Towards a general framework for tactile sensing

Carmelo Sferrazza
Institute for Dynamic Systems and Control, ETH Zurich

I. Background

To fulfill their potential in the manufacturing and retail sectors of the modern world, autonomous machines must be able to perceive and react to contact with their surroundings, both to enhance their capabilities and increase operational safety. To this end, **my research focuses on the contact sensing problem of robotic systems, pivoting on the development of a vision-based tactile sensing principle that provides rich feedback upon physical interaction with the environment.**

In the last decade, the robotics community has experienced extensive growth in tactile sensing research [13, 21], necessitating at least three fundamental aspects to be investigated: i) The scalability and versatility of modern tactile sensors for generic robotic tasks; ii) the processing of raw tactile data to extract quantities of interest for the tasks at hand; iii) the use of raw or processed tactile data for high-level control and estimation tasks. The remainder of this section describes my contributions to addressing the first two aspects. The third aspect is discussed in Section II, where our latest results and future research directions are discussed. Research videos are linked via footnotes, listed at the end of the references.

1) Vision-based tactile sensors for arbitrary surfaces: With the objective of inferring distributed contact information, I designed and fabricated a tactile sensor [8] that is potentially scalable to larger and more complex surfaces. The sensing technique is based on an internal camera that tracks fluorescent particles [1], which are densely and randomly distributed within a soft elastomer. This is in contrast to most vision-based tactile sensors [17, 19, 20], which exhibit regular or sparse patterns. Specifically, the denseness of the particle patterns is effective at creating information about the strain of the soft material at each pixel of the resulting image, and their randomness simplifies manufacture. The simplicity of the fabrication technique enables a straightforward design of sensing surfaces with different shapes and geometry [1, 14], as it does not make any assumptions about the sensing surface, as long as this can be fully captured by an internal camera. My research has shown how larger sensing surfaces[2] can be covered with the use of multiple cameras and efficient data-processing strategies [14].

2) Learning to estimate the contact force distribution: Vision-based tactile sensors provide rich, qualitatively interpretable tactile images. However, extracting relevant physical quantities from such images is challenging, as this generally requires modeling the behavior of the soft materials involved, as well as the optics of the internal camera. Leveraging the high-resolution information provided by the sensor, I developed a **machine learning-based data processing framework** [10]. This is essentially a calibration procedure that utilizes a deep neural network to map the raw tactile images to the three-dimensional force distribution applied to the sensing surface, with the spatial resolution of a human fingertip.

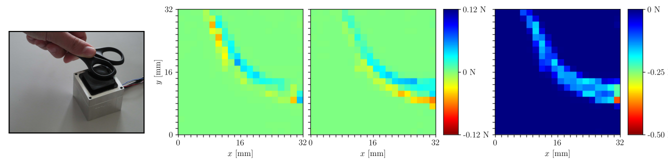

Fig. 1: The shear (first two maps) and pressure (third) components of the force distribution measured during generic contact by the tactile sensor I developed.

3) Calibrating vision-based tactile sensors in simulation: The main concerns typically raised about learning-based approaches to data processing, such as the one I developed for extracting force distribution measurements, are data efficiency and cross-task generalization. To address these issues, I have proposed a solution to generate rich and highly accurate training data in a simulation [11], [9] based on finite-element models and including the camera projection. Hence, the deep neural network described above is entirely trained with synthetic data, avoiding the need for real-world data collection. Employing a strategy based on classical computer vision, the neural network is transferred from simulation to reality, where high sensing accuracy in real time[3] is retained across a variety of contact scenarios, see for example Fig. 1. In addition, I proposed a technique to transfer the calibration mapping across real-world sensors of the same type without further training.

4) Discussion: The data processing framework leverages on the fact that, **by design, the tactile sensor I developed is well-suited for simulation.** In fact, the randomness of the particle patterns tracked by the internal camera requires data processing algorithms that are robust to the pattern variations, such as those introduced when generating synthetic tactile images. In addition, the dense spread of particles showed to considerably improve the contact localization accuracy [8] compared to tracking markers placed at sparser locations on the sensing surface. Finally, the emphasis of my work on estimating general physical quantities, as opposed to directly using the raw camera output, also targets a reusability issue in tactile sensing, where the software and algorithms developed are typically tailored to the specific sensors employed. The estimation of the force distribution provides a physical abstraction from the raw tactile data that aims to facilitate the transfer of high-level robotics algorithms across systems that rely on different tactile sensing principles (e.g., resistive, capacitive, vision-based), which, as such, provide different raw tactile data (e.g., electrical current, capacitance, pixel intensities). The next section provides examples of where such a general abstraction may be leveraged for downstream tasks.

II. Current and future research

1) Tactile control and estimation for dexterous manipulation: To date, most of the achievements using tactile sensing in robotics have focused on showcasing the potential of such a sensory feedback on specific applications [4, 12, 16] based on task-dependent strategies and heuristics. However, the potential of high-resolution tactile sensors for dexterous manipulation can only be fulfilled by developing tailored feedback control strategies that efficiently incorporate the tactile information and can cope with the complex effects arising during non-trivial motions. In this regard, the force distribution abstraction proved to be very practical for planning higher-level robotic tasks. In [1], I led the research on a high-performance manipulation task, where a parallel-jaw gripper provided with tactile sensing and moved by a linear motor aims to swing a pendulum up to a vertical position. Such a proof-of-concept system is similar to a classical inverted pendulum scenario, but rather than being attached to a pivot point, the pendulum (which has unknown physical characteristics) is free to escape the grip at any time. We achieved the task with an off-the-shelf reinforcement learning algorithm, entirely trained on a simulator of the system that we developed. The simulator was based on efficient finite-element and contact modeling techniques, which leveraged the availability of distributed forces. The resulting control policy was successfully transferred to the real-world system[4] without further adaptation.

In such a low-dimensional setting, knowledge about the system was exploited to extract relevant quantities (total force and angular information) from the force distribution, which were then used as inputs to the policy. However, for generic tasks, it may be more challenging to identify such key information beforehand. Therefore, I plan to investigate strategies based on representation learning [15] to automatically extract such specialized features from the general force distribution readings depending on the task of interest. Accurate simulations will remain crucial to alleviate the training requirements of such techniques. As an alternative to representation learning, I will investigate classical filtering techniques, which are relatively unexplored in the field of tactile sensing and may facilitate control tasks by systematically condensing high-resolution tactile information into a state representative of the system considered.

2) Real-time force monitoring for grasping applications: Humans are able to maintain a firm grasp on objects by constantly monitoring slippage and refining the necessary gripping force without damaging such objects. Similarly, tactile sensors show promise to estimate slippage during robotic manipulation tasks to achieve safe and reliable operations. The three-dimensional force distribution provided by the sensor I developed enables the estimation of the stick ratio, defined as the ratio between the sticking and the slipping regions of the sensing surface in contact with an external object. This ratio provides an indication of incipient slip, which can be employed to anticipate the actual slippage of the object. In a recent work [2], this was able to considerably improve the performance in slip prediction[5] compared to standard approaches in a variety of scenarios, including those involving rotational slippage. The force distribution abstraction is fully leveraged in this approach, providing a means to bypass the deformation of the sensor, solely based on the estimated force field. Building on these results, I plan to develop strategies that can adjust the gripping force in real-time based on the available slippage information. In addition, I will further study the case of unsuccessful grasps [3], where I will investigate novel approaches to exploit the history of tactile data and accordingly correct the grasping strategies for the subsequent attempts.

3) Interdisciplinary connections: Only machines enabled with multiple and complementary sensing modalities will be able to ultimately approach a human-level versatility in complex interactive tasks, such as those involving grasping and manipulation. While robotic systems have typically planned most of these tasks pivoting on the availability of rich visual feedback from the environment [7], tactile feedback remains of the utmost importance when it comes to reducing uncertainty during interactions with small or fragile objects, for fine manipulation tasks, or in conditions where visual information deteriorates or is insufficient, such as when coping with the occlusions naturally caused by a grasping motion. In addition to providing force and torque information during contact with external bodies, the force distribution readings also provide local information about the pose of such bodies. The systematic fusion of this information with that provided by vision on a larger scale (such as shown in the context of classical force sensors [6]), will be a necessary step to address the next challenges in the development of autonomous systems that interact with the environment.

However, the applications of tactile sensing extend beyond the development of autonomous robots and have a great potential for human-machine interaction. Haptic feedback technologies are continuously improving the way in which humans interact with machines. As an example, surgical robots can be teleoperated remotely using such technologies, ultimately targeting underdeveloped regions where specialized medical care is not directly available. In fact, surgeons can greatly benefit from touch and haptic feedback [5] during teleoperations, to get a direct sense of the forces they are exerting on the patient, in addition to visual feedback. Finding a connection between distributed tactile forces (as measured through the sensor discussed in this statement) and the corresponding haptic stimuli may provide medical specialists with rich force feedback on their hands, which would further advance their teleoperation capabilities. Similarly, applications of such a connection may find a place in augmented reality systems, by allowing the user to feel haptic stimuli while exploring a virtual environment.

Not only robots and machines benefit from the development of an artificial sense of touch. In fact, the achievements in the miniaturization of tactile sensors, and their scalability to large and arbitrary surfaces [14], will make possible the development of electronic skins that may find applications in smart prosthetic systems [18]. Leveraging advances in neuroscience research, the potential conversion of the tactile feedback to corresponding stimulations in humans may open possibilities to restore tactile sensations to people who have lost limbs.

References

[1] Thomas Bi, Carmelo Sferrazza, and Raffaello D'Andrea. Zero-shot sim-to-real transfer of tactile control policies for aggressive swing-up manipulation. *IEEE Robotics and Automation Letters*, 2021.

[2] Pietro Griffa, Carmelo Sferrazza, and Raffaello D'Andrea. Leveraging distributed contact force measurements for slip detection: a physics-based approach enabled by a data-driven tactile sensor. In *Proceedings of the IEEE international conference on robotics and automation (ICRA)*, 2022.

[3] Francois R. Hogan, Maria Bauza, Oleguer Canal, Elliott Donlon, and Alberto Rodriguez. Tactile regrasp: Grasp adjustments via simulated tactile transformations. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2963–2970. IEEE, 2018.

[4] Francois R. Hogan, Jose Ballester, Siyuan Dong, and Alberto Rodriguez. Tactile dexterity: Manipulation primitives with tactile feedback. In *Proceedings of the IEEE international conference on robotics and automation (ICRA)*, pages 8863–8869, 2020.

[5] Jacqueline K. Koehn and Katherine J. Kuchenbecker. Surgeons and non-surgeons prefer haptic feedback of instrument vibrations during robotic surgery. *Surgical endoscopy*, 29(10):2970–2983, 2015.

[6] Michelle A. Lee, Yuke Zhu, Krishnan Srinivasan, Parth Shah, Silvio Savarese, Li Fei-Fei, Animesh Garg, and Jeannette Bohg. Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 8943–8950, 2019.

[7] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.

[8] Carmelo Sferrazza and Raffaello D'Andrea. Design, Motivation and Evaluation of a Full-Resolution Optical Tactile Sensor. *Sensors*, 19(4), 2019.

[9] Carmelo Sferrazza and Raffaello D'Andrea. Sim-to-Real for High-Resolution Optical Tactile Sensing: From Images to Three-Dimensional Contact Force Distributions. *Soft Robotics*, 2021.

[10] Carmelo Sferrazza, Adam Wahlsten, Camill Trueeb, and Raffaello D'Andrea. Ground Truth Force Distribution for Learning-Based Tactile Sensing: A Finite Element Approach. *IEEE Access*, 7:173438–173449, 2019.

[11] Carmelo Sferrazza, Thomas Bi, and Raffaello D'Andrea. Learning the sense of touch in simulation: a sim-to-real strategy for vision-based tactile sensing. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4389–4396, 2020.

[12] Yu She, Shaoxiong Wang, Siyuan Dong, Neha Sunil, Alberto Rodriguez, and Edward Adelson. Cable manipulation with a tactile-reactive gripper. In *Proceedings of Robotics: Science and Systems (RSS)*, 2020.

[13] Kazuhiro Shimonomura. Tactile Image Sensors Employing Camera: A Review. *Sensors*, 19(18), 2019.

[14] Camill Trueeb, Carmelo Sferrazza, and Raffaello D'Andrea. Towards vision-based robotic skins: a data-driven, multi-camera tactile sensor. In *Proceedings of the IEEE International Conference on Soft Robotics (RoboSoft)*, pages 333–338, 2020.

[15] Michael Tschannen, Olivier Bachem, and Mario Lucic. Recent advances in autoencoder-based representation learning. *Bayesian Deep Learning Workshop at NeurIPS*, 2018.

[16] Chen Wang, Shaoxiong Wang, Branden Romero, Filipe Veiga, and Edward Adelson. SwingBot: Learning Physical Features from In-hand Tactile Exploration for Dynamic Swing-up Manipulation. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5633–5640, 2020.

[17] Benjamin Ward-Cherrier, Nicholas Pestell, Luke Cramphorn, Benjamin Winstone, Maria Elena Giannaccini, Jonathan Rossiter, and Nathan F. Lepora. The TacTip Family: Soft Optical Tactile Sensors with 3D-Printed Biomimetic Morphologies. *Soft Robotics*, 5(2):216–227, 2018.

[18] Yuanzhao Wu, Yiwei Liu, Youlin Zhou, Qikui Man, Chao Hu, Waqas Asghar, Fali Li, Zhe Yu, Jie Shang, Gang Liu, et al. A skin-inspired tactile sensor for smart prosthetics. *Science Robotics*, 3(22):eaat0429, 2018.

[19] Akihiko Yamaguchi and Christopher G. Atkeson. Combining finger vision and optical tactile sensing: Reducing and handling errors while cutting vegetables. In *Proceedings of the IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 1045–1051, 2016.

[20] Wenzhen Yuan, Siyuan Dong, and Edward H. Adelson. GelSight: High-Resolution Robot Tactile Sensors for Estimating Geometry and Force. *Sensors*, 17(12), 2017.

[21] Liang Zou, Chang Ge, Z. Jane Wang, Edmond Cretu, and Xiaoou Li. Novel Tactile Sensor Technology and Smart Tactile Sensing Systems: A Review. *Sensors*, 17 (11), 2017.

[1] Video 1: `https://youtu.be/yuwFYrgjMZM`
[2] Video 2: `https://youtu.be/lbavqAlKl98`
[3] Video 3: `https://youtu.be/dvOk2XrSmLE`
[4] Video 4: `https://youtu.be/In4jkaHzJLc`
[5] Video 5: `https://youtu.be/YeotGbKVWcY`