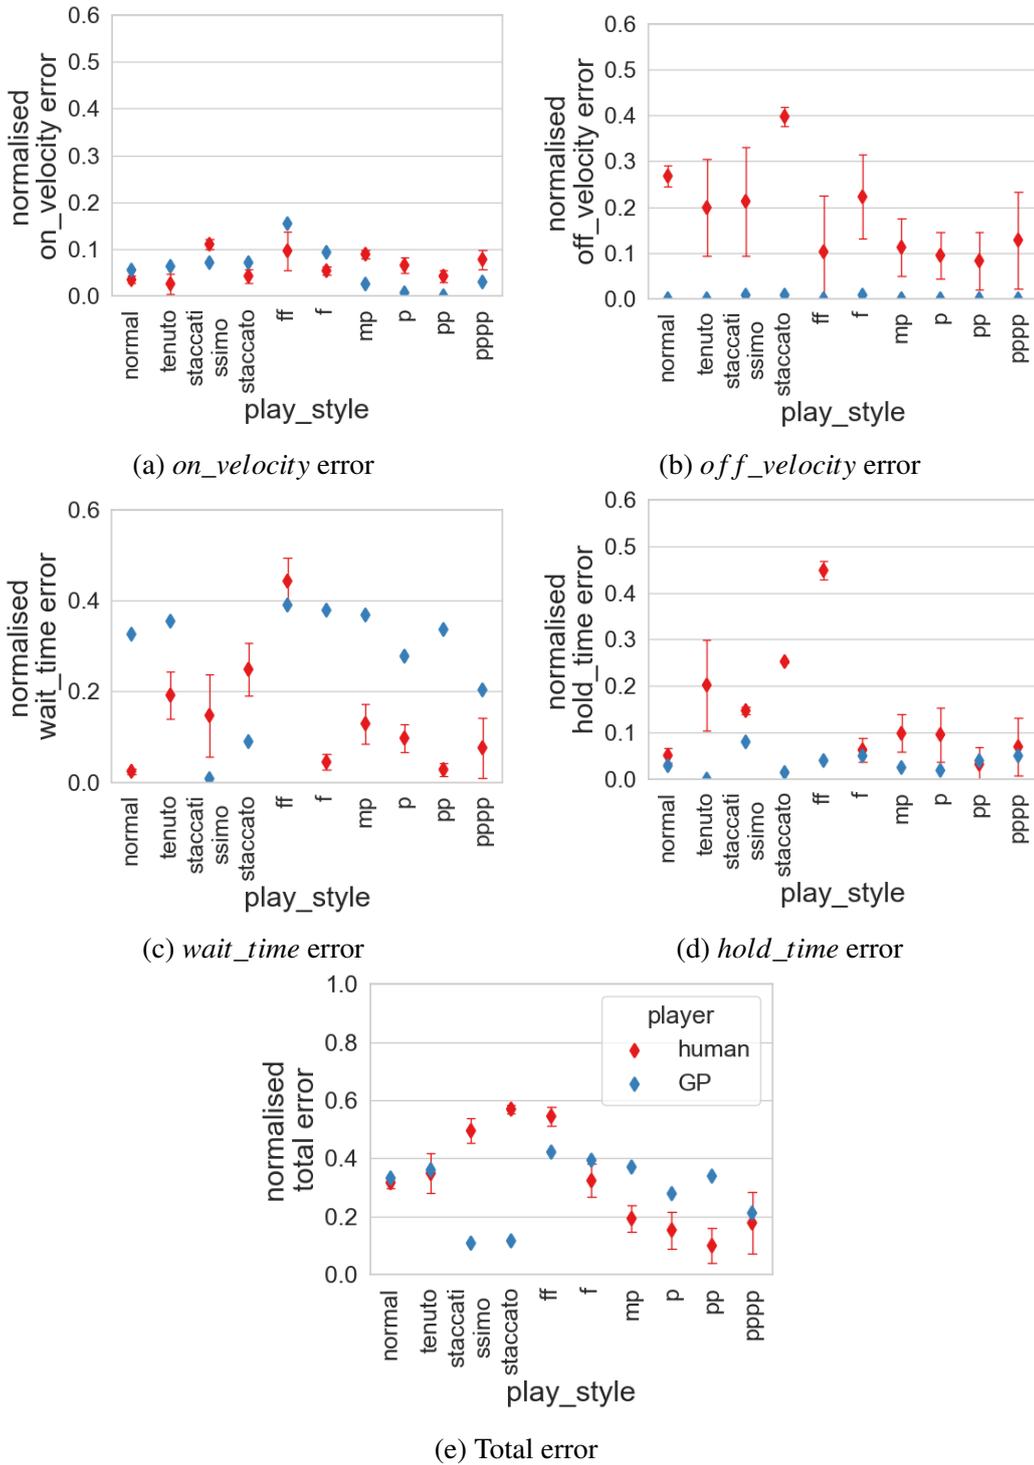


Fig. 9.9 The Comparison of the playing score of the robot optimized to play the 10 different playing styles after 3125 learning epochs, a human player playing the same, and the computer generated outputs.

playing error results higher than the human of 0.1 to 0.4 normalized units in most playing styles, as it is unable to play a melody consistently at 40 BPM with the required waiting time between notes. This is likely due to the inherent delays in the robot's online control when switching commands between key-presses, a consequence of the chosen parameterization and robot key-press control in the experiments, while the human player has a good grasp of rhythm and plays each note at consistent intervals. The robot performs more consistently than the human player in terms of *hold_time* error, with errors lower than 0.1 normalized units across playing styles as shown in Fig. 9.9d, due to the precise clock control during the holding phase of the key-press. In terms of *wait_time* and *hold_time*, the human player's style errors are higher than the robot's for *tenuto*, *staccatissimo*, *staccato* and *ff*, with error value differences varying from 0.1 to 0.5 normalized units. The timing of these articulation styles are exaggerated by the human for greater impact and variation in expression, thus deviating further from the MuseScore generated ground truth. Also note that the error for *wait_time* and *hold_time* (Fig. 9.9c and Fig. 9.9d) show similar trends for the human player, due to the aforementioned good grasp of rhythm; there is no delay between notes and a longer *wait_time* is always compensated by a shorter *hold_time* for that note played. On the other hand, the robot's delays between key-press commands are strongly reflected only in the *wait_time* error. Finally, the overall normalized error by play-style shown in Fig. 9.9e shows an interesting picture. After 3125 learning epochs, the robot is able to perform similarly to the human player for *normal*, *tenuto* and *pppp* playing styles, with normalized error differences lower than 0.01 units. The robot achieves lower errors ranging from 0.1 to 0.4 normalized units for the styles of *staccatissimo*, *staccato* and *ff*, largely due to its accurate *off_velocity*, *wait_time* and *hold_time* performance. The robot, however, performs with errors between 0.1 and 0.2 normalized units higher than the human player for the styles of *f*, *mp*, *p*, and *pp*, largely due to its poor *wait_time* performance.

9.4 Discussion and Conclusion

We investigate the ability for a robot to play the piano according to 10 different playing styles, like a human player. We propose a GP-based framework for the robot to incrementally model the relationship between the control actions, to the resulting sound output, and learn appropriate controllers to play according to each music style. We show that the relationship between control and sound is non-linear in nature, and that different control parameters are not independent with respect to the generated note from the corresponding key-press. The GP-based model can faithfully capture the relationship between control and generated music output, outperforming simpler linear model.

To be able to play different playing styles faithfully it is necessary for the robot to explore its action space, so to find appropriate key-press for each style. The resulting combinatorial explosion in parametric search presents itself as an issue. A second advantage of the proposed GP-based framework is its ability to quickly converge to appropriate controllers for each style. In fact, we observe the GP convergence to be a factor of 3-times faster than linear models, with respect to the learning of the playing styles considered in this work.

The main limitation of the approach lies with the drawbacks of GP modelling. As the model takes into account every single sample to compute the fit, it can eventually be computationally expensive to fit the control to MIDI relationship. This, can in part be obviated by methods which do not need a full kernel representation, and by the dismissal of points far away in time with a sliding learning windows [258].

Finally we compare the ability of the robot to approximate each of the playing styles, with respect to an expert human player. We show the comparison sheds some light to several interesting aspects of robotic piano playing. The robot is capable of performing comparatively to the human player in the styles of *normal*, *tenuto*, *staccato* and *pppp*, largely due to the precise control at low speeds, and clock waiting times. The human-player, however, exhibits a much more dynamic and varied playing, which allows them to achieve lower style error to the MuseScore generated playing styles in *ff*, *f*, *mp*, and *p*. These styles, in fact, require higher downward key-press speeds and dynamic playing.

The dynamic and varied behaviour exhibited by the human player is one of the many advantages complex tools like human hands can possess. Partly, the lack of dynamism is indeed due to the stiffness and simplicity of the robotic end-effector. With the advent of soft-robotics and continuum robots, however, these limitations can be revoked, and the next generation of robots might indeed be able to move away from stiff and *hard-robotics* solutions, towards a softer human-like touch [42, 97]. These experiments shed some light into the limitations of robotic-piano playing, and the issues to be faced when attempting to go beyond monotonic piano playing.

This work has further shown how the complex actions play a fundamental role in determining the quality of the soft interactions arising between the robot end-effector (finger) and a piano instrument. One limitation of the current approach lies in the parameterization of the robot control action as well as the design of the morphology parameters. As the morphology and control parameters were ultimately human designed, the fully automation of robot morphing and control is still a far goal. Future work should focus on releasing some of the constraints and biases imposed by human design and aim at automating the generation of solutions, which can be pruned and assessed probabilistically with the proposed framework. A second important issue lies with the dismissal of the concepts introduced in

Chapter 6, where in these experiments the finger design and finger properties were set beforehand. Within the concepts expressed in *SomComp*, both morphology and action co-exist to allow the robot to *condition* the physical stimuli appropriately for learning. A framework which jointly considers both morphology and action is then necessary to fully embrace these concepts. The next chapter will investigate the development of one such framework.

Chapter 10

Morphology-Action Co-Optimization: Complex Tactile Object Feature Recognition

In Chapter 6, we show how the making of the body has a fundamental role in shaping the information extracted from the world. The shape of the sensor, the material it is made of, as well as its location in the robot body, all contribute to extracting particular information from the world while filtering out others. In Chapter 7 we have shown how appropriate actions can allow the robot to condition the physical stimuli at will. In both Chapters it is also shown how there is a particular morphology as well as a particular action which allow the robot to solve a specific task. The actions of the robot however, heavily depend on the morphology of

Reference Publication

This chapter was adapted from an article titled “**Efficient Bayesian Exploration for Soft Morphology-Action Co-Optimization**” [232], published in the 3rd IEEE International Conference on Soft Robotics (RoboSoft 2020). The article was written in collaboration with Dr Perla Maiolino and Dr Fumiya Iida and proposes a mathematical framework to use jointly optimize both the morphology and the action of a robot to improve on discrimination tasks. Dr Fumiya Iida and Dr Perla Maiolino contributed with the conceptualization of the topics, design of the experiments and the writing of the article. As first author in the article, my contribution includes conceptualization of the topics, the design and execution of the experiments, the robot control, the formulation of the mathematics for the framework, data analysis and the writing of the article.

the robot itself, as with a different body the robot must behave differently to achieve similar results in a task. Conversely, if a robot were to be given limiting motor constraints to solve a particular task, then there must be a morphology that is optimal for the robot's achievement of its goals through the limiting motion. The idea of morphology & motor-coordination iterative optimization is born from the need of the two paradigms to co-exist in a theory where they can improve concurrently. This chapter shows one such way to achieve the joint optimization of both morphology and action, and merges the two into a unique theory via Bayesian Exploration.

10.1 Introduction

10.1.1 Morphology-Action Co-Optimization

In the light of the experiments in Chapter 7 and Chapter 8, let's envision a scenario where a robot is placed in front of a soft phantom, containing hard inclusions of different types in various locations. In this scenario, an initial probing end-effector, equipped with a tactile sensor, could be designed and used for probing the phantom during the experiments. Like in previous chapters, the robot could find the motion strategy which allow the information retrieved from the tactile sensor to be most discriminative amongst different types of hard inclusions. Upon finding the motion strategy which maximizes the structure of the information, it is then possible to look at the taxel activation pattern to discriminate which taxels were more or less *useful* in the discrimination. The morphology can then be redesigned to better capture the discriminative taxel activation pattern observed. Repeating the experiments with the new morphology generates a new optimal motion strategy. At this point, it is possible to re-iterate, and stop only when a convergence of motion strategy vs. morphological design is observed. Algorithm 1 shows the pseudo-code for the morphology & motor-coordination iterative optimization concept in the example of robotics palpation.

```

Initialize probing end-effector design;
Probe the phantom;
Find optimal motion strategy;
while not converged do
    | Re-design probing end-effector based on taxel activation;
    | Probe the phantom;
    | Find optimal motion strategy;
end
return Morphology & motion strategy;

```

Algorithm 1: Morphology & motor-coordination iterative optimization

Convergence could here be when, for example, the observed most discriminative taxels observed do not change in two consecutive experiments. The probing end effector re-design could instead be both physical, e.g. a probe shape or material change, virtual, e.g. select the location and number of taxels most appropriate for the interaction at each given time, or a combination of both. In this Chapter we focus on the co-optimization of morphology and action together, discarding the ordering constraints imposed in algorithm 1.

10.1.2 Bayesian Exploration for Morphology-Action Co-Optimization

As previously explained, the sensor morphology and action control in robots can affect the characteristics of the sensed stimuli, and consequently the way in which robots ‘perceive’ the world [137, 171]. In this context, the intelligent use of (changing) morphology and control can be fundamental in aiding sensory perception tasks, like object discrimination [106, 227].

To date, the role of sensor morphology and action in encoding and categorizing touch stimuli remains a significant challenge. The interpretation of the sensor signals to discriminate between a set of stimuli or to perform object recognition has relied mainly on supervised machine learning techniques [72, 113, 252], burdening solutions with the need of expensive computation and large amount of labelled data. In Chapter 6 it was shown how through the use of elastomeric filters as an interface layer between a tactile sensor and the environment, it was possible to perform complex object discrimination with simple clustering analysis, voiding the need for complex learning procedures, and offloading part of the task resolution to the body. The concept of changing morphology has previously been explored [159, 205, 216], mainly in the context of growth. Here, however, we focus on driving the change in sensor morphology and action through sensory perception, thus endowing the robot with the ability to autonomously explore its own morphing and motor control abilities and adapt to different categorization tasks.

The main contribution of this chapter is to propose a conceptual framework where, through the use of Bayesian Exploration, a robotic agent is capable of exploring the perceptual tactile consequences of both changing sensor morphology and robot control action concurrently. Bayesian Exploration has previously been proposed in [64], and applied for accurate identification of textures and objects in [274]. In this work, additionally to the robot control action, the framework also accounts for the parametric exploration of the robot’s soft morphing abilities, to improve detection in complex tactile object discrimination scenarios. To demonstrate this approach, we have developed a set of 8 objects presenting three main surface feature differences, i.e.: geometric (edged vs. non-edged), texture (smooth vs. rough) and elasticity (stiff vs. soft) (Fig. 9.2).

Firstly, we show how without appropriate morphology, discrimination is often highly non-linear or impossible. Secondly, we show that through the developed framework, the robot is capable of reason probabilistically about the consequence of its own actions, as well as its own morphology, to its sensory perception. The robot is thus capable of meaningfully search its own morphing and action abilities, and avoid the need for expensive systematic search methods. To our knowledge, this is the first application of Bayesian Exploration to enable morphing based on sensor stimuli, and marks a step towards the creation of robots capable of using morphology to actively aid in discrimination tasks.

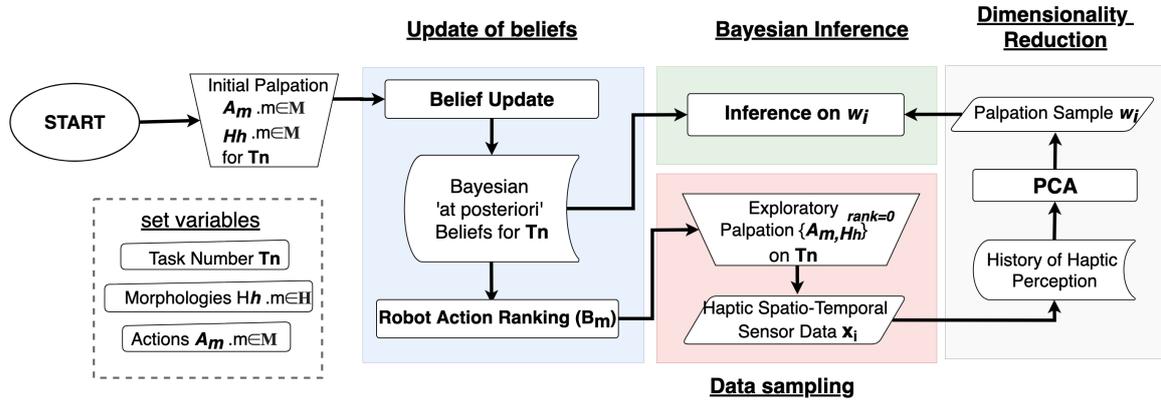


Fig. 10.1 Flowchart of the developed framework.

The chapter is organized as follows: In Section 10.2 we describe the methods in this chapter, including the implemented morphological Bayesian Exploration procedure in Section 10.2.1, the sensor technology in Section 10.2.3 and the set-up for the experiments in Section 10.2.2. In Section 10.3 the experimental results are presented. Finally, Section 10.4 we provide a discussion and a conclusion.

10.2 Methods and Experimental Set-up

10.2.1 Morphological Bayesian Exploration framework

Bayesian Exploration is an iterative procedure, which can drive the exploration of the robot's morphing and action parametric space within a pre-set task. The proposed framework is comprised of 4 stages: data sampling, dimensionality reduction, Bayesian inference and update of beliefs, and exploratory action identification (Fig. 10.1). In the last phase, the Bayesian exploratory action identification algorithm implemented is an extension of the one first proposed in [64], to account for morphology exploration during experiments.

Data Sampling

Let X be an $(N \times D)$ matrix, where each unique D dimensional row in the matrix is a sequence of tactile images for a touched object, sampled at a constant time interval. The value of N is initially 0 and for each 'experiment iteration' $N = N + K$ where K is the number of classes, or object features to discriminate against.

Dimensionality Reduction

After gathering tactile evidence for different objects and obtaining the tactile image sequences matrix X , Principal Component Analysis is used to reduce the dimensionality of the high dimensional spatiotemporal touch evidence.

After obtaining the tactile image sequences matrix \mathbf{X} , like in previous chapters we use equations 6.1, 6.2, and 6.3 to create a scatter matrix \mathbf{S} of \mathbf{X} , and factorize it into matrices \mathbf{Q} and Λ . Similarly to Chapter 6, the \mathbf{Q} matrix is such that each column q_j corresponds to an eigenvector of \mathbf{S} , and each element λ_{jj} in the diagonal matrix Λ is its corresponding eigenvalue. Furthermore, we proceed to form a 2D projection matrix \mathbf{P} with the two eigenvectors corresponding to the two highest eigenvalues in Λ , and use equations 6.4 and 6.5 to form an $(N \times 2)$ matrix \mathbf{W} , where each row in the matrix is a 2-dimensional *encoding* of a tactile image sequence for a touched object. Similarly to chapter 6, reducing the number of dimensions is beneficial in maintaining only the relevant information for object discrimination. In this context, one dimension might be too low to capture the separation of classes across the 8 different objects, while dimensions higher than 3 might not induce robust clusters (due to the fluctuations of the tactile sensor response over time across experiments).

Bayesian inference and update of beliefs

As the robot touches an object, the type of surface under touch C , the type of robot control action A and the sensor morphology H generate an observable sensor measurement w_i . The likelihood that a specific surface $C_k \in C$ has generated the haptic observation w_i can thus be computed as:

$$P(C_k|w_i, A_m, H_h) = \frac{P(w_i|C_k, A_m, H_h)P(C_k)}{P(w_i|A_m, H_h)} \quad (10.1)$$

H_h is a particular morphology and A_m is a specific touch control action. According to the central limit theorem we can approximate the conditional probability of observing w_i , with the probability density function $p(w_i|C_k, A_m, H_h)$, defined by a mean $\vec{\mu}_{k,m,h}$ and a covariance matrix $\Sigma_{k,m,h}$:

$$P(w_i|C_k, A_m, H_h)P(C_k) \propto p(w_i|C_k, A_m, H_h) = \frac{1}{\sqrt{(2\pi)^2 |\Sigma_{k,m,h}|}} e^{-\frac{1}{2}(\vec{x} - \vec{\mu}_{k,m,h})^T \Sigma_{k,m,h}^{-1} (\vec{x} - \vec{\mu}_{k,m,h})} \quad (10.2)$$

We will refer to the set of densities for all morphology-action pairs as the belief state of the robot

As the robot forms a belief state, it is possible to perform Bayesian inference with respect to a specific morphology-action pair by simply evaluating a new, unseen, sample w_j under the Gaussian densities $p(w_i|A_m, H_h)$, for each object under that morphology-action pair. The Gaussian with the highest density at w_j is the most probable class for the unseen sample under consideration.

Exploratory action identification

We use Bayesian Exploration to identify the exploratory morphology-action pair necessary to update the beliefs of the robot [64]. The estimate of the morphology and the control action which is most likely to discriminate best amongst different object features is the one which minimizes the discriminatory confusion amongst all possible classes under a specific morphology-action pair. One possible measure of confusion between probability density functions is the amount of overlap between them. Like previously we use the Bhattacharyya coefficient [203] to compute a confusion probability matrix $\Psi_{ks,m,h}$ for each possible exploratory action control A and morphology H . Each element in $\Psi_{ks,m,h}$ is a mutual confusion between any two classes C_k and C_s , and can be computed as:

$$\Psi_{ks,m,h} = \int \sqrt{p(w_i|C_k, A_m, H_h)p(w_i|C_s, A_m, H_h)} \quad (10.3)$$

To make the computation possible within the framework we assume normal probability densities in the belief state, reducing the computation to:

$$\Psi_{ks,m,h} = \sqrt{\frac{2\vec{\sigma}_{k,m,h}^2 \vec{\sigma}_{s,m,h}^2}{\vec{\sigma}_{k,m,h}^2 + \vec{\sigma}_{s,m,h}^2}} e^{-\frac{(\bar{\mu}_{k,m,h} - \bar{\mu}_{s,m,h})^2}{4\vec{\sigma}_{k,m,h}^2 + 4\vec{\sigma}_{s,m,h}^2}} \quad (10.4)$$

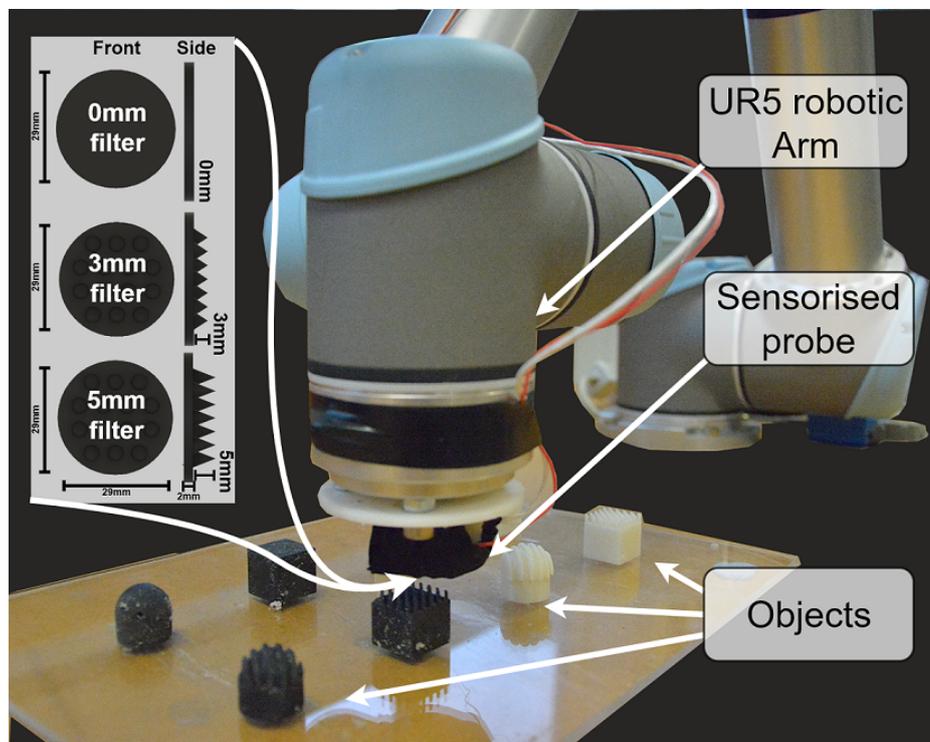
where $\vec{\sigma}_{k,m,h}$ is the diagonal vector of $\Sigma_{k,m,h}$. The Ψ probability confusion matrix can be used to find the benefit of making an exploratory action A_m with a sensor morphology H_h . We define two different benefit estimation equations: an unbiased benefit estimation $\hat{\mathbf{B}}_{m,h}$, and a biased exploratory benefit estimation $\mathbf{B}_{m,h}$. The unbiased benefit estimation for action A_m and morphology H_h can be computed as:

$$\hat{\mathbf{B}}_{m,h} = \sum_k^K \frac{P(C_k)^2}{\sum_s^K \Psi_{ks,m,h} P(C_s)} \quad (10.5)$$

And its value will be higher for control actions with class probability density functions with least overlap. The ‘confusion’ of using a sensor morphology H_h when making an exploratory action A_m is thus $\widehat{\mathbf{B}}_{m,h}$. Furthermore, we define the biased benefit estimation as:

$$\mathbf{B}_{m,h} = 1 - (1 - \widehat{\mathbf{B}}_{m,h})^{\frac{1}{n_{m,h}}} \tag{10.6}$$

where $\frac{1}{n_{m,h}}$ is the number of times the robot used morphology H_h and action A_m to touch the objects during experiments. The biased benefits are discounted by the number of times the morphology-action pair has already been performed during action exploration, to discourage



(a) Robot set-up

	Object 0	Object 1	Object 2	Object 3	Object 4	Object 5	Object 6	Object 7
Edged	✓	✓	✓	✓	✗	✗	✗	✗
Smooth	✓	✓	✗	✗	✓	✓	✗	✗
Stiff	✗	✓	✗	✓	✗	✓	✗	✓

(b) Designed objects and object properties.

Fig. 10.2 Experimental set-up for Morphology-Action Co-Optimization.

excessive exploitation and eventually encourage the exploratory update of belief states under less exploited morphology-actions. The equations and scores were modified to account for the co-optimization of, not only the action, but also the morphology of the robot.

10.2.2 Experimental Set-Up

We set up experiments to allow the robot to improve its feature discriminative abilities by co-optimizing morphology and robot control action. The touch experiments were performed using a UR5 robot arm, equipped with a probe sensorised with a capacitive tactile sensor array (Section 10.2.3). Fig. 10.2 shows the experimental setup developed for both the experiments. To modify the sensor morphology, three different dielectric elastomeric layers were explored, each 3D printed with VeroBlack PolyJet Rubber, and presenting a thin circular layer of $2mm$, as well as conical protrusion, spaced $2mm$ from each other, and with varying height of $0mm$, $3mm$ and $5mm$ (Fig. 10.2a).

We thus create a set of 8 objects, differing in three sets of features. Fig. 10.2b shows the objects designed for the experiments. Each object is circumscribed by a $20mm \times 20mm \times 30mm$ cuboid, while we simulate roughness by reproducing $3mm$ protrusion spaced at $2mm$ onto the object's top surface. Following the object design it is possible to classify the objects over three sets of different salient features, i.e.: round vs edged objects, objects with rough vs smooth surfaces, and stiff vs non-stiff. The non-stiff objects (objects 0, 2, 4, 6 Fig. 10.2b) were 3D printed with VeroBlack rubber. The stiff objects (objects 1, 3, 5, 7 Fig. 10.2b) were 3D printed with rigid PLA material.

Each touch experiment consisted of 2 seconds of contact between the sensorised robot end-effector with a target object. We manually taught the robotic arm the x-y location of each object within its work-space, and set the robot starting position with the end-effector aligned normally to the upper surface of the objects. The end-effector was thus driven downward until a touch event was detected by the capacitive tactile sensor at its extremity, whereby the touch experiment would begin (Fig. 10.3a). The robot was controlled in Cartesian coordinates at $\approx 60Hz$, acting upon the X, Y and Z tool axis simultaneously. Distinct sinusoidal displacements profiles were generated for every axis, each of which was controlled in amplitude and frequency parameters, thus a total of 6 parameters were used to control the robot for each touch experiment, i.e. A_x , A_y , A_z , ω_x , ω_y and ω_z (Fig. 10.3a).

The A_x , A_y , ω_x and ω_y parameters were set to allow the robot to rub the surface of the objects on a $10mm$ radius from the center, thus $A_x = 3$, $A_y = 3$, $\omega_x \in [-0.0025, 0, 0.0025]$ and $\omega_y \in [-0.0025, 0, 0.0025]$.

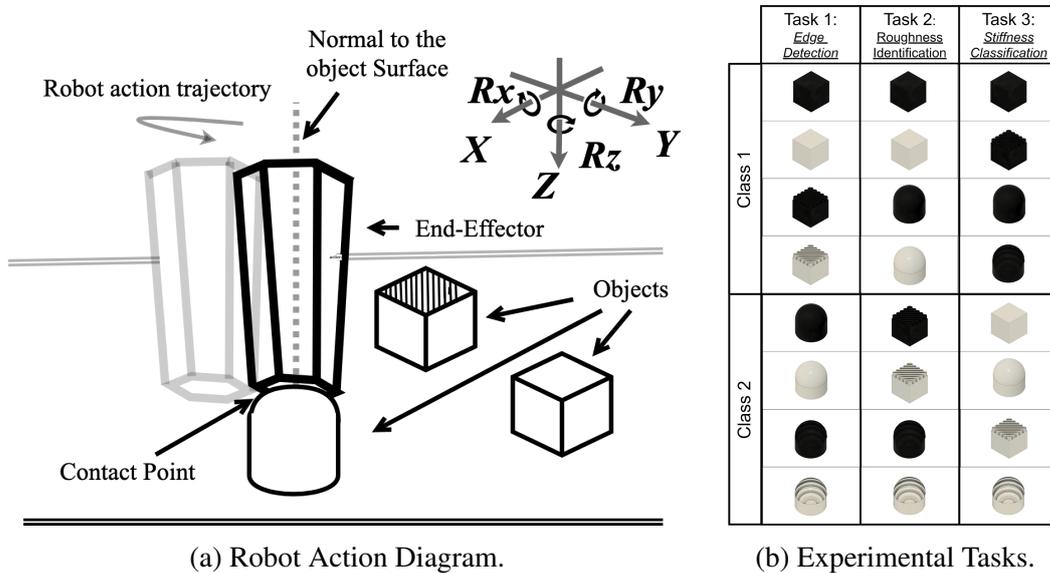


Fig. 10.3 Robot control action and experimental tasks.

		Parameters	Test Acc. (%) (Avg. Test Acc.)
Task 1	Action Control	$A_x = 3, \omega_x = 1$ $A_y = 3, \omega_y = -0.0025$ $A_z = 1, \omega_z = 0.001$	80 % (31.1 %)
	Morphogy	5mm	
Task 2	Action Control	$A_x = 3, \omega_x = 1$ $A_y = 3, \omega_y = -0.0025$ $A_z = 1, \omega_z = 0.001$	100% (47.2%)
	Morphogy	3mm	
Task 3	Action Control	$A_x = .05, \omega_x = 1$ $A_y = 1, \omega_y = 1$ $A_z = 1, \omega_z = 0.001$	75% (29.81%)
	Morphogy	3mm	

Table 10.1 The table shows the highest test accuracy achieved through the best performing morphology-action pairs after gathering 10 sample evidence of three object for each task category, and testing on 10 tactile samples for a new object.

10.2.3 Sensor Technology

To endow the robot with tactile sensing ability we mount a capacitive tactile sensor, developed in [152, 154] to a custom 3D printed end-effector. The sensor has been integrated into a number of existing robotic systems which exploit sensory-motor co-ordination [229, 230]. The utilized module has a layered structure consisting in a Flexible Printed Circuit Board (FPCB), a dielectric layer and conductive lycra which act as common ground plane for all the taxels and constitute the second plate of the capacitor. The FPCB hosts 7 tactile elements (Taxels), corresponding each to the first plate of a capacitor, and a Capacitance to Digital Converter (CDC AD747 from Analog Devices). The sensor technology and architecture is as described in Chapter 4.

10.3 Results

10.3.1 Morphology and Action for Object Classification

We assess whether any meaningful filtering can be performed by changing sensor morphology, so to be able to classify each object based on the three tasks: round vs edged objects, objects with rough vs smooth surfaces, and stiff vs non-stiff objects. Systematic touch experiments are performed by varying the robot morphology and action control via every possible action-morphology pair. Each experiment is performed 10 different times, to provide sample evidence for the density distributions in the robot belief state. We can thus test the ability of the robot to classify objects based on the respective task features by forming the density distributions on three of the objects within each feature set, and performing Bayesian inference (Section 10.2.1) on a random left-out object within it. Table 10.1 shows the accuracy achieved for all attempted discriminative tasks. The best morphology-action pairs can achieve accuracy higher than 75% on all tasks.

More interestingly, without an appropriate combination of morphology and motor control, it is almost impossible to discriminate objects based on their geometrical, surface roughness or stiffness properties, as shown by the average performance per task by any one pair. Fig. 10.4 shows the ranked morphology-actions with respect to the unbiased $\hat{\mathbf{B}}_{m,h}$ benefit estimator for Task 2, roughness identification. The figure shows the relationship between the probability densities formed in the robot belief state, under the developed framework. Highly ranked morphology-action pairs (e.g. Rank 0 and Rank 1) show Gaussian distributions which more easily discriminate between different features, while lower ranked pairs present more distributional overlap, and thus higher degrees of discriminative confusion. Interestingly, although action control can reduce the distributional overlap, the morphology ultimately

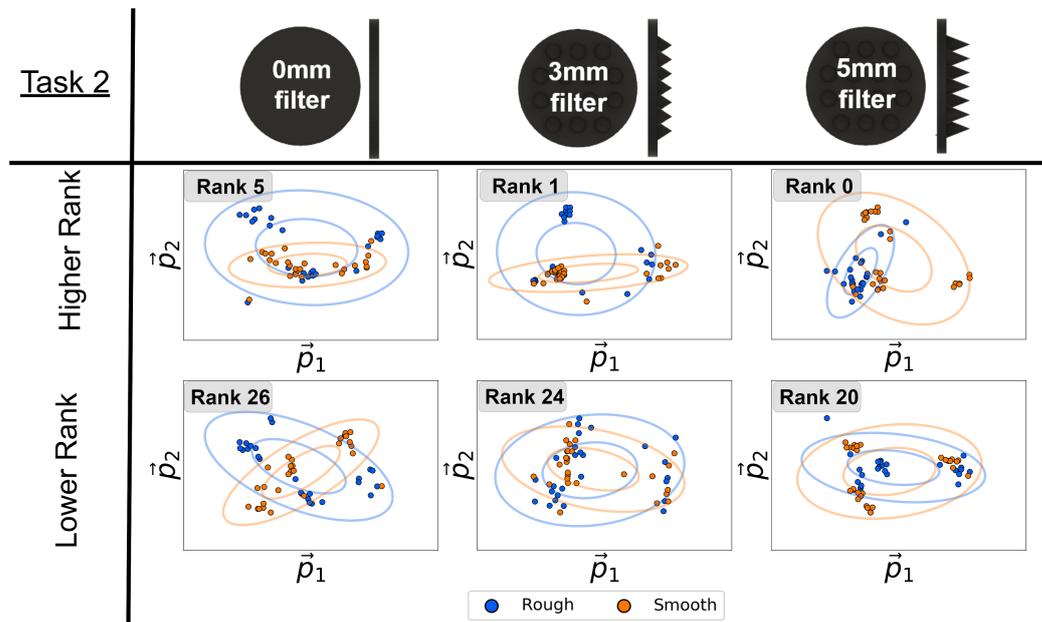


Fig. 10.4 The robot belief state after 200 iterations for Task 2 (Roughness Identification), under 6 different morphology-action pairs.

enables accurate classification (e.g. the distributions of pair Rank 5 vs. pair Rank 0). Varying the morphology, in fact, ‘filters’ the tactile response [227], inducing sensory differences between objects of varying surface roughness, and enabling discrimination. Fig. 10.5 shows the ranked morphology-action pairs after approximately 200 iterations on Task 2. The distributional differences between highly and lowly ranked motion are evident within the row sensor data, with morphology-action pairs inducing the sensor data for objects of different classes to be increasingly more distinct (e.g. Rank 0 vs. Rank 26).

10.3.2 Morphological Bayesian Exploration

We test the Bayesian Exploration framework for morphology-action optimization by running non-systematic experiments and comparing the results to the previous findings. Under the Bayesian exploration framework, the robot is made to touch each object under every morphology-action pair only once, to form an initial belief state. From then on, the robot decides which morphology-action pair to gather additional evidence for, based on the biased estimator ($\mathbf{B}_{m,h}$). Fig. 10.6 shows the maximal accuracy achieved by the robot during the systematic and Bayesian exploratory experiments. The figure shows how Bayesian exploration consistently outperforms the systematic search over the robot morphology-action parameters, finding good configurations in about half the time necessary to systematic methods.

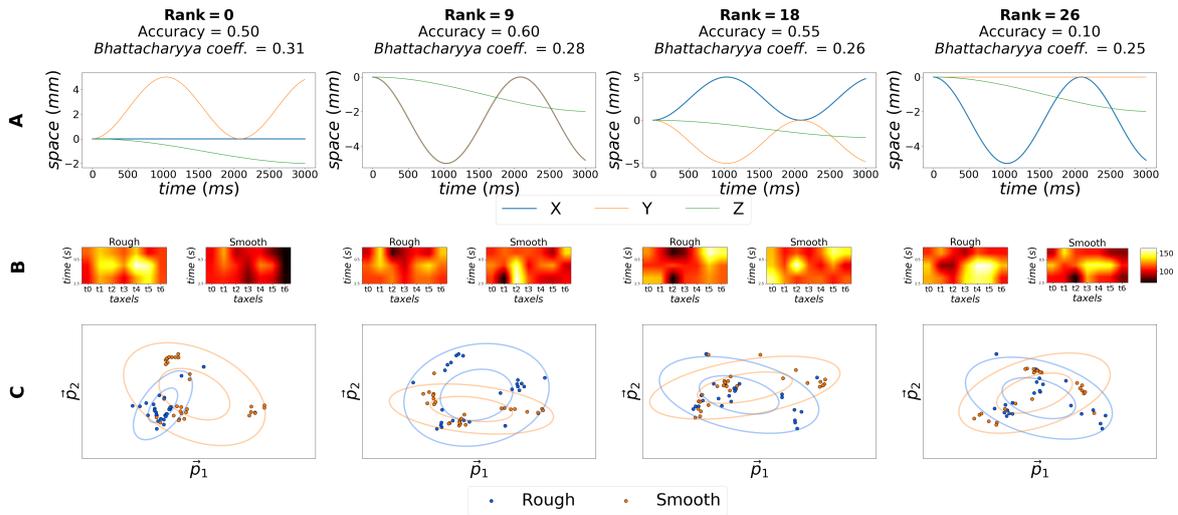


Fig. 10.5 The ranked morphology-action pairs after approximately 200 iterations on Task 2. Row A shows the action control employed by the robot. Row B shows example raw time series sensor data for each class within Task 2, where the taxels (x-axis) show brighter or darker shades over time (y-axis) depending on the sensed pressure. Finally, row C shows the 2D Gaussians in the belief state of the robot under each action-morphology ranked pair.

The fast configuration finding is due to the confusion-driven exploration based on sensor evidence. Assuming distributional consistency amongst sensor values generated under the same conditions, the robot will make informed decisions on which evidence to gather to discriminate the objects with the least amount of confusion at each iteration. The lower accuracy values per task in Table 10.1 Section 10.3.1 suggest the possibility of overfitting on the objects under touch. Early stopping through cross validational methods can here be used to halt robot training and prevent overfitting.

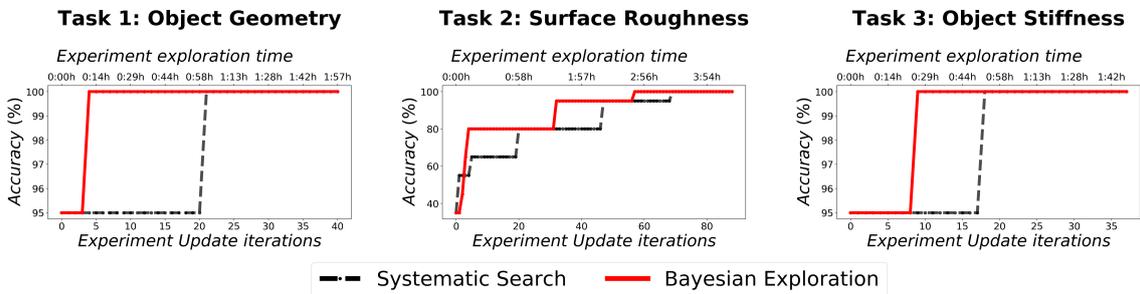


Fig. 10.6 The figure shows the highest running accuracy for robot throughout the experiments on each task. The developed Bayesian exploration framework consistently out-performs systematic search of morphology-action parameters.

10.4 Discussion and Conclusion

The importance of morphology and sensory-motor coordinated action has been emphasized in the past few decades. In this work, we proposed a Bayesian Exploration framework for a robot to co-optimize the morphology and robot control action to perform object discrimination based on salient features. We show that appropriate control action can aid in object discrimination tasks [227]. More radically, we show that morphology is necessary to enable classification in complex scenario, as it is the case for tactile roughness estimation with our sensor. The appropriate morphing of the dielectric layer, in fact, induces the sensor response to the touching of rough and smooth surfaces to be easily classifiable through Bayesian inference methods. Depending on the employed sensor morphology, instead, almost independently from the control action, the extreme overlap of sensor evidence makes discrimination poor, if not impossible at times.

The Bayesian Exploration framework allows for a reduction in the exploration of parameters, and thus facilitates the real-world parametric exploration of the robot morphing and action capabilities. We show that the robot is capable of finding good morphology-action configurations in approximately half the time necessary to systematic search approach, for each of the attempted tasks.

This chapter is the last of a set of chapters involving the employment of morphology and action to influence the ability of the robot to learn, and perform particular tasks. As such, the chapter combines both morphology and action into a unique framework where neither would be able to solve the task at hand, but the combination of both can achieve the desired results. This approach, however, has several limitation. One, for example, is the necessity of parametrization and discretization, mentioned in previous chapters. Another lies in the temporal equity of morphology and action changes. As both are allowed to change at the same time, progress can only be made by exhaustive search of both until a good joint solution is found. Additional benefits might be achieved by inducing changes in morphology and action in different temporal timescales. More details of these issues and future directions will be provided in the next chapter.

Chapter 11

Conclusion and Future Work

11.1 Conclusion

Abandoning the rigidity constraints imposed in the robotics of the previous century, Soft Robotics suggests a new way of thinking about robotics systems. Leveraging on the possibility of achieving new bodily properties and behaviours, soft robots must learn to leverage on the complexity of their own body, and the underlying complex interactions with the environment surrounding them, to their advantage. This complexity can be leveraged via the use of learning processes, however learning must be grounded in the real world, where learning processes depend first and foremost on the quality of the physical stimuli arising between the soft system-environment interactions. Throughout the Chapters in this thesis we have discussed four principles, which constitute *Soft Morphological Computation (SoMComp)*, i.e.: Soft Morphology, Soft Actuation, and Soft Sensing and Soft Proprioception. These principles allow a robot to structure the sensory stimuli purposefully, so as to improve perception and learning.

After demonstrating a use case of complexity in the soft interactions between a robot and its environment in Chapter 3, we treat the principle of soft proprioception in Chapter 4. A robot, like a biological organism, must have an idea of *bodily self*, or otherwise knowledge of the properties of its own body (Soft Proprioception), so to be able to influence its perceptual information. Soft Proprioception becomes a core part of *SoMComp*, where with the knowledge of its own body as well as its own actions, the robot can appropriately condition the physical stimuli, arising from the interaction with the environment, to its advantage. In other words, to perform *SoMComp* Soft Proprioception is first necessary. Soft Proprioception, however, must not necessarily be bestowed upon the robot by human design. But rather the robot can achieve Soft Proprioception autonomously, and use this at later stages within the framework of *SoMComp*. In the chapter we propose a sensorization method to

achieve proprioception via tactile sensor arrays, without a known model. The robot, although initially without knowledge of its own body, can achieve proprioception through movement akin to ‘twitching’ in infants, where random jerky motions reinforce the robot understanding of its own body given sensor evidence. We show the method is precise and reliable, but further issues relate to scaling to larger systems, with even more complex dynamics.

After introducing the concept of morphology and its importance in achieving appropriate manipulation of soft and delicate objects in Chapter 5, in Chapter 6 we discuss the concept Soft Morphology. One of the two principal ways to condition soft interactions in *SoMComp*, in fact, relies on the ability of the robots of the future to change their bodily properties. Changes can here be of countless types, growth, tool use or body-parts replacements are all examples of possible such changes. In Soft Robotics, changes in stiffness, elasticity, or changes due to heating, deterioration or self-healing, can all contribute to bodily changes a robot should take into consideration. But if changes within the body of a robot are unavoidable, and indeed necessary, then why not change bodily properties purposefully to aid in solving required tasks? The paradigm of Soft Morphology revolves around the concept that morphology can influence sensory stimuli before such stimuli reach any classic information processing stage, and that this influence can be useful for computation. In the chapter we use unsupervised methods to show how morphological computation can indeed aid in sensory perception. We thus show how a complex tactile classification task can be performed with simple unsupervised clustering methods, if appropriate changes are induced in the tactile apparatus. This goes against common approaches in the field, where more and more complex machine learning techniques are used to solve complex problems. Complexity in the computation, here, is instead off-loaded to the body itself. The brain, in this case, would be free to instead resolve higher cognitive tasks.

In Chapter 7 and 8 we treat the topic of Soft Actuation. In fact, changes in the morphology of an agent are not the only way to influence physical stimuli. If once the morphology changes, the stimuli are influenced accordingly, it is also true that the way a robot acts in the world also has a similar influence. Another powerful mechanism thus comes from the physical interaction between a robot and the environment the robot is situated in, as well as itself. In Chapter 7 and Chapter 8 we take into consideration a complex tactile sensory perception task, i.e. medical palpation. A robot is thus made to palpate a silicon phantom organ with hard spherical shaped inclusions as small as $5mm$ in diameter. To enable the robot to perform classification it is necessary for the robot to autonomously find a palpation strategy capable of enhancing its own tactile sensing capabilities to perform the task at hand. In Chapter 7 we show how coordinated action, indeed, influences sensory perception and enables palpation. In Chapter 8 we also show how complex actions are necessary to condition information

appropriately in complex scenarios. Moreover, we devise a probabilistic framework for a robot to explore high dimensional action spaces efficiently, based on sensor evidence. It is therefore the sensory perception itself which drives the robot to change its interaction strategy, and ultimately efficiently find a motor action capable of enhancing perception to solve the palpation task.

In Chapter 9 we move away from classification to show how the concepts can be used instead in a regression case. Here, a robot is made to approximate 10 different playing styles via appropriate key-stroke interactions with a piano instrument. The complexity of the interactions between a non-rigid finger and the instrument, to produce complex and dynamic music pattern is a complex one, but it is addressed by a proposed probabilistic framework, based on Gaussian Processes modelling. The framework can capture the complex relationship between the music patterns and the complex key-stroke actions generating them, affectively allowing the robot to out-perform a seasoned human player on 4 out of 10 styles, and perform comparatively on 2 additional styles. The work further addresses the concept of action for perception, where perception is now sound-feedback, the action is continuous, and the task is regression, and not classification.

Finally, we merge together the two concepts of Soft Actuation and Soft Morphology in Chapter 10, further addressing the problem of high-dimensional space search. If it is true that a robot should be able to change its own bodily properties according to the task to solve, then its interaction strategy with the environment and itself must change with the changing body, else the improvement achieved through the changes in one may be hindered by the inflexibility of the other. It is here that we propose a probabilistic approach to efficiently reason about the changes in perception due to both changes in the morphology and the action strategy employed by the robot to solve a perception task. We focus on tactile perception for the discrimination of objects based on salient features. The chosen features are shown to be otherwise undetectable if not by appropriate changes in both morphology and action, and the efficient autonomous exploration of both.

11.2 Future Work

The concept of *SoMComp* proposed in this thesis hinges on the idea that if a robot has an idea of bodily self (soft proprioception), then the robot should change its bodily properties (soft morphology) or actions (soft actuation) to actively affect the soft interactions arising within its environment. These changes, in turn, influence each other, and thus a continuous morphing and development is iteratively necessary to achieve stable optimal behaviours. The chapters throughout this thesis have shown how it is possible to intervene at the level of soft

interactions, to induce statistical regularities in the stimuli to make the resolutions of the robot task simpler, faster or more robust. Like shown in core use-cases, the appropriate use of these concepts through the developed frameworks can improve, and indeed enable, robots to solve complex tasks, requiring augmented sensory capabilities or complex post-processing and learning paradigms. Several directions are interesting whilst considering each of the principles proposed in this thesis.

11.2.1 Ubiquitous Proprioception

Soft Proprioception is set to be the necessary, if not sufficient condition for *SoMComp* to take place. In Chapter 6, for example, proprioceptive knowledge included robot kinematics and the properties of the soft filters. In chapter 7, Chapter 8 and 10 kinematics and knowledge of the end-effector properties were necessary to devise appropriate actions to achieve *SoMComp*.

As such, achieving proprioception in Soft Robotics is fundamental. The physical dynamics of soft materials is such that simulation is hard, and often impossible. In Chapter 4, a model-free sensorization approach was proposed to achieve proprioception of a continuum soft robot, and tactile sensing was used as the main technology towards this goal. The approach proposed, however, was also limited *because* it did not consider any bodily physical constrain. For example, let's take a similar scenario where the soft continuum robot had been composed by several sections, each similar to the one described for the experiments. The tendons actuating the continuum soft robot would be more in number, each pulling a section at different heights within the finger. In this scenario, the complexity would have been high enough that perhaps only several tens of hours of continuous “motor bubbling” would have achieved an appropriate mapping between the end-effector position and the tactile images retrieved at the base of the finger. Moreover, sensors at different levels would most likely be necessary to discriminate between complex robot poses.

One of most significant ways to obviate these issues is by appropriately combine model free and model-based methods, to both leverage on human knowledge of the laws of the world through models, and the flexibility of model-free approaches, where the model themselves are moulded based on sensory evidence, and can change and adapt over time. In the Soft Robotics of the future it is also going to be fundamental to achieve ubiquitous proprioception, where sensing is distributed redundantly throughout the body of soft robots, and learning can be used to self-organize the physical stimuli to achieve proprioception at different scales. To achieve ubiquitous proprioception, it is necessary to make progress in several areas. Firstly, it is necessary to design sensors capable of coexisting in soft structures, and which can stretch and bend at will. Much research has been done in this direction, as detailed in Chapter 2. However, technologies are yet to reach a breakthrough, where seamless integration of sensors

within a soft body is possible. Secondly, it is necessary to merge model-based and model-free control schemes into a unique framework. In this framework, proprioception should benefit from human design, by embedding human knowledge within a model, however the model should change over time, as dictated by an ‘online’ learning framework.

11.2.2 A Comprehensive Theory of Robotic Learning

The concept of *SoMComp* is a powerful tool, one which can improve robot learning by taking embodiment into account. The way we have discussed the achievement of this goal, however, is through the conditioning of physical stimuli by the robot, while leveraging on existing learning frameworks to achieve the robot task.

One of the most significant step forwards from this, is the development, instead, of a comprehensive theory of learning in robotics. This theory should embrace embodiment as the means through which the information can be conditioned appropriately for learning. The closest existing framework is perhaps reinforcement learning, where the robot actions and repercussions in the world are considered. Reinforcement learning, however, does not often consider the physical embodiment of the robot, where morphology and soft interactions have a deeply rooted effect on the quality of the physical stimuli the robot can perceive. A comprehensive theory of learning would embed the *SoMComp* concepts within the learning loops, where the robot is concurrently thinking about “how to do a task” (traditional learning), and “what to do to learn it better” (*SoMComp*).

A concept related to this theory concerns learning timescales. If it is possible for a robot to change its morphology, as well as its action, to condition information appropriately, then it is not necessarily true that these changes must happen concurrently. In fact, it is possible that a change in morphology alone would improve the ability of the robot to perform a task, while the same morphology with an additional change in action would worsen it. In Chapter 10 this was obviated by Bayesian Exploration, which would eventually be exhaustive in its search for new jointly optimal action-morphologies. In a more generic scenario, however, it might be necessary to induce temporal differences in the learning of morphology and actions. For example, in robotic manipulation it may be more ‘expensive’ to change the physical characteristics of a body (i.e. change the viscosity of the fingertips), while a change in grasping action would be less ‘expensive’. The robot could improve grasping of an object by learning on two different timescales. The first, a faster timescales where changes in action dictate immediate benefits from robot grasping. The second, a slower timescales where small morphological changes are induced in the manipulator to aid the action change. This would allow changes in the morphology to co-exist with changes in grasping action, and neither to influence each other negatively.

Finally, general object grasping and manipulation is one problem where the concepts of *SoMComp* can be most useful. This is mainly because of the physical soft nature of the interactions arising in real-world grasping scenarios, and for the need of using appropriate morphology and action to achieve appropriate grasping.

11.2.3 Limitations and Future Remarks

Moving forward, it is important to consider several factors to make progress in this area of research. Two of the biggest limitations of current frameworks, including the ones described in this thesis, are the need for both parametrization and discretization of continuous domains. We have encountered this throughout the chapters. In Chapter 6, for example, the properties of the ‘soft’ filter to use for tactile object discrimination needed to be first parametrised in width, and then discretized into specific widths to explore. In Chapter 7 and 8 the robotic trajectories needed to be parametrized along the axis of motion of the robot, and the search over each parameter was performed over discrete values of the same. In Chapter 10 both the properties of the dielectric used and the action strategy needed to be designed appropriately, parametrized and searched discretely. The automatic design and parametrization is an important future research topic, one which would allow systems to adapt their morphology and action appropriately without, or with minimum, human intuition and intervention. To this extent, several research directions are relevant, including, for example, Evolutionary Robotics. This current directions and future limitations of this research direction were discussed at length in Chapter 2.

11.2.4 Additional Work Beyond This Thesis

Several related publications could have been included in this thesis. The content of each chapter was chosen mainly to keep the content succinct while providing a rounded accounted of the concept of *SoMComp* and its applications. Additional publications include [86], [101], [99] and [233].

In [86], we propose a method to sensorize a soft phantom for the purpose of medical palpation. This method is shown to be able to locate the point of contact, as well quantitative information about the palpation trajectory employed, including the force applied, and the type of trajectory (e.g. rotational or normal). A sensorized phantom can take the concepts treated in this thesis one step further, where the optimization of trajectories is not egocentric, i.e. related only on the robot’s perception of self, but also on the perception of the states of entities around the robot. This is analogous to a medical practitioner adjusting their palpation technique based on the body language of the patient, which may experience pain

or discomfort. As a co-author in the paper my contribution included data collection, robot control for multi-axis palpation strategy, and paper review. The article is forthcoming in Journal IEEE Transactions on Robotics.

In [99] we provide approaches to solve real-world problems by using morphology and the robot's bodily properties as the means through which to achieve progress. The accounts provided in the paper are the results of our entry to the "World Robotics Summit", an international robotics competition held in Tokyo in 2018. In the article we describe several new ways to approach common hard problems. For example, we show how to achieve robot picking, placing, manipulation and use of small objects, such as screws, by taking advantage of the adhesion properties of some viscous materials (e.g. grease). As co-author in the paper my contribution included the participation to the competition, the design and implementation of the learning and vision technologies for the competition, data collection, data analysis and article writing. The article was published in the Journal Intelligent Service Robotics.

In [101], we provide a review of current trends in the area of agricultural robotics and focus on future technologies which can revolutionize the field. This is with particular focus on the difficulties of this area, mainly due to the soft interactions arising in harvesting, and post-harvest manipulation in agricultural settings. In this context, soft robotics technologies may be key in achieving compliance, avoid damage and perform these tasks appropriately, and as such, *SoMComp* can be key to achieving appropriate robotics solutions. The article is a forthcoming book chapter in a field robotics book to be published by Elsevier. As a co-author in the review chapter my contributions include the design of the chapter structure, and the writing of the chapter.

Finally, in [233] we propose a method to achieve self-supervision in robotics, by exploiting the dynamics of the objects within the field of vision of the robot. A human operator is made to act upon the object in a conveyor belt, by removing any one object based on a specific property. The robot can observe the conveyor, and upon human intervention automatically assess human intention, and take over the selection process for the objects. One of the set-backs of the methods used in some of the chapters in this thesis lied with the need of humans to provide explicit supervision. This is a common problem in Machine Learning as well. This work focuses on technologies to move away from supervised learning to technologies which instead take into account the world surrounding the robot, such that self-supervision is possible.

References

- [1] Judith A. Abbott. Quality measurement of fruits and vegetables. *Postharvest Biology and Technology*, 15(3):207 – 225, 1999. ISSN 0925-5214. doi: [https://doi.org/10.1016/S0925-5214\(98\)00086-6](https://doi.org/10.1016/S0925-5214(98)00086-6).
- [2] Ehud Ahissar and Eldad Assa. Perception as a closed-loop convergence process. *Elife*, 5:e12830, 2016.
- [3] Katrina Armstrong, Elizabeth A Handorf, Jinbo Chen, and Mirar N Bristol Demeter. Breast cancer risk prediction and mammography biopsy decisions: a model-based study. *American journal of preventive medicine*, 44(1):15–22, 2013.
- [4] Minoru Asada, Karl F MacDorman, Hiroshi Ishiguro, and Yasuo Kuniyoshi. Cognitive developmental robotics as a new paradigm for the design of humanoid robots. *Robotics and Autonomous systems*, 37(2-3):185–193, 2001.
- [5] Minoru Asada, Koh Hosoda, Yasuo Kuniyoshi, Hiroshi Ishiguro, Toshio Inui, Yuichiro Yoshikawa, Masaki Ogino, and Chisato Yoshida. Cognitive developmental robotics: A survey. *IEEE transactions on autonomous mental development*, 1(1):12–34, 2009.
- [6] W Ross Ashby. Requisite variety and its implications for the control of complex systems. In *Facets of systems science*, pages 405–417. Springer, 1991.
- [7] Elif Ayvali, Rangaprasad Arun Srivatsan, Long Wang, Rajarshi Roy, Nabil Simaan, and Howie Choset. Using bayesian optimization to guide probing of a flexible environment for simultaneous registration and stiffness mapping. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 931–936. IEEE, 2016.
- [8] C Wouter Bac, Eldert J Henten, Jochen Hemming, and Yael Edan. Harvesting robots for high-value crops: State-of-the-art review and challenges ahead. *Journal of Field Robotics*, 31(6):888–911, 2014.
- [9] C Wouter Bac, Jochen Hemming, BAJ Tuijl, Ruud Barth, Ehud Wais, and Eldert J Henten. Performance evaluation of a harvesting robot for sweet pepper. *Journal of Field Robotics*, 34(6):1123–1139, 2017.
- [10] Knarik Bagdasarian, Marcin Szwed, Per Magne Knutsen, Dudi Deutsch, Dori Derdikman, Maciej Pietr, Erez Simony, and Ehud Ahissar. Pre-neuronal morphological processing of object location by individual whiskers. 16:622, 04 2013.

- [11] Gianluca Baldassarre and Marco Mirolli. *Intrinsically motivated learning in natural and artificial systems*. Springer, 2013.
- [12] Andrew G Barto. Intrinsic motivation and reinforcement learning. In *Intrinsically motivated learning in natural and artificial systems*, pages 17–47. Springer, 2013.
- [13] DN Beal, FS Hover, MS Triantafyllou, JC Liao, and George V Lauder. Passive propulsion in vortex wakes. *Journal of Fluid Mechanics*, 549:385–402, 2006.
- [14] Jim H Belanger and Barry A Trimmer. Combined kinematic and electromyographic analyses of proleg function during crawling by the caterpillar *manduca sexta*. *Journal of Comparative Physiology A*, 186(11):1031–1039, 2000.
- [15] Nikolai Bernstein. The co-ordination and regulation of movements. *The co-ordination and regulation of movements*, 1966.
- [16] Luc Berthouze and Max Lungarella. Motor skill acquisition under environmental perturbations: On the necessity of alternate freezing and freeing of degrees of freedom. *Adaptive behavior*, 12(1):47–64, 2004.
- [17] Lynn S Bickley, Peter G Szilagy, and Barbara Bates. *Bates' guide to physical examination and history taking*. Lippincott Williams & Wilkins, 2009.
- [18] Elaine Biddiss and Tom Chau. Electroactive polymeric sensors in hand prostheses: Bending response of an ionic polymer metal composite. *Medical engineering & physics*, 28(6):568–578, 2006.
- [19] C Blanes, M Mellado, C Ortiz, and A Valera. Technologies for robot grippers in pick and place operations for fresh fruits and vegetables. *Spanish Journal of Agricultural Research*, 9(4):1130–1141, 2011.
- [20] G Blanpied, W Bramlage, D Dewey, R LaBelle, L Massey Jr, G Mattus, W Stiles, and A Watada. A standardized method for collecting apple pressure test data. 1978.
- [21] Jeannette Bohg, Karol Hausman, Bharath Sankaran, Oliver Brock, Danica Kragic, Stefan Schaal, and Gaurav S Sukhatme. Interactive perception: Leveraging action in perception and perception in action. *IEEE Transactions on Robotics*, 33(6):1273–1291, 2017.
- [22] Anne M Bosch, Alfons GH Kessels, Geerard L Beets, Jan D Rupa, Dick Koster, Jos MA van Engelshoven, and Maarten F von Meyenfeldt. Preoperative estimation of the pathological breast tumour size by physical examination, mammography and ultrasound: a prospective study on 105 invasive tumours. *European journal of radiology*, 48(3):285–292, 2003.
- [23] Malcom C Bourne. Fruit texture—an overview of trends and problems. *J. of Texture Studies*, 10(1):83–94, 1979.
- [24] Valentino Braitenberg. *Vehicles: Experiments in synthetic psychology*. MIT press, 1986.

- [25] Rodney A Brooks. Elephants don't play chess. *Robotics and autonomous systems*, 6 (1-2):3–15, 1990.
- [26] Eric Brown, Nicholas Rodenberg, John Amend, Annan Mozeika, Erik Steltz, Mitchell R Zakin, Hod Lipson, and Heinrich M Jaeger. Universal robotic gripper based on the jamming of granular material. *Proceedings of the National Academy of Sciences*, 107(44):18809–18814, 2010.
- [27] Thomas Buhrmann, Ezequiel Alejandro Di Paolo, and Xabier Barandiaran. A dynamical systems account of sensorimotor contingencies. *Frontiers in psychology*, 4:285, 2013.
- [28] John Butter, Thomas H Grant, Mari Egan, Marsha Kaye, Diane B Wayne, Violeta Carrión-Carire, and William C McGaghie. Does ultrasound training boost year 1 medical student competence and confidence when learning abdominal examination? *Medical education*, 41(9):843–848, 2007.
- [29] Marcello Calisti, Michele Giorelli, Guy Levy, Barbara Mazzolai, B Hochner, Cecilia Laschi, and Paolo Dario. An octopus-bioinspired solution to movement and manipulation for soft robots. *Bioinspiration & biomimetics*, 6(3):036002, 2011.
- [30] Florence Charles, Phrutiya Nilprapruck, David Roux, and Huguette Sallanon. Visible light as a new tool to maintain fresh-cut lettuce post-harvest quality. *Postharvest Biology and Technology*, 135:51–56, 2018.
- [31] Shi Cheng and Zhigang Wu. Microfluidic electronics. *Lab on a Chip*, 12(16):2782–2791, 2012.
- [32] SI Cho, SJ Chang, YY Kim, and KJ An. Ae-automation and emerging technologies: Development of a three-degrees-of-freedom robot for harvesting lettuce using machine vision and fuzzy logic control. *Biosystems Engineering*, 82(2):143–149, 2002.
- [33] C. Chorley, C. Melhuish, T. Pipe, and J. Rossiter. Development of a tactile sensor based on biologically inspired edge encoding. In *2009 Int. Conf. on Advanced Robotics*, pages 1–6, June 2009.
- [34] M Cianchetti, A Arienti, M Follador, B Mazzolai, P Dario, and C Laschi. Design concept and validation of a robotic arm inspired by the octopus. *Materials Science and Engineering: C*, 31(6):1230–1239, 2011.
- [35] Matteo Cianchetti, Federico Renda, Alessia Licofonte, and Cecilia Laschi. Sensorization of continuum soft robots for reconstructing their spatial configuration. In *Biomedical Robotics and Biomechatronics (BioRob), 2012 4th IEEE RAS & EMBS International Conference on*, pages 634–639. IEEE, 2012.
- [36] Matteo Cianchetti, Tommaso Ranzani, Giada Gerboni, Iris De Falco, Cecilia Laschi, and Arianna Menciassi. Stiff-flop surgical manipulator: mechanical design and experimental characterization of the single module. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3576–3581. IEEE, 2013.
- [37] Andy Clark and Rick Grush. Towards a cognitive robotics. *Adaptive Behavior*, 7(1): 5–16, 1999.

- [38] Kevin Cleary and Charles Nguyen. State of the art in surgical robotics: clinical applications and technology challenges. *Computer Aided Surgery*, 6(6):312–328, 2001.
- [39] Steve Collins, Andy Ruina, Russ Tedrake, and Martijn Wisse. Efficient bipedal robots based on passive-dynamic walkers. *Science*, 307(5712):1082–1085, 2005.
- [40] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.
- [41] L. Cramphorn, B. Ward-Cherrier, and N. F. Lepora. Addition of a biomimetic fingerprint on an artificial fingertip enhances tactile spatial acuity. *IEEE Robotics and Automation Letters*, 2(3):1336–1343, July 2017. ISSN 2377-3766. doi: 10.1109/LRA.2017.2665690.
- [42] Utku Çulha and Fumiya Iida. Enhancement of finger motion range with compliant anthropomorphic joint design. *Bioinspiration & biomimetics*, 11(2):026001, 2016.
- [43] Utku Culha, Surya G Nurzaman, Frank Clemens, and Fumiya Iida. Svas3: strain vector aided sensorization of soft structures. *Sensors*, 14(7):12748–12770, 2014.
- [44] Ravinder S Dahiya, Giorgio Metta, Maurizio Valle, and Giulio Sandini. Tactile sensing—from humans to humanoids. *IEEE transactions on robotics*, 26(1):1–20, 2010.
- [45] R.S. Dahiya, G. Metta, M. Valle, and G. Sandini. Tactile sensing-from humans to humanoids. *IEEE Trans. on Robotics*, 26(1):1–20, 2010.
- [46] Lee Danisch, Kevin Englehart, and Andrew Trivett. Spatially continuous six degree of freedom position and orientation sensor. *Sensor Review*, 19(2):106–112, 1999.
- [47] Vytautas Daniulaitis, M Osama Alhalabi, Haruhisa Kawasaki, and Yuji Tanaka. Medical palpation of deformable tissue using physics-based model for haptic interface robot (hiro). In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566)*, volume 4, pages 3907–3911. IEEE, 2004.
- [48] Paolo Dario and Massimo Bergamasco. An advanced robot system for automated diagnostic tasks through palpation. *IEEE Transactions on Biomedical Engineering*, 35(2):118–126, 1988.
- [49] Sheyda Davaria, Farshid Najafi, MJ Mahjoob, and SM Motahari-Bidgoli. Design and fabrication of a robotic tactile device for abdominal palpation. In *2014 Second RSI/ISM International Conference on Robotics and Mechatronics (ICRoM)*, pages 339–344. IEEE, 2014.
- [50] Richard Dawkins et al. *The extended phenotype*, volume 8. Oxford University Press Oxford, 1982.
- [51] Frédérique De Vignemont. Body schema and body image—pros and cons. *Neuropsychologia*, 48(3):669–680, 2010.

- [52] Cosimo Della Santina, Robert K Katzschmann, Antonio Biechi, and Daniela Rus. Dynamic control of soft robots interacting with the environment. In *2018 IEEE International Conference on Soft Robotics (RoboSoft)*, pages 46–53. IEEE, 2018.
- [53] Evan S Dellon, Robin Mourey, and A Lee Dellon. Human pressure perception values for constant and moving one-and two-point discrimination. *Plastic and reconstructive surgery*, 90(1):112–117, 1992.
- [54] Michal Karol Dobrzynski, Ramon Pericet-Camara, and Dario Floreano. Contactless deflection sensor for soft robots. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pages 1913–1918. IEEE, 2011.
- [55] Stephane Doncieux, Nicolas Bredeche, Jean-Baptiste Mouret, and Agoston E Gusz Eiben. Evolutionary robotics: what, why, and where to. *Frontiers in Robotics and AI*, 2:4, 2015.
- [56] Gerald M Edelman. *Neural Darwinism: The theory of neuronal group selection*. Basic books, 1987.
- [57] M Eder, Florian Hisch, and Helmut Hauser. Morphological computation-based control of a modular, pneumatically driven, soft robotic arm. *Advanced Robotics*, 32(7): 375–385, 2018.
- [58] Esraa El Hariri, Nashwa El-Bendary, Aboul Ella Hassanien, and Amr Badr. Automated ripeness assessment system of tomatoes using pca and svm techniques. In *Computer Vision and Image Processing in Intelligent Syst. and Multimedia Technologies*, pages 101–130. IGI global, 2014.
- [59] Yahya Elsayed, Augusto Vincensi, Constantina Lekakou, Tao Geng, CM Saaj, Tommaso Ranzani, Matteo Cianchetti, and Arianna Menciassi. Finite element analysis and design optimization of a pneumatically actuating silicone module for robotic surgery applications. *Soft Robotics*, 1(4):255–262, 2014.
- [60] Shereen Ezzat, Dennis A Sarti, Delver R Cain, and Glenn D Braunstein. Thyroid incidentalomas: prevalence by palpation and ultrasonography. *Archives of internal medicine*, 154(16):1838–1840, 1994.
- [61] Chin-Shyurng Fahn, Cheng-Feng Tsai, and Yung-Wei Lin. Development of a novel two-hand playing piano robot. In *Proceedings of the 4th International Conference of Control, Dynamic Systems, and Robotics (CDSR'17)*. CDSR, 2017.
- [62] Qingchun Feng, Wei Zou, Pengfei Fan, Chunfeng Zhang, and Xiu Wang. Design and test of robotic harvesting system for cherry tomato. *Int. J. of Agricultural and Bio. Eng.*, 11:96–100, 2018.
- [63] Charles M Ferguson. *Inspection, auscultation, palpation, and percussion of the abdomen*. 1990.
- [64] Jeremy A Fishel and Gerald E Loeb. Bayesian exploration for intelligent identification of textures. *Frontiers in neurorobotics*, 6:4, 2012.

- [65] Peter J Fleming and Robin C Purshouse. Evolutionary algorithms in control systems engineering: a survey. *Control engineering practice*, 10(11):1223–1241, 2002.
- [66] Dario Floreano, Ramon Pericet-Camara, Stéphane Viollet, Franck Ruffier, Andreas Brückner, Robert Leitel, Wolfgang Buss, Mohsine Menouni, Fabien Expert, Raphaël Juston, et al. Miniature curved artificial compound eyes. *Proceedings of the National Academy of Sciences*, 110(23):9267–9272, 2013.
- [67] Jerry A. Fodor. Representations: Philosophical essays on the foundations of cognitive science. *British Journal for the Philosophy of Science*, 34(2):175–182, 1983.
- [68] Alan Fogel. Theoretical and applied dynamic systems research in developmental science. *Child Development Perspectives*, 5(4):267–272, 2011.
- [69] Kevin C Galloway, Yue Chen, Emily Templeton, Brian Rife, Isuru S Godage, and Eric J Barth. Fiber optic shape sensing for soft robotics. *Soft robotics*, 2019.
- [70] Carlos J Garci’a, Rocio Garci’a-Villalba, Maria I Gil, and Francisco A Tomas-Barberan. Lc-ms untargeted metabolomics to explain the signal metabolites inducing browning in fresh-cut lettuce. *Journal of agricultural and food chemistry*, 65(22):4526–4535, 2017.
- [71] Lucas A Garibaldi, Barbara Gemmill-Herren, Raffaele D’Annolfo, Benjamin E Graeub, Saul A Cunningham, and Tom D Breeze. Farming approaches for greater biodiversity, livelihoods, and food security. *Trends in ecology & evolution*, 32(1):68–80, 2017.
- [72] P. Gastaldo, L. Pinna, L. Seminara, M. Valle, and R. Zunino. A tensor-based approach to touch modality classification by using machine learning. *Robotics and Autonomous Systems*, 63(Part 3):268 – 278, 2015. ISSN 0921-8890. doi: <https://doi.org/10.1016/j.robot.2014.09.022>. Advances in Tactile Sensing and Touch-based Human Robot Interaction.
- [73] Shineng Geng, Youyu Wang, Cong Wang, and Rongjie Kang. A space tendon-driven continuum robot. In Ying Tan, Yuhui Shi, and Qirong Tang, editors, *Advances in Swarm Intelligence*, pages 25–35, Cham, 2018. Springer International Publishing. ISBN 978-3-319-93818-9.
- [74] R Brent Gillespie, Bo Yu, Robert Grijalva, and Shorya Awtar. Characterizing the feel of the piano action. *Computer Music Journal*, 35(1):43–57, 2011.
- [75] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256, 2010.
- [76] EJ Go, F Del Pozo, JA Quiles, MT Arredondo, H Rahms, M Sanz, P Cano, et al. A telemedicine system for remote cooperative medical imaging diagnosis. *Computer Methods and Programs in Biomedicine*, 49(1):37–48, 1996.
- [77] Eugene Curtis Goldfield. *Emergent forms: Origins and early development of human action and perception*. Oxford University Press on Demand, 1995.

- [78] Rick Grush. In defense of some cartesian'assumptions concerning the brain and its operation. *Biology and Philosophy*, 18(1):53–93, 2003.
- [79] Maria Guijarro, Gonzalo Pajares, Isabel Riomoros, PJ Herrera, XP Burgos-Artizzu, and Angela Ribeiro. Automatic segmentation of relevant textures in agricultural images. *Computers and Electronics in Agriculture*, 75(1):75–83, 2011.
- [80] James C Gwilliam, Zachary Pezzementi, Erica Jantho, Allison M Okamura, and Steven Hsiao. Human vs. robotic tactile sensing: Detecting lumps in soft tissue. In *2010 IEEE Haptics Symposium*, pages 21–28. IEEE, 2010.
- [81] Brian Hall, Monroe W Strickberger, et al. *Strickberger's evolution*. Jones & Bartlett Learning, 2008.
- [82] Fumio Hara and Rolf Pfeifer. *Morpho-functional machines: The new species: Designing embodied intelligence*. Springer Science & Business Media, 2003.
- [83] D Mike Hardin Jr. Acute appendicitis: review and update. *American family physician*, 60(7):2027, 1999.
- [84] Elliot W Hawkes, Laura H Blumenschein, Joseph D Greer, and Allison M Okamura. A soft robot that navigates its environment through growth. *Science Robotics*, 2(8): ean3028, 2017.
- [85] Shigehiko Hayashi, Katsunobu Ganno, Yukitsugu Ishii, and Itsuo Tanaka. Robotic harvesting system for eggplants. *Japan Agricultural Research Quarterly: JARQ*, 36(3):163–168, 2002.
- [86] Liang He, Nicolas Herzig, Simon de Lusignan, Luca Scimeca, Perla Maiolino, Fumiya Iida, and Thrishantha Nanayakkara. An abdominal phantom with tunable stiffness nodules and force sensing capability for palpation training. *IEEE Transactions on Robotics*, forthcoming.
- [87] Nicolas Herzig, Perla Maiolino, Fumiya Iida, and Thrishantha Nanayakkara. A variable stiffness robotic probe for soft tissue palpation. *IEEE Trans. Robot.*, 3(2): 1168–1175, 2018.
- [88] Matej Hoffmann, Hugo Marques, Alejandro Arieta, Hidenobu Sumioka, Max Lungarella, and Rolf Pfeifer. Body schema in robotics: a review. *IEEE Transactions on Autonomous Mental Development*, 2(4):304–324, 2010.
- [89] Matej Hoffmann, Zdeněk Straka, Igor Farkaš, Michal Vavrečka, and Giorgio Metta. Robotic homunculus: Learning of artificial skin representation in a humanoid robot motivated by primary somatosensory cortex. *IEEE Transactions on Cognitive and Developmental Systems*, 10(2):163–176, 2017.
- [90] Bianca S Homberg, Robert K Katzschmann, Mehmet R Dogar, and Daniela Rus. Haptic identification of objects using a modular soft robotic gripper. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pages 1698–1705. IEEE, 2015.

- [91] Bianca S Homberg, Robert K Katzschmann, Mehmet R Dogar, and Daniela Rus. Robust proprioceptive grasping with a soft robot hand. *Autonomous Robots*, pages 1–16, 2018.
- [92] Qilin Hua, Junlu Sun, Haitao Liu, Rongrong Bao, Ruomeng Yu, Junyi Zhai, Caofeng Pan, and Zhong Lin Wang. Skin-inspired highly stretchable and conformable matrix networks for multifunctional sensing. *Nature communications*, 9(1):244, 2018.
- [93] J. Hughes and F. Iida. Localized differential sensing of soft deformable surfaces. In *2017 IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 4959–4964, May 2017. doi: 10.1109/ICRA.2017.7989576.
- [94] JAE Hughes, P Maiolino, and Fumiya Iida. An anthropomorphic soft skeleton hand exploiting conditional models for piano playing. *Science Robotics*, 3(25):eaau3098, 2018.
- [95] Josie Hughes and Fumiya Iida. Localized differential sensing of soft deformable surfaces. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4959–4964. IEEE, 2017.
- [96] Josie Hughes and Fumiya Iida. Tactile sensing applied to the universal gripper using conductive thermoplastic elastomer. *Soft robotics*, 2018.
- [97] Josie Hughes, Utku Culha, Fabio Giardina, Fabian Guenther, Andre Rosendo, and Fumiya Iida. Soft manipulators and grippers: a review. *Frontiers in Robotics and AI*, 3:69, 2016.
- [98] Josie Hughes, Luca Scimeca, Ioana Ifrim, Perla Maiolino, and Fumiya Iida. Achieving robotically peeled lettuce. *IEEE Trans. Robot.*, 2018.
- [99] Josie Hughes, Kieran Gilday, Luca Scimeca, Soham Garg, and Fumiya Iida. Flexible, adaptive industrial assembly: driving innovation through competition. *Intelligent Service Robotics*, pages 1–10, 2019.
- [100] Josie Hughes, Perla Maiolino, Thrishantha Nanayakkara, and Fumiya Iida. Sensorized phantom for characterizing large area deformation of soft bodies for medical applications. In *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*, pages 278–284. IEEE, 2020.
- [101] Josie Hughes, Simon Birell, Luca Scimeca, and Fumiya Iida. *Field Robotics for Harvesting*. Elsevier, forthcoming.
- [102] Jennifer CT Hui and Katherine J Kuchenbecker. Evaluating the biotac’s ability to detect and characterize lumps in simulated tissue. In *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*, pages 295–302. Springer, 2014.
- [103] Jennifer CT Hui, Alexis E Block, Camillo J Taylor, and Katherine J Kuchenbecker. Robust tactile perception of artificial tumors using pairwise comparisons of sensor array readings. In *2016 IEEE Haptics Symposium (HAPTICS)*, pages 305–312. IEEE, 2016.

- [104] Abid Hussain, Hongbin Pu, and Da-Wen Sun. Innovative nondestructive imaging techniques for ripening and maturity of fruits—a review of recent applications. *Trends in Food Sci. & Technology*, 2017.
- [105] Fumiya Iida and Cecilia Laschi. Soft robotics: challenges and perspectives. *Procedia Computer Science*, 7:99–102, 2011.
- [106] Fumiya Iida and Surya G Nurzaman. Adaptation of sensor morphology: an integrative view of perception from biologically inspired robotics perspective. *Interface focus*, 6(4):20160016, 2016.
- [107] Fumiya Iida and Rolf Pfeifer. Sensing through body dynamics. *Robotics and Autonomous Systems*, 54(8):631–640, 2006.
- [108] Nick Jakobi, Phil Husbands, and Inman Harvey. Noise and the reality gap: The use of simulation in evolutionary robotics. In *European Conference on Artificial Life*, pages 704–720. Springer, 1995.
- [109] Lin Jen-Chang, Hsin-Cheng Li, Kuo-Cheng Huang, and Shu-Wei Lin. Design of piano-playing robotic hand. *IAES International Journal of Robotics and Automation*, 3(2):118, 2014.
- [110] SN Jha, K Narsaiah, AD Sharma, Manpreet Singh, Sunil Bansal, and R Kumar. Quality parameters of mango and potential of non-destructive techniques for their measurement—a review. *J. of Food Sci. and Technology*, 47(1):1–14, 2010.
- [111] Allen Jiang, Georgios Xynogalas, Prokar Dasgupta, Kaspar Althoefer, and Thrishantha Nanayakkara. Design of a variable stiffness flexible manipulator with composite granular jamming and membrane coupling. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2922–2927. IEEE, 2012.
- [112] Mark H Johnson. The neural basis of cognitive development. *Handbook of child psychology*, 2:1–49, 1998.
- [113] Mohsen Kaboli, Alex Long, and Gordon Cheng. Humanoids learn touch modalities identification via multi-modal robotic skin and robust tactile descriptors. *Advanced Robotics*, 29(21):1411–1425, 2015. doi: 10.1080/01691864.2015.1095652.
- [114] Adel A Kader. Fruit maturity, ripening, and quality relationships. In *Int. Symp. Effect of Pre- & Postharvest factors in Fruit Storage 485*, pages 203–208, 1997.
- [115] Adel A Kader and Rosa Sonya Rolle. *The role of post-harvest management in assuring the quality and safety of horticultural produce*, volume 152. Food & Agriculture Org., 2004.
- [116] Keren Kapach, Ehud Barnea, Rotem Mairon, Yael Edan, and Ohad Ben-Shahar. Computer vision for fruit harvesting robots—state of the art and challenges ahead. *International Journal of Computational Vision and Robotics*, 3(1-2):4–34, 2012.
- [117] Ichiro Kato, Sadamu Ohteru, Katsuhiko Shirai, Toshiaki Matsushima, Seinosuke Narita, Shigeki Sugano, Tetsunori Kobayashi, and Eizo Fujisawa. The robot musician 'wabot-2' (waseda robot-2). *Robotics*, 3(2):143–155, 1987.

- [118] Ichiro Kato, Koichi Koganezawa, and Atsuo Takanishi. Automatic breast cancer palpation robot: Wapro-4. *Advanced Robotics*, 3(4):251–261, 1988.
- [119] Robert K Katzschmann, Joseph DelPreto, Robert MacCurdy, and Daniela Rus. Exploration of underwater life with an acoustically controlled soft robotic fish. *Science Robotics*, 3(16):ear3449, 2018.
- [120] Sangbae Kim, Cecilia Laschi, and Barry Trimmer. Soft robotics: a bioinspired evolution in robotics. *Trends in biotechnology*, 31(5):287–294, 2013.
- [121] Jelizaveta Konstantinova, Allen Jiang, Kaspar Althoefer, Prokar Dasgupta, and Thrishantha Nanayakkara. Implementation of tactile sensing for palpation in robot-assisted minimally invasive surgery: A review. *IEEE Sensors Journal*, 14(8):2490–2501, 2014.
- [122] Jelizaveta Konstantinova, Min Li, Gautam Mehra, Prokar Dasgupta, Kaspar Althoefer, and Thrishantha Nanayakkara. Behavioral characteristics of manual palpation to localize hard nodules in soft tissues. *IEEE Transactions on Biomedical Engineering*, 61(6):1651–1659, 2014.
- [123] Jelizaveta Konstantinova, Min Li, Gautam Mehra, Prokar Dasgupta, Kaspar Althoefer, and Thrishantha Nanayakkara. Behavioral characteristics of manual palpation to localize hard nodules in soft tissues. *IEEE Transactions on Biomedical Engineering*, 61(6):1651–1659, 2014.
- [124] Jelizaveta Konstantinova, Giuseppe Cotugno, Prokar Dasgupta, Kaspar Althoefer, and Thrishantha Nanayakkara. Palpation force modulation strategies to identify hard regions in soft tissue organs. *PloS one*, 12(2):e0171706, 2017.
- [125] Konrad P Körding and Daniel M Wolpert. Bayesian decision theory in sensorimotor control. *Trends in cognitive sciences*, 10(7):319–326, 2006.
- [126] Patricia K Kuhl. Language, mind, and brain: Experience alters perception. *The new cognitive neurosciences*, 2:99–115, 2000.
- [127] Shinya Kusuda, Satoshi Sawano, and Satoshi Konishi. Fluid-resistive bending sensor having perfect compatibility with flexible pneumatic balloon actuator. In *Micro Electro Mechanical Systems, 2007. MEMS. IEEE 20th International Conference on*, pages 615–618. IEEE, 2007.
- [128] Keerthy Kusumam, Tomáš Krajník, Simon Pearson, Grzegorz Cielniak, and Tom Duckett. Can you pick a broccoli? 3d-vision based detection and localisation of broccoli heads in the field. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 646–651. IEEE, 2016.
- [129] Michael F Land and Dan-Eric Nilsson. *Animal eyes*. Oxford University Press, 2012.
- [130] Cecilia Laschi and Matteo Cianchetti. Soft robotics: new perspectives for robot bodyware and control. *Frontiers in bioengineering and biotechnology*, 2:3, 2014.
- [131] Cecilia Laschi, Matteo Cianchetti, Barbara Mazzolai, Laura Margheri, Maurizio Follador, and Paolo Dario. Soft robot arm inspired by the octopus. *Advanced Robotics*, 26(7):709–727, 2012.

- [132] Frank A Lederle and David L Simel. Does this patient have abdominal aortic aneurysm? *Jama*, 281(1):77–82, 1999.
- [133] Sukhan Lee. *Recent Progress in Robotics: Viable Robotic Service to Human: An Edition of the Selected Papers from the 13th International Conference on Advanced Robotics*, volume 370. Springer Science & Business Media, 2008.
- [134] Bin Li, Ye Shi, Adam Fontecchio, and Yon Visell. Mechanical imaging of soft tissues with a highly compliant tactile sensing array. *IEEE Transactions on Biomedical Engineering*, 65(3):687–697, 2017.
- [135] SH Li, QY Zeng, YL Xiao, SY Fu, and BL Zhou. Biomimicry of bamboo bast fiber with engineering composite materials. *Materials Science and Engineering: C*, 3(2): 125–130, 1995.
- [136] Yen-Fang Li and Chi-Yi Lai. Intelligent algorithm for music playing robot—applied to the anthropomorphic piano robot control. In *2014 IEEE 23rd International Symposium on Industrial Electronics (ISIE)*, pages 1538–1543. IEEE, 2014.
- [137] Lukas Lichtensteiger. The Need to Adapt and Its Implications for Embodiment. In Fumiya Iida, Rolf Pfeifer, Luc Steels, and Yasuo Kuniyoshi, editors, *Embodied Artificial Intelligence: International Seminar, Dagstuhl Castle, Germany, July 7-11, 2003. Revised Papers*, pages 98–106. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004. ISBN 978-3-540-27833-7.
- [138] Angelica Lim, Takeshi Mizumoto, Toru Takahashi, Tetsuya Ogata, and Hiroshi G Okuno. Programming by playing and approaches for expressive robot performances. In *IROS Workshop on Robots and Musical Expressions*, 2010.
- [139] Angelica Lim, Tetsuya Ogata, and Hiroshi G Okuno. Towards expressive musical robots: a cross-modal framework for emotional gesture, voice and music. *EURASIP Journal on Audio, Speech, and Music Processing*, 2012(1):3, 2012.
- [140] Huai-Ti Lin, Gary G Leisk, and Barry Trimmer. Goqbot: a caterpillar-inspired soft-bodied rolling robot. *Bioinspiration & biomimetics*, 6(2):026007, 2011.
- [141] Jen-Chang Lin, Hsing-Hsin Huang, Yen-Fang Li, Jen-Chao Tai, and Li-Wei Liu. Electronic piano playing robot. In *2010 International Symposium on Computer, Communication, Control and Automation (3CA)*, volume 2, pages 353–356. IEEE, 2010.
- [142] Hod Lipson and Jordan B Pollack. Automatic design and manufacture of robotic lifeforms. *Nature*, 406(6799):974, 2000.
- [143] Jizhan Liu, Pingping Li, and Hanping Mao. Mechanical and kinematic modeling of assistant vacuum sucking and pulling operation oftomato fruits in robotic harvesting. *Transactions of the ASABE*, 58(3):539–550, 2015.
- [144] S.P. Lloyd. Least squares quantization in pcm. *IEEE Trans. Inform. Theory*, 28(2): 129–137, 1982. doi: 10.1109/TIT.1982.1056489.

- [145] Gerald E Loeb and Jeremy A Fishel. Bayesian action&perception: Representing the world in the brain. *Frontiers in neuroscience*, 8:341, 2014.
- [146] Nanshu Lu and Dae-Hyeong Kim. Flexible and stretchable electronics paving the way for soft robotics. *Soft Robotics*, 1(1):53–62, 2014.
- [147] Chiara Lucarotti, Calogero Maria Oddo, Nicola Vitiello, and Maria Chiara Carrozza. Synthetic and bio-artificial tactile sensing: A review. *Sensors*, 13(2):1435–1466, 2013.
- [148] Henrik Hautop Lund, John Hallam, and Wei-Po Lee. Evolving robot morphology. In *Proceedings of 1997 IEEE International Conference on Evolutionary Computation (ICEC'97)*, pages 197–202. IEEE, 1997.
- [149] M. Lungarella and Olaf Sporns. Information self-structuring: Key principle for learning and development. volume 2005, pages 25–30, 08 2005. ISBN 0-7803-9226-4. doi: 10.1109/DEVLRN.2005.1490938.
- [150] Max Lungarella, Giorgio Metta, Rolf Pfeifer, and Giulio Sandini. Developmental robotics: a survey. *Connection science*, 15(4):151–190, 2003.
- [151] Max Lungarella, Teresa Pegors, Daniel Bulwinkle, and Olaf Sporns. Methods for quantifying the informational structure of sensory and motor data. *Neuroinformatics*, 3(3):243–262, 2005.
- [152] Marco Maggiali, G Cannata, Perla Maiolino, Giorgio Metta, Marco Randazzo, and Giulio Sandini. Embedded distributed capacitive tactile sensor. 07 2008.
- [153] P Maiolino, Fabia Galantini, F Mastrogiovanni, G Gallone, G Cannata, and Federico Carpi. Soft dielectrics for capacitive sensing in robot skins: Performance of different elastomer types. *Sensors and Actuators A: Physical*, 226:37–47, 2015.
- [154] Perla Maiolino, Marco Maggiali, Giorgio Cannata, Giorgio Metta, and Lorenzo Natale. A flexible and robust large scale capacitive tactile system for robots. *IEEE Sensors Journal*, 13(10):3910–3917, 2013.
- [155] Carmel Majidi. Soft robotics: a perspective—current trends and prospects for the future. *Soft Robotics*, 1(1):5–11, 2014.
- [156] Maja J Matarić. Situated robotics. *Encyclopedia of cognitive science*, 2006.
- [157] Alexander G. de G. Matthews, Mark van der Wilk, Tom Nickson, Keisuke. Fujii, Alexis Boukouvalas, Pablo León-Villagrà, Zoubin Ghahramani, and James Hensman. GPflow: A Gaussian process library using TensorFlow. *Journal of Machine Learning Research*, 18(40):1–6, apr 2017. URL <http://jmlr.org/papers/v18/16-537.html>.
- [158] Craig Mautner and Richard K Belew. Evolving robot morphology and control. *Artificial Life and Robotics*, 4(3):130–136, 2000.
- [159] Barbara Mazzolai, Alessio Mondini, Paolo Corradi, Cecilia Laschi, Virgilio Mattoli, Edoardo Sinibaldi, and Paolo Dario. A miniaturized mechatronic system inspired by plant roots for soil exploration. *IEEE/ASME Transactions on Mechatronics*, 16(2): 201–212, 2010.

- [160] Barbara Mazzolai, Lucia Beccai, and Virgilio Mattoli. Plants as model in biomimetics and biorobotics: new perspectives. *Frontiers in bioengineering and biotechnology*, 2: 2, 2014.
- [161] Sharon McDonald, Debbie Saslow, and Marianne H Alciati. Performance and reporting of clinical breast examination: a review of the literature. *CA: a cancer journal for clinicians*, 54(6):345–361, 2004.
- [162] RC Miall, D Jo Weir, Daniel M Wolpert, and JF Stein. Is the cerebellum a smith predictor? *Journal of motor behavior*, 25(3):203–216, 1993.
- [163] Aslan Miriyev, Kenneth Stack, and Hod Lipson. Soft material for soft actuators. *Nature communications*, 8(1):596, 2017.
- [164] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [165] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [166] Nur Badariah Ahmad Mustafa, Nurashikin Ahmad Fuad, Syed Khaleel Ahmed, Aidil Azwin Zainul Abidin, Zaipatimah Ali, Wong Bing Yit, and Zainul Abidin Md Sharrif. Image processing of an agriculture produce: Determination of size and ripeness of a banana. In *Inf. Technology, 2008. ITSIM 2008. Int. Symp. on*, volume 1, pages 1–7. IEEE, 2008.
- [167] Joseph T Muth, Daniel M Vogt, Ryan L Truby, Yiğit Mengüç, David B Kolesky, Robert J Wood, and Jennifer A Lewis. Embedded 3d printing of strain sensors within highly stretchable elastomers. *Advanced Materials*, 26(36):6307–6312, 2014.
- [168] Cota Nabeshima, Yasuo Kuniyoshi, and Max Lungarella. Adaptive body schema for robotic tool-use. *Advanced Robotics*, 20(10):1105–1126, 2006.
- [169] Kohei Nakajima, Helmut Hauser, Rongjie Kang, Emanuele Guglielmino, Darwin G Caldwell, and Rolf Pfeifer. A soft body as a reservoir: case studies in a dynamic model of octopus-inspired soft robotic arm. *Frontiers in computational neuroscience*, 7:91, 2013.
- [170] Kohei Nakajima, Tao Li, Helmut Hauser, and Rolf Pfeifer. Exploiting short-term memory in soft body dynamics as a computational resource. *Journal of The Royal Society Interface*, 11(100):20140437, 2014.
- [171] Kohei Nakajima, Helmut Hauser, Tao Li, and Rolf Pfeifer. Information processing via physical soft body. 5:10487, 05 2015.
- [172] Kohei Nakajima, Helmut Hauser, Tao Li, and Rolf Pfeifer. Information processing via physical soft body. *Scientific reports*, 5:10487, 2015.
- [173] V Eyarkai Nambi, K Thangavel, and D Manohar Jesudas. Scientific classification of ripening period and development of colour grade chart for indian mangoes (*mangifera indica* l.) using multivariate cluster analysis. *Scientia Horticulturae*, 193:90–98, 2015.

- [174] Shruti Nambiar and John TW Yeow. Conductive polymer-based sensors for biomedical applications. *Biosensors and Bioelectronics*, 26(5):1825–1832, 2011.
- [175] Allen Newell and Herbert A Simon. Computer science as empirical inquiry: Symbols and search. *PHILOSOPHY OF PSYCHOLOGY*, page 407, 1975.
- [176] Tien Thanh Nguyen, Koenraad Vandevoorde, Erdal Kayacan, Josse De Baerdemaeker, and Wouter Saeys. Apple detection algorithm for robotic harvesting using a rgb-d camera. In *Int. Conf. of Agricultural Eng., Zurich, Switzerland*, 2014.
- [177] Kirk A Nichols and Allison M Okamura. Methods to segment hard inclusions in soft tissue during autonomous robotic palpation. *IEEE Transactions on Robotics*, 31(2): 344–354, 2015.
- [178] Stefano Nolfi, Dario Floreano, Orazio Miglino, and Francesco Mondada. How to evolve autonomous robots: Different approaches in evolutionary robotics. In *Artificial life iv: Proceedings of the fourth international workshop on the synthesis and simulation of living systems*, pages 190–197. MIT Press, 1994.
- [179] Stefano Nolfi, Dario Floreano, and Director Dario Floreano. *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT press, 2000.
- [180] Surya G Nurzaman, Utku Culha, Luzius Brodbeck, Liyu Wang, and Fumiya Iida. Active sensing system with in situ adjustable sensor morphology. *PLoS One*, 8(12): e84090, 2013.
- [181] Ailish O’Halloran, Fergal O’malley, and Peter McHugh. A review on dielectric elastomer actuators, technology, applications, and challenges. *Journal of Applied Physics*, 104(7):9, 2008.
- [182] Lars Olsson, Chrystopher L Nehaniv, and Daniel Polani. Sensory channel grouping and structure from uninterpreted sensor data. In *Evolvable Hardware, 2004. Proceedings. 2004 NASA/DoD Conference on*, pages 153–160. IEEE, 2004.
- [183] J Kevin O’Regan and Alva Noë. A sensorimotor account of vision and visual consciousness. *Behavioral and brain sciences*, 24(5):939–973, 2001.
- [184] Pierre-Yves Oudeyer, Frdric Kaplan, and Verena V Hafner. Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2):265–286, 2007.
- [185] Claudio Pacchierotti, Domenico Prattichizzo, and Katherine J Kuchenbecker. Cutaneous feedback of fingertip deformation and vibration for palpation in robotic surgery. *IEEE Transactions on Biomedical Engineering*, 63(2):278–287, 2015.
- [186] Sue Taylor Parker and Michael L McKinney. *Origins of intelligence: The evolution of cognitive development in monkeys, apes, and humans*. JHU Press, 2012.
- [187] Martin J Pearson, Ben Mitchinson, J Charles Sullivan, Anthony G Pipe, and Tony J Prescott. Biomimetic vibrissal sensing for robots. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1581):3085–3096, 2011.

- [188] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [189] Rolf Pfeifer and Christian Scheier. Sensory—motor coordination: The metaphor and beyond. *Robotics and autonomous systems*, 20(2-4):157–178, 1997.
- [190] Rolf Pfeifer and Christian Scheier. *Understanding intelligence*. MIT press, 1999.
- [191] Rolf Pfeifer, Fumiya Iida, and Josh Bongard. New robotics: Design principles for intelligent systems. *Artificial life*, 11(1-2):99–120, 2005.
- [192] Rolf Pfeifer, Fumiya Iida, and Gabriel Gómez. Morphological computation for adaptive behavior and cognition. In *International Congress Series*, volume 1291, pages 22–29. Elsevier, 2006.
- [193] Rolf Pfeifer, Max Lungarella, and Fumiya Iida. Self-organization, embodiment, and biologically inspired robotics. *science*, 318(5853):1088–1093, 2007.
- [194] Rolf Pfeifer, Max Lungarella, Olaf Sporns, and Yasuo Kuniyoshi. On the information theoretic implications of embodiment—principles and methods. In *50 years of artificial intelligence*, pages 76–86. Springer, 2007.
- [195] Rolf Pfeifer, Max Lungarella, and Fumiya Iida. The challenges ahead for bio-inspired soft robotics. *Communications of the ACM*, 55(11):76–87, 2012.
- [196] Rolf Pfeifer, Fumiya Iida, and Max Lungarella. Cognition from the bottom up: on biological inspiration, body morphology, and soft materials. *Trends in cognitive sciences*, 18(8):404–413, 2014.
- [197] Rolf Pfeifer et al. On the role of morphology and materials in adaptive behavior. *From animals to animats*, 6:23–32, 2000.
- [198] Jean Piaget. *The psychology of intelligence*. Routledge, 2003.
- [199] Jean Piaget and Margaret Cook. *The origins of intelligence in children*, volume 8. International Universities Press New York, 1952.
- [200] Tony J Prescott, Martin J Pearson, Ben Mitchinson, J Charles W Sullivan, and Anthony G Pipe. Whisking with robots. *IEEE robotics & automation magazine*, 16(3):42–50, 2009.
- [201] P Puangmali, K Althoefer, LD Seneviratne, D Murphy, and P Dasgupta. State-of-the-art in force and tactile sensing for minimally invasive surgery. *IEEE Sensors Journal*, 4(8):371–381, 2008.
- [202] Simon S Rabinowitz and Eric Dean. Abdominal examination: Overview, preparation, technique, Nov 2017. URL <https://emedicine.medscape.com/article/1909183-overview>.
- [203] S Ray. On a theoretical property of the bhattacharyya coefficient as a feature evaluation criterion. *Pattern recognition letters*, 9(5):315–319, 1989.

- [204] Steven I Rich, Robert J Wood, and Carmel Majidi. Untethered soft robotics. *Nature Electronics*, 1(2):102, 2018.
- [205] John Rieffel, Davis Knox, Schuyler Smith, and Barry Trimmer. Growing and evolving soft robots. *Artificial life*, 20(1):143–162, 2014.
- [206] M SB Shah Rizam, AR Farah Yasmin, MY Ahmad Ihsan, and K Shazana. Non-destructive watermelon ripeness determination using image processing and artificial neural network (ann). *World Academy of Sci., Eng. and Technology*, 38:542–546, 2009.
- [207] Philippe Rochat. Self-perception and action in infancy. *Experimental brain research*, 123(1-2):102–109, 1998.
- [208] Hugo Rodrigue, Wei Wang, Min-Woo Han, Thomas JY Kim, and Sung-Hoon Ahn. An overview of shape memory alloy-coupled actuators and robots. *Soft robotics*, 4(1): 3–15, 2017.
- [209] Hugo Rodrigue, Wei Wang, Dong-Ryul Kim, and Sung-Hoon Ahn. Curved shape memory alloy-based soft actuators and application to soft gripper. *Composite Structures*, 176:398–406, 2017.
- [210] John A Rogers, Takao Someya, and Yonggang Huang. Materials and mechanics for stretchable electronics. *science*, 327(5973):1603–1607, 2010.
- [211] Jonathan Rossiter and Helmut Hauser. Soft robotics—the next industrial revolution. *IEEE Robot. Autom. Mag*, 23:17–20, 2016.
- [212] Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.
- [213] D Caleb Rucker and Robert J Webster. Mechanics of continuum robots with external loading and general tendon routing. In *Experimental Robotics*, pages 645–654. Springer, 2014.
- [214] Daniela Rus and Michael T Tolley. Design, fabrication and control of soft robots. *Nature*, 521(7553):467, 2015.
- [215] Ali Sadeghi, Alice Tonazzini, Liyana Popova, and Barbara Mazzolai. Robotic mechanism for soil penetration inspired by plant root. In *2013 IEEE International Conference on Robotics and Automation*, pages 3457–3462. IEEE, 2013.
- [216] Ali Sadeghi, Alice Tonazzini, Liyana Popova, and Barbara Mazzolai. A novel growing device inspired by plant root soil penetration behaviors. *PloS one*, 9(2):e90139, 2014.
- [217] Carl E Sams. Preharvest factors affecting postharvest texture. *Postharvest Bio. and Technology*, 15(3):249–254, 1999.
- [218] A Sarvazyan, V Egorov, and N Sarvazyan. Tactile sensing and tactile imaging in detection of cancer. In *Biosensors and molecular technologies for cancer diagnostics*, pages 337–352. Taylor & Francis, Abingdon, UK, 2012.

- [219] Christian Scheier and Rolf Pfeifer. The embodied cognitive science approach. In *Dynamics, Synergetics, Autonomous Agents: Nonlinear Systems Approaches to Cognitive Psychology and Cognitive Science*, pages 159–179. World Scientific, 1999.
- [220] G Scheltjens, MM Diaz, J Brancart, G Van Assche, and B Van Mele. A self-healing polymer network based on reversible covalent bonding. *Reactive and Functional Polymers*, 73(2):413–420, 2013.
- [221] R Schiessel, M Wunderlich, and F Herbst. Local recurrence of colorectal cancer: effect of early detection and aggressive surgery. *British journal of surgery*, 73(5): 342–344, 1986.
- [222] Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- [223] A. Schmitz, P. Maiolino, M. Maggiali, L. Natale, G. Cannata, and G. Metta. Methods and Technologies for the Implementation of Large-Scale Robot Tactile Sensors. *IEEE Trans. on Robotics*, 27(3):389–400, June 2011. ISSN 1552-3098. doi: 10.1109/TRO.2011.2132930.
- [224] Luca Scimeca and Fumiya Iida. *Soft Robotics: A Developmental Approach*. MIT Press, forthcoming.
- [225] Luca Scimeca, Josie Hughes, Perla Maiolino, Liang He, Thrishantha Nanayakkara, and Fumiya Iida. A bayesian framework for multi-axis soft-body palpation. *Soft Robotics*, .
- [226] Luca Scimeca, Cheryn Ng, and Fumiya Iida. Gaussian process inference modelling of dynamic robot control for expressive piano playing. *PloS one*, .
- [227] Luca Scimeca, Perla Maiolino, and Fumiya Iida. Soft morphological processing of tactile stimuli for autonomous category formation. In *2018 IEEE International Conference on Soft Robotics (RoboSoft)*. IEEE, 2018.
- [228] Luca Scimeca, Perla Maiolino, and Fumiya Iida. Soft morphological processing of tactile stimuli for autonomous category formation. IEEE International Conference on Soft Robotics, Robosoft, 2018.
- [229] Luca Scimeca, Josie Hughes, Perla Maiolino, and Fumiya Iida. Model-free soft-structure reconstruction for proprioception using tactile arrays. *IEEE Robotics and Automation Letters*, 2019.
- [230] Luca Scimeca, Perla Maiolino, Daniel Cardin-Catalan, Angel P del Pobil, Antonio Morales, and Fumiya Iida. Non-destructive robotic assessment of mango ripeness via multi-point soft haptics. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 1821–1826. IEEE, 2019.
- [231] Luca Scimeca, Perla Maiolino, Ed Bray, and Fumiya Iida. Structuring of tactile sensory information for category formation in robotics palpation. *Autonomous Robots*, pages 1–17, 2020.

- [232] Luca Scimeca, Perla Maiolino, and Fumiya Iida. Efficient bayesian exploration for soft morphology-action co-optimization. In *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*, pages 639–644. IEEE, 2020.
- [233] Luca Scimeca, Perla Maiolino, and Fumiya Iida. Self-supervised learning through scene observation for selective item identification in conveyor belt applications. In *Annual Conference Towards Autonomous Robotic Systems*. Springer, forthcoming.
- [234] Thomas C Scott-Phillips, Kevin N Laland, David M Shuker, Thomas E Dickins, and Stuart A West. The niche construction perspective: a critical appraisal. *Evolution*, 68(5):1231–1243, 2014.
- [235] Ruby T Senie, Paul Peter Rosen, Martin L Lesser, and David W Kinne. Breast self-examination and medical examination related to breast cancer stage. *American Journal of Public Health*, 71(6):583–590, 1981.
- [236] Sangok Seok, Cagdas D Onal, Robert Wood, Daniela Rus, and Sangbae Kim. Peristaltic locomotion with antagonistic actuators in soft robotics. In *2010 IEEE International Conference on Robotics and Automation*, pages 1228–1233. IEEE, 2010.
- [237] Chaoyang Shi, Xiongbiao Luo, Peng Qi, Tianliang Li, Shuang Song, Zoran Najdovski, Toshio Fukuda, and Hongliang Ren. Shape sensing techniques for continuum robots in minimally invasive surgery: A survey. *IEEE Transactions on Biomedical Engineering*, 64(8):1665–1678, 2017.
- [238] M. Shimojo. Mechanical filtering effect of elastic cover for tactile sensor. *IEEE Trans. Robot. Automat.*, 13(1):128–132, February 1997. ISSN 1042-296X. doi: 10.1109/70.554353.
- [239] Jun Shintake, Bryan Schubert, Samuel Rosset, Herbert Shea, and Dario Floreano. Variable stiffness actuator for soft robotics using dielectric elastomer and low-melting-point alloy. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1097–1102. IEEE, 2015.
- [240] Robert W Shumaker, Kristina R Walkup, and Benjamin B Beck. *Animal tool behavior: the use and manufacture of tools by animals*. JHU Press, 2011.
- [241] KO Siegenthaler, A Künkel, G Skupin, and M Yamamoto. Ecoflex® and ecovio®: biodegradable, performance-enabling plastics. In *Synthetic biodegradable polymers*, pages 91–136. Springer, 2011.
- [242] Eric Singer, Jeff Feddersen, Chad Redmon, and Bil Bowen. Lemur’s musical robots. In *Proceedings of the 2004 conference on New interfaces for musical expression*, pages 181–184, 2004.
- [243] DC Slaughter. Nondestructive maturity assessment methods for mango. *University of California, Davis*, pages 1–18, 2009.
- [244] Nantachai Sornkarn and Thrishantha Nanayakkara. The efficacy of interaction behavior and internal stiffness control for embodied information gain in haptic perception. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on*, pages 2657–2662. IEEE, 2016.

- [245] Nantachai Sornkarn, Prokar Dasgupta, and Thrishantha Nanayakkara. Morphological computation of haptic perception of a controllable stiffness probe. *PLoS one*, 11(6): e0156982, 2016.
- [246] Olaf Sporns. *Embodied cognition*. MIT Press, 2003.
- [247] Olaf Sporns and Max Lungarella. Evolving coordinated behavior by maximizing information structure. In *Artificial life X: proceedings of the tenth international conference on the simulation and synthesis of living systems*, volume 10, page 323. Citeseer, 2006.
- [248] Li Sun, Gerardo Aragon-Camarasa, Simon Rogers, Rustam Stolkin, and J Paul Siebert. Single-shot clothing category recognition in free-configurations with application to autonomous clothes sorting. *arXiv preprint arXiv:1707.07157*, 2017.
- [249] S Taccola, A Zucca, F Greco, B Mazzolai, and V Mattoli. Electrically driven dry state actuators based on pedot: Pss nanofilms. In *EuroEAP 2013, International Conference on Electromechanically Active Polymer (EAP) Transducers & Artificial Muscles*, 2013.
- [250] G. Taga, R. Takaya, and Y. Konishi. Analysis of general movements of infants towards understanding of developmental principle for motor control. In *IEEE SMC'99 Conference Proceedings. 1999 IEEE International Conference on Systems, Man, and Cybernetics (Cat. No.99CH37028)*, volume 5, pages 678–683 vol.5, Oct 1999. doi: 10.1109/ICSMC.1999.815633.
- [251] Ali Talasaz and Rajnikant V Patel. Telerobotic palpation for tumor localization with depth estimation. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 463–468. IEEE, 2013.
- [252] David Silvera Tawil, David Rye, and Mari Velonaki. Interpretation of the modality of touch on an artificial arm covered with an eit-based sensitive skin. *The International Journal of Robotics Research*, 31(13):1627–1641, 2012. doi: 10.1177/0278364912455441.
- [253] Russell H Taylor, Arianna Menciassi, Gabor Fichtinger, Paolo Fiorini, and Paolo Dario. Medical robotics and computer-integrated surgery. In *Springer handbook of robotics*, pages 1657–1684. Springer, 2016.
- [254] Seppe Terryn, Joost Brancart, Dirk Lefeber, Guy Van Assche, and Bram Vanderborght. Self-healing soft pneumatic robots. *science Robotics*, 2(9):eaan4268, 2017.
- [255] Esther Thelen and Linda B Smith. *A dynamic systems approach to the development of cognition and action*. MIT press, 1996.
- [256] Hamish Thomson and DMA Francis. Abdominal-wall tenderness: a useful sign in the acute abdomen. *The Lancet*, 310(8047):1053–1054, 1977.
- [257] Michael E. Tipping and Christopher M. Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, January 1999. ISSN 1467-9868. doi: 10.1111/1467-9868.00196.

- [258] Michalis Titsias. Variational learning of inducing variables in sparse gaussian processes. In *Artificial Intelligence and Statistics*, pages 567–574, 2009.
- [259] Michael T Tolley, Robert F Shepherd, Michael Karpelson, Nicholas W Bartlett, Kevin C Galloway, Michael Wehner, Rui Nunes, George M Whitesides, and Robert J Wood. An untethered jumping soft robot. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 561–566. IEEE, 2014.
- [260] Anthony James Topper and Thomas Maloney. Piano-playing robotic arm, Nov 2019. URL <https://digitalcommons.wpi.edu/mqp-all/7087/>.
- [261] R Blythe Towal, Brian W Quist, Venkatesh Gopal, Joseph H Solomon, and Mitra JZ Hartmann. The morphology of the rat vibrissal array: a model for quantifying spatiotemporal patterns of whisker-object contact. *PLoS computational biology*, 7(4): e1001120, 2011.
- [262] Ana Luisa Trejos, Jagadeesan Jayender, MT Perri, Michael D Naish, Rajnikant V Patel, and RA Malthaner. Robot-assisted tactile sensing for minimally invasive tumor localization. *The International Journal of Robotics Research*, 28(9):1118–1133, 2009.
- [263] Edwin Royden Trueman. *Locomotion of soft-bodied animals*. Edward Arnold, 1975.
- [264] Manos Tsakiris. My body in the brain: a neurocognitive model of body-ownership. *Neuropsychologia*, 48(3):703–712, 2010.
- [265] Paul Tuffield and Hugo Elias. The shadow robot mimics human actions. *Industrial Robot: An International Journal*, 30(1):56–60, 2003.
- [266] Michael T Turvey. Coordination. *American psychologist*, 45(8):938, 1990.
- [267] SA Velcro. Improvements in or relating to a method and a device for producing a velvet type fabric. *Swiss patent*, 721338, 1955.
- [268] D Vogt, Yigit Menguc, Yong-Lae Park, Michael Wehner, RK Kramer, Carmel Majidi, LP Jentoft, Yaroslav Tenzer, RD Howe, and Robert J Wood. Progress in soft, flexible, and stretchable sensing systems. In *Proceedings of the International Workshop on Research Frontiers in Electronics Skin Technology at ICRA*, volume 13, 2013.
- [269] Ian D Walker, Darren M Dawson, Tamar Flash, Frank W Grasso, Roger T Hanlon, Binyamin Hochner, William M Kier, Christopher C Pagano, Christopher D Rahn, and Qiming M Zhang. Continuum robot arms inspired by cephalopods. In *Unmanned Ground Vehicle Technology VII*, volume 5804, pages 303–315. International Society for Optics and Photonics, 2005.
- [270] Robert J Webster, Jasenka Memisevic, and Allison M Okamura. Design considerations for robotic needle steering. In *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, pages 3588–3594. IEEE, 2005.
- [271] Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.

- [272] Robert A Wilson and Lucia Foglia. Embodied cognition. *The Stanford Encyclopedia of Philosophy*, 2011.
- [273] Daniel M Wolpert, Kenji Doya, and Mitsuo Kawato. A unifying computational framework for motor control and social interaction. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 358(1431):593–602, 2003.
- [274] Danfei Xu, Gerald E Loeb, and Jeremy A Fishel. Tactile identification of objects using bayesian exploration. In *2013 IEEE International Conference on Robotics and Automation*, pages 3056–3061. IEEE, 2013.
- [275] Sajjad Yaghoubi, Negar Ali Akbarzadeh, Shadi Sadeghi Bazargani, Sama Sadeghi Bazargani, Marjan Bamizan, and Maryam Irani Asl. Autonomous robots for agricultural tasks and farm assignment and future trends in agro robots. *International Journal of Mechanical and Mechatronics Engineering*, 13(3):1–6, 2013.
- [276] Hong Kai Yap, Hui Yong Ng, and Chen-Hua Yeow. High-force soft printable pneumatics for soft robotic applications. *Soft Robotics*, 3(3):144–158, 2016.
- [277] Ping-Lang Yen. Palpation sensitivity analysis of exploring hard objects under soft tissue. In *Advanced Intelligent Mechatronics, 2003. AIM 2003. Proceedings. 2003 IEEE/ASME International Conference on*, volume 2, pages 1102–1106. IEEE, 2003.
- [278] Fatih Yildiz. Initial preparation, handling, and distribution of minimally processed refrigerated fruits and vegetables. In *Minimally processed refrigerated fruits and vegetables*, pages 53–92. Springer, 2017.
- [279] Osman Dogan Yirmibesoglu, John Morrow, Steph Walker, Walker Gosrich, Reece Cañizares, Hansung Kim, Uranbileg Daalkhajav, Chloe Fleming, Callie Branyan, and Yigit Menguc. Direct 3d printing of silicone elastomer soft robots and their performance comparison with molded counterparts. In *2018 IEEE International Conference on Soft Robotics (RoboSoft)*, pages 295–302. IEEE, 2018.
- [280] Karl E Zelik and Arthur D Kuo. Human walking isn’t all hard work: evidence of soft tissue contributions to energy dissipation and return. *Journal of Experimental Biology*, 213(24):4257–4264, 2010.
- [281] Dan Zhang, Jianhe Lei, Beizhi Li, Daniel Lau, and Craig Cameron. Design and analysis of a piano playing robot. In *2009 International Conference on Information and Automation*, pages 757–761. IEEE, 2009.
- [282] Jordan Zlatev and Christian Balkenius. Introduction: Why ‘epigenetic robotics’ ? *Lund University Cognitive Studies*, pages 1–4, 2004.