Whole-Body Robot Skins for Near-Body Perception and Safe Interaction

by

Caleb Sebastian Escobedo

B.S., Trinity University, 2019

M.S., University of Colorado Boulder, 2025

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Computer Science

2026

Committee Members:

Alessandro Roncone, Chair

Nikolaus Correll

Christoffer Heckman

Zackory Erickson

Michael Gleicher

Escobedo, Caleb Sebastian (Ph.D., Computer Science)

Whole-Body Robot Skins for Near-Body Perception and Safe Interaction

Thesis directed by Prof. Alessandro Roncone

Human skin is a distributed sensory organ that lets us perceive proximity, contact, and the structure of the world across the entire body, an ability so fundamental to acting safely among people that robots cannot truly share our spaces without it. Yet robots still rely on sparse, task-specific sensors that leave most of their body blind, and the whole-body skins that could close this gap remain hand-built, one-off systems confined to a handful of labs. This dissertation argues that the central challenge in physical human-robot interaction is making whole-body sensing something that can be fabricated, grounded on the body, and deployed repeatably across any robot. I address this challenge at its foundation, in fabrication. I present a method for producing whole-body robotic skins through multi-material additive manufacturing (3D printing), integrating diverse sensing modalities into rigid and compliant structures that conform directly to any robot surface. Because sensor locations are set in the digital model and preserved by the print, a robot knows where its sensing data originates on its own body without per-sensor recalibration. Using this method, I fabricate skins spanning capacitive, time-of-flight, hybrid, and magnetic sensing and deploy them across a humanoid, a quadruped, and a robot arm. Building on this foundation, I show how the distributed sensing on these skins is calibrated and integrated into perception of the near-body space, and how object-aware controllers use that perception to anticipate and avoid contact during physical interaction. Together, these contributions turn whole-body sensing from a one-off engineering effort that only a few labs can take on into a capability any roboticist can reproduce and extend. What has held whole-body sensing back is what each system demands to build: specialized hardware knowledge, custom wiring, hand-placed sensors, and months of integration for a result that rarely transfers to another robot. Removing that burden reframes whole-body perception as shared infrastructure for robotics, a foundation that researchers in manipulation,

learning, humanoids, and human-robot interaction can build on without reinventing the hardware, and a concrete step toward robots that leave their isolated cells to work safely alongside people in the everyday physical world.

Dedication

To my older brother Francisco "Pancho" Enrique Escobedo for keeping our family together and for sharing his love of technology with me.

Acknowledgements

This work would not have been possible without the love and support of my advisor, family, friends, and collaborators. I would first like to thank my advisor, Alessandro Roncone, for his support throughout this entire process. Thank you for having confidence in my ability to pursue original research, for providing me with the time and resources necessary to explore ambitious ideas, and for encouraging me to build a team around my research. Your guidance and trust allowed me to grow as both a researcher and a mentor. I would also like to thank my mom, Linda, for constantly encouraging me, as a dyslexic child, to continue exploring my passion for technology. I thank my younger brother, Diego, for being my partner as we both pursued careers in computer science. I am also grateful to my sister-in-law, Dr. Morgan Fisler, and my nephew, Marcelo Vicente Escobedo, for giving me the motivation and support to take chances and pursue robotics research that may one day positively impact Marcelo's life. I would also like to thank my grandma, Elvira "Bibi" Ruiz, who helped raise me from the time Diego was born, as well as my godparents, Rosaisela and Enrique, and their children, for their enduring love and support. I am deeply grateful to my friends Josh, Amy, Jake, Hudson, and Remington McMillen for supporting me throughout this journey. I thank my partner, Mary Martin, for undertaking the challenges of a PhD alongside me and for her constant love and encouragement. I also thank my friend Haden, whose support helped me restart my education and begin the path that ultimately led me here.

Lastly, I would like to thank my dad, Abram Escobedo, who passed away before I began this journey. His life and passing played an important role in my decision to pursue this path, and his constant reminders that he was proud of me have stayed with me throughout it.

Contents

Tables

Table

Figures

Figure

Chapter 1

Introduction

The next generation of robots will not stay behind fences. They are being asked to work in homes, hospitals, warehouses, and other spaces built for and occupied by people, where motion is dynamic, surfaces are cluttered, and the people nearby act in ways no plan can anticipate. In these spaces a robot interacts with the world using its whole body, not just a gripper at the end of its arm; any part of its body can meet an obstacle, a surface, or a person at any moment. Whether a robot can operate there safely comes down to how much of its own body it can sense, and how well it can act on what it senses before contact turns into harm.

Human skin is the benchmark for what robots are missing. It is distributed, compliant, and information-rich, sensing contact and the nearby world at once while adapting to the body at every scale. Robots have no equivalent, and building one is increasingly recognized as foundational to machines that can perceive and act through physical contact with the world [37].

Research in tactile and proximity sensing has produced strong results in manipulation, collision avoidance, and physical interaction. The harder problem has been extending that sensing across the entire body. Hand fabrication is still common, and many designs rely on small sensing units, custom PCBs, platform-specific attachment, or a single sensing modality. These choices can produce capable demonstrations, but they make skins difficult to scale, reproduce, and move between robots.

Proximity sensing is central to this problem. The region within roughly 20 cm of a robot's surface is where close interaction is decided, yet it is largely unobserved by the sensors robots normally carry: cameras are occluded by the robot's own body, and force and tactile sensors report contact only after it has happened. Sensing this near-body space directly from the robot's surface is what lets a robot slow, redirect, or prepare for contact before it becomes a collision.

The cost is paid in research progress. Whole-body skins are built as one-off systems for a single robot, task, or sensor type, and are rarely shared in a form another lab can reuse.

This dissertation focuses on that gap. The central question is how to build whole-body robot skins that can be placed where sensing matters, localized accurately enough on the body to use their signals, and adapted to new robot platforms without rebuilding the process each time. This cannot be solved by sensor design alone. It requires a foundation that connects fabrication, sensing, calibration, and control into a process other researchers can adopt, adapt, and extend.

My work builds toward this argument through three threads that move from hardware to application. First, *fabrication* shows how multi-material additive manufacturing brings whole-body robot skins together into repeatable, form-fitting hardware, integrating diverse sensing modalities, including SPAD-based optical proximity sensing, conductive-filament capacitive sensing, hybrid sensing, and magnetic sensing, while preserving where each sensor sits on the robot body. Second, *perception* shows how the distributed sensing on these skins is calibrated and localized on the robot body, and how onboard proximity sensing and third-person sensors cover different parts of the environment and become more powerful when used together. Third, *application* shows how this perception supports real-time, object-aware control that reduces contact forces and transitions smoothly between avoidance and contact.

This dissertation shows that multi-material additive manufacturing is a foundation for whole-body robot skins, producing skins that are repeatable, scalable, and compatible with diverse sensing modalities across robot platforms. The larger goal is to move whole-body sensing from a sequence of one-off systems, each rebuilt from scratch, toward a shared capability the field can adopt and extend, so that giving a robot a sense of its whole body becomes a routine starting point rather than a research project of its own.

1.1 Research Questions

The research in this dissertation is published across a series of peer-reviewed papers on scalable fabrication methods for robotic skins, whole-body sensing, and contact-aware control. It is organized around three research questions that follow a hardware-to-application progression: the first establishes the fabrication foundation, the second grounds distributed sensing on the robot body, and the third uses that sensing for safe physical interaction.

In Chapter 2, I examine how multi-material additive manufacturing can make whole-body robotic skins reproducible and scalable. This chapter centers on a fabrication process that turns robot-specific digital skin geometries into functional, deployable skins with known sensor placement, integrating capacitive, time-of-flight, hybrid, and magnetic sensing across a humanoid, a quadruped, and a robot arm. Here I address **(RQ1) How can multi-material additive manufacturing enable the reproducible fabrication of whole-body robotic skins with diverse sensing modalities?**

In Chapter 3, I turn to how the distributed sensing on these skins is made usable. Sensors only inform control when the robot knows where each reading originates on its body and can integrate those readings with external perception. This chapter addresses **(RQ2) How can distributed onboard sensing be calibrated and integrated to support reliable nearby-space perception?**

Finally, in Chapter 4, I examine how nearby-space perception is put to use. Object-aware controllers turn perception into motion, anticipating and avoiding contact while preserving task execution, and admitting contact when it is useful. This chapter addresses **(RQ3) How can this perception be utilized by object-aware controllers to enable safe, kinematically feasible contact anticipation and avoidance?**

Chapter	Research Question
Chapter 2 Fabrication	**RQ1**: How can multi-material additive manufacturing enable the reproducible fabrication of whole-body robotic skins with diverse sensing modalities?
Chapter 3 Perception	**RQ2**: How can distributed onboard sensing be calibrated and integrated to support reliable nearby-space perception?
Chapter 4 Application	**RQ3**: How can this perception be utilized by object-aware controllers to enable safe, kinematically feasible contact anticipation and avoidance?

1.2 Background

In this section I overview the state of the art (SotA) in Robot Skin, Sensing, 3D Printing, and Contact anticipation in robotics applications. This thesis focuses on the combination of these subtopics into a singular bottleneck focused on the integration of onboard sensing to generic robot platforms. The primary takeaway from each of the sections respectively is 1. RobotSkin: there is no generic method for equipping a robot with arbitrary onboard sensing such that the sensor is localized on the body, calibrated to sense a specific stimulus, or make for the specific robot's morphology. 2. Sensing: Current SotA sensing does not transfer to robotic systems due to difficulties in integration, specifically for full body coverage of large surface areas, one-off systems do not transfer in terms of fabrication cost, time, or integration into the robot system. 3. 3D printing: Fabrication of digitally designed robot skins need to be instantiated into a physical form, here I leverage the SotA in 3D printing to create multi-material skins that can have embedded wiring, rigid and soft materials, known locations to mount third-party sensors, and reduce the overall burden placed on the researcher. 4. Contact Anticipation: We then look at the SotA in robot contact anticipation for robotic manipulators, their limitations due to sensing, and their ability to avoid objects in fully observed settings. A limiting factor for these systems in the real-world is their estimation of the environment directly around their body, an area that is normally completely unobserved by commonly included sensor suites such as cameras, force/torque sensors, or tactile sensors placed on the surface. We are primarily concerned with the close-proximity area around the robot's body which we define as less than 20 cm from the robot surface.

1.2.1 Robot Skin

The focus on robot whole-body sensing has largely focused on the integration of tactile sensing into whole-body robot systems, to mimic human tactile sensations and understand the extent to which humans rely on touch in their daily lives [12]. Reviews of tactile skin emphasize that the harder problem is often not the sensing element itself but effectively utilizing skin at scale: integrating, wiring, and interpreting large numbers of sensors distributed across a body [13]. Robots are not limited to the sensing capabilities that are available to the human form. In this thesis humans and animals serve as a model of what can be done when sensor information is leveraged for action prediction in a variety of cases [63]. When following humans as the model for sensorization of robot systems, the focus has largely been on tactile and small scale sensor suites for physical manipulation such as fingertip or robot gripper sensors that are used to explicitly make contact with an object [61, 74, 48, 94, 26]. These types of sensors, specifically visuotactile sensors have been shown to give exceedingly high resolution for the forces and geometric representations of the object being touched [18, 19]. Recent SotA fingertip sensors are introducing super-human levels of touch through multi-modal sensor integration [49].

Moving from the fingertip toward whole-body coverage, flexible fabric-based sensors have been developed as a conformable skin for humanoid robots, characterizing how a textile substrate can host distributed tactile sensing over large body areas [34]. Beyond a single fabric substrate, whole-body electronic skins have been realized on humanoid platforms. A modular, self-organizing, multi-modal skin can cover a humanoid's upper body and automatically discover the layout of its own sensor cells [58, 59, 10], and other modular skins colocate computation with sensing to report proximity, contact, and force at low bandwidth for collision avoidance and affective touch recognition [36], building on distributed soft skins that sense and localize textures [35]. A large-scale, flexible capacitive tactile system of this kind covers the iCub humanoid with conformable patches, built on methods developed for implementing large-scale robot tactile sensors [8, 77, 55], and it is the iCub's capacitive skin that underlies the whole-body representation work discussed next. Whole-body skin has also been used to represent the space around the robot rather than only the point of contact. By anchoring a visual receptive field to each skin taxel and recording whether an approaching object ends in contact, a robot learns a distributed, visuo-tactile representation of the space around its body that drives both avoidance and reaching [73]. This was developed into a biologically grounded model of a margin of safety around the whole body, adding motor equivalence so that any body part can act as the controlled endpoint [72]. The representation was then turned into a compact controller for real-time avoidance during human-robot interaction [66]. When placing sensing on the body of a robot, one of the main issues with integration into real-world systems is self-detection. Prior work with ToF sensors has focused on detection of objects near the robot, while not detecting the robot itself. In this work they collect an ambient histogram from the ToF sensor, which manifests as a DC offset based on the ambient light and sensor itself. They use data from the robot with no objects nearby to subtract out the self-detection signal showing that these sensors can be used to find only objects near the robot surface [79]. Whole-body sensing has also been pursued through dense visual coverage rather than skin, distributing many cameras across a robot to achieve whole-body dexterity [96].

1.2.2 Sensing

Robotics applications commonly rely on mature, widely available sensing technologies: photon-based time-of-flight (ToF), capacitance, acoustics, and vision. This section covers the operating principles of each and their use in robotics. One of the primary sensing modalities used within this thesis is photonic ToF sensing which gives a depth range estimation based on the amount of photons captured over a discrete time histogram. Here we detail applications and advancements in this kind of ToF range sensors. Low-cost single-photon avalanche diode (SPAD) cameras are extending this photon-counting depth sensing toward broadly accessible 3D vision [60]. Photon-based ToF sensors have been used to determine the 3D geometry of known objects from *very few* proximity depth measurements [80]. Alternatively to recovering geometry, the raw temporal data captured by these sensors' transient histograms encode information that standard distance estimates discard.

It has been shown that using full transient histograms rather than single-point distance summaries improves detection of geometric deviations such as obstacles or surface discontinuities [81]. Their method fits a Gaussian Mixture Model to histograms captured from a known planar surface, then flags new measurements as deviations when they fall outside the learned distribution. A key challenge they identify is the ambiguity between geometry and surface reflectance: a histogram bin's intensity depends on both how much surface area lies at that depth and the albedo of that surface, making the two indistinguishable from a single measurement without prior knowledge of the scene. Optical proximity sensing can also be embedded in compliant material to recover contact force alongside distance. Commodity infrared proximity sensors set in a transparent elastomer measure range as an object approaches, and on contact the sensor reads the elastomer's deformation as an estimate of applied force, giving both pre-touch proximity and contact force from one element [68].

Capacitive methods sense the change an object induces in an electric field, and appear in two forms. In self-capacitance sensing, a single electrode's capacitance to ground shifts as an object approaches, so the same element reports both pre-touch proximity and contact; this has been used to interface collaborative robots with a flexible sensor that pairs self-capacitance with force sensing [28]. In mutual-capacitance sensing, the coupling between transmit and receive electrode pairs changes with touch, which has been embedded as an electrode matrix in a compliant, human-like artificial skin for physical human-robot interaction [84]. Because the measured field depends on a target's permittivity as well as its distance, capacitive sensors can also recognize the material of a nearby object rather than only its range [16, 2]. This combination of proximity and contact sensitivity has made capacitive sensing attractive for safe human-robot interaction, from enforcing safety margins around a manipulator [62] to close-range assistive tasks such as robot-assisted dressing and bathing [21].

Active acoustic methods emit sound and interpret the returning field. AuraSense introduced leaky surface waves for full-surface, blindspot-free proximity detection on off-the-shelf arms using only a single transmitter and a few receivers [24], and AmbiSense extended this acoustic-field approach to also estimate the bearing of an object around the sensor [75]. Active acoustics has also been applied inside soft pneumatic actuators, where emitted sound is interpreted to sense contact and deformation of the actuator itself [99]. Surface-based acoustic methods instead listen to the waves produced by contact, as in AcouSkin, which localizes contact across a full surface using acoustic waves [47, 46, 51]. Acoustic transduction can also be used for pre-touch sensing at the fingertip, detecting an object just before contact to trigger reactive pre-grasping [74, 76]. Acoustic sensing is attractive because a small number of transducers can cover a large surface, but it is sensitive to material properties and environmental acoustics, which complicates transfer across robot bodies.

Vision-based tactile sensing infers contact from images of a soft surface as it deforms against an object. In GelSight-style sensors, a camera views an elastomer coated with a reflective membrane under directional colored lighting, and photometric stereo reconstructs high-resolution contact geometry while displacement of embedded markers captures force [98]. Related optical sensors instead learn to reconstruct dense 3D shape directly from the imaged deformation, with compact designs targeting tasks such as grasping multi-scale objects from flat surfaces [17]. Recent work scales this imaging principle from the fingertip to entire robot links, placing a wide-angle camera inside a hollow link to view the deforming inner surface of a soft skin covering it and recovering contact location and force across the whole link [86, 54].

1.2.3 3D printing

Consumer multi-material 3D printers can now combine rigid structure, compliant layers, and conductive traces in a single print [31]. Conductive polymer composites, such as carbon-filled or graphene-filled thermoplastics, enable piezoresistive and capacitive sensing directly from printed material. Piezoresistive elements change resistance under mechanical strain; capacitive elements use deformable dielectric layers between conductive plates to sense pressure. These mechanisms require no post-process assembly when the conductive and structural materials are co-printed. This co-printing of function into structure reflects a broader vision of robotic materials that couple sensing, actuation, computation, and communication within the fabricated body itself [56]. Recent reviews document the rapid growth of 3D-printed electromechanical sensors across these mechanisms, comparing their performance and pointing to increasingly mature fabrication pipelines [95]. Integration strategies range from multi-material printing, where conductive and non-conductive filaments are deposited in the same job, to conductor infusion, where channels are printed and later filled with conductive ink, to hybrid approaches that pause the print to insert off-the-shelf components. Robot skins need all three properties: structure to mount on the robot, compliance to conform to surfaces and absorb contact, and sensing to detect proximity or pressure. Building these separately means wiring sensors into a structure after the fact, which introduces placement uncertainty and limits how reliably the skin can be reproduced. When structure, compliance, and sensing come from the same print, sensor locations are fixed by the geometry file and the result is repeatable across fabrication. Fully 3D-printed tactile sensors have been demonstrated directly on complex, curved geometries rather than flat test coupons, printing conformal, multi-directional sensing onto shapes such as fingertips and bone surfaces [93, 97].

Computational design pipelines have emerged to generate printable geometries with embedded functionality. Early work introduced pipelines that jointly consider form and sensing during design, producing 3D printed objects with integrated capacitive touchpoints [6]. Subsequent work reduced wiring complexity by optimizing internal conductive paths to create unique resistor-capacitor delays, enabling multiple touch regions to be sensed through a single wire [5]. Our GenTact work extends this co-design approach to robot skins, combining procedural geometry generation with multi-material printing to produce form-fitting tactile sensors that conform to specific robot links [40]. Capacitive touch sensing has also been demonstrated directly on general 3D surfaces using mutual capacitance, distributing transmit and receive electrodes over arbitrary printed geometry, in contrast to the self-capacitance sensing used primarily though this thesis [67].

Beyond HCI, design tools are enabling more expressive control over printed geometry and material distribution. Volumetric multi-material compilers allow designers to define both shape and material gradients using implicit functions, scaling to printing systems that require billions of voxels [89, 88, 87]. Generative AI methods can automatically segment meshes into functional and aesthetic components, enabling stylization without compromising structural integrity [25]. These tools shift the bottleneck from what printers can physically deposit toward what designers can specify and what software can compile.

The trajectory of this work points toward fully functional prints with geometry and material distributions that are difficult to imagine today. Functionally graded lattices, simulation-informed structures, and embedded electronics are already possible on research-grade systems. As these capabilities become accessible on consumer hardware, fabrication pipelines will need to keep pace. Chapter 2 develops one such pipeline for robot skins, using multi-material printing to produce whole-body skins with tactile, proximity, hybrid, and magnetic sensing modalities.

1.2.4 Contact Anticipation

Work in the early 1990s demonstrated that proximity sensors covering the entire arm could enable real-time obstacle detection without prior knowledge of the environment [92, 53, 11]. These systems wrapped arrays of infrared sensors around the robot body and fed distance readings directly into motion planners, producing collision-free paths in uncertain, dynamic environments. Later work introduced modular sensor units that combined capacitive and magnetic sensing with local signal processing, allowing the same hardware to be deployed across different robot platforms without redesigning the skin for each application [85]. This early push toward modularity and reusability foreshadowed a persistent challenge in robot skin research: building sensing systems that transfer across robots rather than remaining one-off prototypes, which is addressed directly in this thesis.

Modern approaches formulate collision avoidance as constrained optimization rather than reactive sensor-driven motion. One class of methods computes repulsive vectors from external depth cameras and applies them at control points along the robot body [27]. Another uses onboard proximity skins with capacitive and time-of-flight sensors to constrain approach velocities toward obstacles while leaving motion away from obstacles unrestricted [14]. In industrial settings, distributed distance sensors have similarly been used to scale robot speed as a function of measured proximity to nearby people, enforcing safety constraints during shared operation [4]. On humanoids, egocentric depth sensing distributed on the body has been used for whole-body collision avoidance and motion planning in cluttered space [39]. Reference governor methods filter commanded references using external motion capture or depth cameras to ensure state and input constraints are satisfied in real time without requiring online optimization [57]. A related line of work poses motion generation as a real-time weighted multi-objective optimization. RelaxedIK synthesizes accurate and feasible arm motion by trading off pose tracking against feasibility objectives [70], CollisionIK extends this by incorporating collision avoidance as a weighted term, assuming obstacle positions are known from external perception [71], and RangedIK generalizes the formulation to ranged-goal tasks that free redundant degrees of freedom for secondary objectives [90]. Other reactive controllers work directly in the velocity domain: a purely reactive, manipulability-maximizing controller resolves motion while avoiding joint limits and singularities [32], and NEO casts avoidance as a quadratic program with explicit collision-avoidance constraints for real-time manipulator control [33]. Dynamical systems approaches offer reactive motion generation through modulation and control barrier functions, often relying on joint kinematics for self-collision or external sensing for environment collision [44]. Beyond pure avoidance, proximity sensing can also be used to proactively adjust robot impedance before contact occurs, reducing impact forces when collision is unavoidable [15].

Recent work has turned to learning-based methods. Neural networks can learn implicit signed distance functions in joint space, providing differentiable collision costs that integrate directly into optimization-based controllers [45]. Configuration-space distance field networks unify self-collision and environment collision into a single geometric representation, enabling zero-shot generalization to unseen environments [9]. Reinforcement learning methods combine potential fields with policy gradients to navigate cluttered 3D workspaces [78]. These approaches shift collision checking from explicit geometric computation to learned representations that scale better with scene complexity.

These methods share a common assumption: contact is a failure condition to be avoided. They succeed when the robot does not touch the obstacle. That framing fits industrial settings where humans and robots are separated. It does not match close-proximity interaction, where contact is sometimes unavoidable and may be useful. A robot that can only avoid obstacles cannot hand an object to a person, respond to a guiding touch, or work collaboratively in cluttered space. A more established line of work handles contact after it occurs: using only proprioceptive sensing, model-based methods detect that a collision has happened, isolate where on the body it occurred, and identify it well enough to trigger a safe reaction [30]. A smaller body of work treats contact as something to interpret rather than prevent. Data-driven methods classify whether a detected contact is intentional or accidental, then switch between compliant admittance control and evasive action [52]. Tactile approaches use hand-shape detection to distinguish voluntary guidance from unintended collisions [29]. These methods show that the same physical event can mean different things depending on context, but they still respond after contact has already occurred.

Contact anticipation reframes the problem further. Instead of classifying contact after it occurs, the robot uses nearby-space perception to prepare for contact, detect when it happens, and respond appropriately. This requires sensing that can detect objects before they touch the robot, control that can modify motion based on proximity, and behavior that transitions smoothly between avoidance and physical interaction. Chapter 4 develops this argument through a contact anticipation framework that uses onboard proximity sensing to anticipate, detect, and react to contact during physical human-robot interaction.

Chapter 2

Multi-Material Additive Manufacturing of Robot Skins

2.1 Overview

Whole-body sensing has stayed the privilege of a few labs for one practical reason: robot skins are too hard to build, which makes fabrication, more than any sensing principle, the step that decides whether the rest of the field can deploy whole-body sensing at all. This dissertation begins with fabrication because it is the foundation everything else depends on. Whole-body robot skins combine sensing with mechanical structure, reliable attachment, known sensor placement, signal routing, and system integration. Multi-material additive manufacturing (3D printing) fulfills these requirements in a single process, fabricating structural, conductive, and compliant features together. This shifts robotic skins away from one-off manual builds and toward systems that can be reproduced, modified, shared, and deployed across robots.

Two properties make this foundation work. First, a single print can fit any robot geometry, producing form-fitting skins for platforms as different as a humanoid, a quadruped, and a robot arm. Second, because sensor locations are set in the digital model and preserved by the print, a robot knows where its sensing data comes from on its own body without per-sensor recalibration.

The work discussed in this chapter builds on collaborative design pipelines that generate robot-specific skin geometries. The contribution emphasized here is the multi-material printing process that turns those geometries into functional skins with conductive regions, structural bodies, compliant layers, embedded routing, sensor mounts, and electrical connection points. This frames fabrication as a practical requirement for deploying whole-body robotic skins beyond iso-

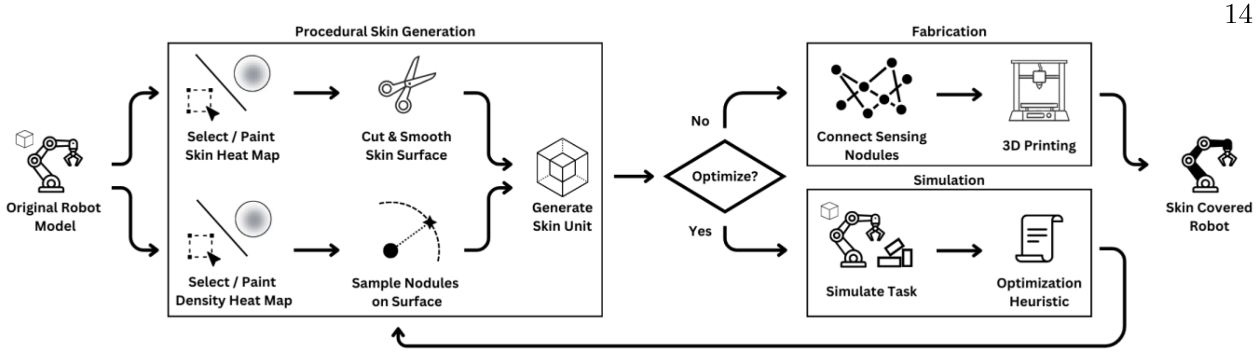

Figure 2.1: The GenTact pipeline for designing form-fitting and adaptable tactile skins is composed of three stages: procedural generation, simulation, and fabrication. The procedural generation stage (left) generates an initial distribution of sensors that are then passed into the simulation stage (bottom right) to be evaluated and improved based on the task they are used for. Finally, the sensors are connected internally in the fabrication stage (top right) to be printed and deployed on the real robot.

lated demonstrations.

Our work, GenTact Toolbox [40], provides the initial context for this fabrication approach. It combines procedural mesh generation, task-driven simulation, and multi-material 3D printing to create form-fitting whole-body tactile skins for robots. As shown in Fig. 2.1, GenTact converts generated skin units into deployable capacitive sensors. This chapter focuses on that fabrication step as a repeatable method for producing robot-specific sensing hardware.

Our subsequent work extends this fabrication approach. Form-fitting printed mounts, for off-the-shelf sensors, demonstrate how fabricated skin units can preserve sensor placement and reduce manual calibration [82]. GenTact-Prox shows that fully 3D-printed skins can integrate tactile and proximity sensing for collision avoidance and body-centered perception [43]. Our multi-modal soft-material skin in Fig. 2.2 combines self-capacitance, time-of-flight sensing, conductive traces, and compliant TPU coverings in a deployable skin architecture [41]. Together, these works show that multi-material additive manufacturing is a practical method for producing deployable robotic skins across different sensing modalities, including a new magnetic skin introduced later in this chapter.

Figure 2.2: The hybrid sensors are composed of three layers: 1) self-capacitance sensors, 2) ToF sensor mounts, and 3) compliant soft covering.

2.1.1 Research Question

> **RQ1:** How can multi-material additive manufacturing enable the reproducible fabrication of whole-body robotic skins with diverse sensing modalities?

RQ1 has two parts. The first is fabrication reliability and transferability: a robot skin should support repeated fabrication, reliable installation on the intended robot body, preserved sensor locations, and connection to a robot system with minimal manual adjustment. The second is fabrication flexibility: a manufacturing method for robot skins should combine structural layers, conductive sensing elements, rigid sensor mounts, compliant coverings, and interfaces to external electronics, and should support more than one sensing modality. Throughout this chapter, the emphasis is on the fabrication pathway that turns robot-specific digital skin geometries into physical, functional, and deployable sensing systems.

2.2 Motivation: From One-Off Robot Skins to Reproducible Fabrication

Whole-body robot skins are often motivated by their ability to extend sensing across the robot body, enabling contact detection, proximity perception, and interaction-aware control. Practical deployment also depends on manufacturability, repeatability, physical compatibility with the robot, and ease of integration.

Fabrication remains a limiting factor for whole-body robot skin research. Custom manual builds, hand wiring, individually placed sensors, and post-hoc calibration keep the field dependent on prototypes that are difficult to compare, replicate, or scale. Multi-material additive manufacturing addresses this limitation through a digital workflow that generates robot-specific geometry, assigns functional materials, and prints the skin as a repeatable physical object.

2.2.1 Limitations of One-Off and Manually Integrated Skins

Many whole-body sensing systems are limited by the difficulty of fabricating and integrating sensors over complex robot surfaces. Robot bodies are not simple flat planes. They include curved, nondevelopable surfaces, moving joints, link-specific geometry, and areas where sensors must not interfere with actuation. Traditional approaches often address this by using modular sensors, flexible electronics, or custom mounts. These methods can be effective, but they frequently require manual assembly and careful sensor placement after fabrication. A recent review of robotic tactile perception makes this concrete: whole-body tactile skins remain costly and complex enough that only a few research groups worldwide have deployed them, and because large-area skins are typically fabricated as flexible planar circuits that are then cut and wrapped onto curved surfaces, the real placement of the sensing elements diverges from the intended design, leaving accurate manual calibration of the thousands of taxels on a complex body effectively infeasible [3]. Manual integration introduces variability in sensor placement, additional calibration requirements, and barriers to reproducibility. Small differences in position or orientation can affect proximity estimation, contact localization, and control. Post-installation calibration can reduce this uncertainty, but it adds time and potential error. Even when the circuit design or sensing principle is published, systems that depend on tacit assembly knowledge remain difficult to reproduce.

The form-fitting sensor mounting work addresses this problem by generating printed mounts that preserve sensor position and orientation relative to the robot body [82]. As shown in Fig. 2.3, a robot model and sensor board model are used to create a skin unit with snap-in mounts, allowing the mounted sensors to reconstruct nearby objects as point clouds. The key manufacturing insight is that the printed artifact itself can constrain sensor placement.

2.2.2 Multi-Material Additive Manufacturing as an Enabling Method

Multi-material additive manufacturing combines functions that are often separated across fabrication steps. Robot skins may require structure, conductive sensing regions, embedded routing, sensor mounts, protective layers, and electrical interfaces. Conventional workflows often distribute these features across circuit boards, molded parts, wires, adhesives, fasteners, and manual assembly. In a multi-material additive manufacturing workflow, many of these features can be fabricated directly into the skin or aligned through the printed geometry.

The GenTact Toolbox demonstrates this approach through multi-material 3D printing of

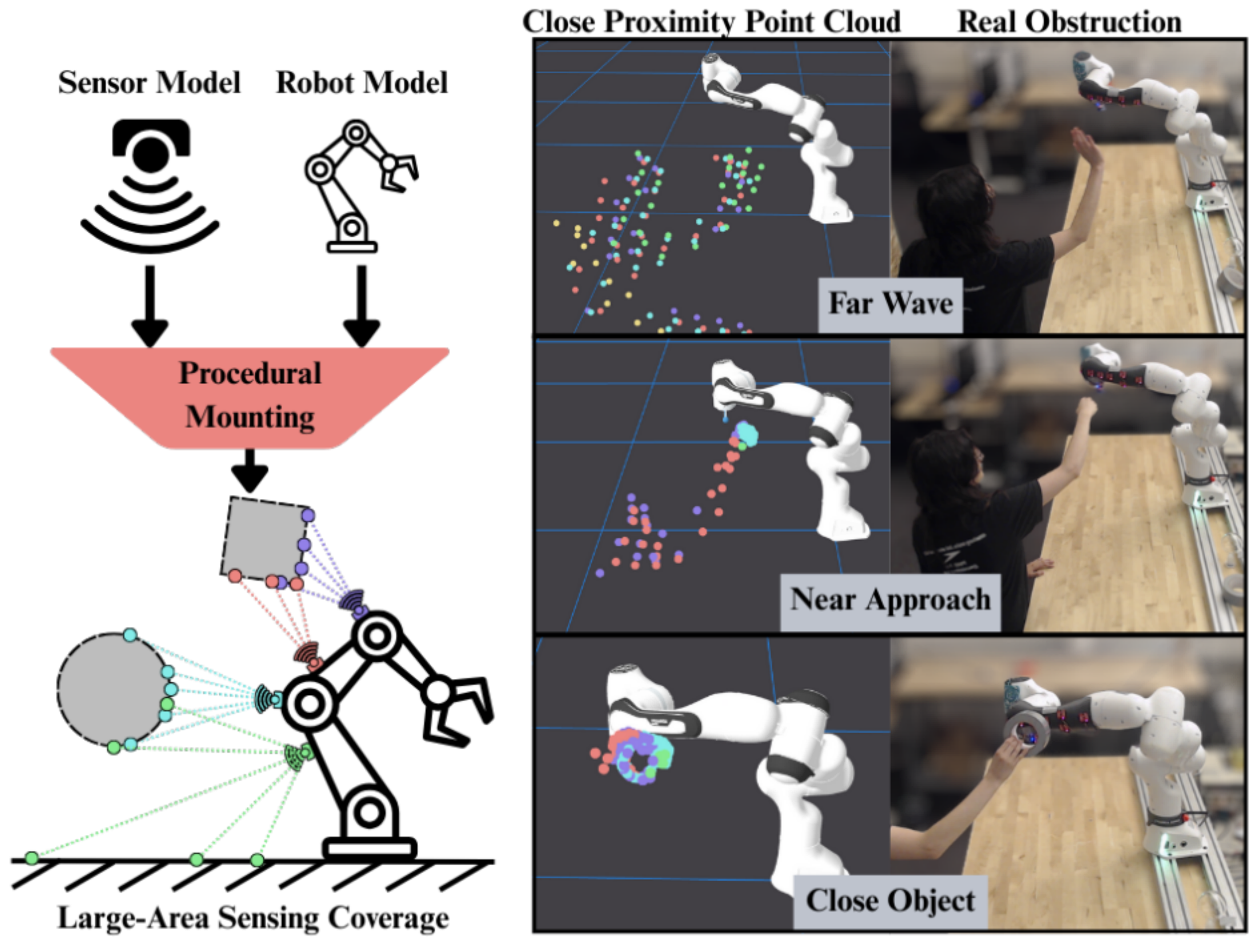

Figure 2.3: (Left) A diagram showing the procedure of digitally designing a skin unit to its deployment on a robot. A model of a sensor board and a robot model are inputted into the procedural mounting step where we use 3D graphics software to fit a skin cover around a specified link. (Right) A comparison of the point cloud representation of detected obstacles versus their position relative to the Franka Research 3 in real life. Different point colors represent depth measurements from different sensors. "Far Wave" shows a person waving their hand approximately 14 cm away from the skin unit. "Near Approach" shows a person putting their fist close to the skin unit. "Close object" shows a roll of tape situated from the sensors at close proximity. In the point cloud reconstruction of the tape roll, the hollow center of the tape is clearly defined.

procedurally generated capacitive tactile skins [40]. In Fig. 2.1, fabrication is the final stage that converts digital skin units into physical sensors. Fig. 2.4 shows this progression for a Franka Research 3 link, where the skin geometry is generated from the robot model, conductive sensing regions are traced into the design, and the physical skin unit is printed and deployed on the robot.

Multi-material additive manufacturing translates digital skin designs into functional robot hardware through shareable files, print parameters, and material choices. This reduces reliance on hand assembly and makes the fabrication process easier to repeat.

2.3 Multi-Material Fabrication of Functional Robotic Skins

This section focuses on how procedurally generated geometries are fabricated as functional skins. Multi-material printing preserves the form-fitting body while incorporating conductive, structural, compliant, and mounting features into the same artifact.

2.3.1 Fabricating Form-Fitting Skin Units

Whole-body skins must conform to the robot surface to avoid shifting, external supports, and sensor-position uncertainty. Form-fitting fabrication addresses this by using the robot's 3D model as the basis for the printed skin unit. In the GenTact pipeline, the robot model constrains the skin geometry so that units can be generated for specific links and printed as modular coverings [40]. Fig. 2.4 shows this process for one FR3 link, from generated geometry to sensing features, fabrication, and attachment.

Form-fitting fabrication preserves the relationship between the digital sensor layout and the robot body after printing. This is especially important for nondevelopable robot surfaces that cannot be flattened without distortion. The form-fitting sensor mounting work uses printed mount geometry to keep ToF sensors in fixed, known locations on the skin unit [82].

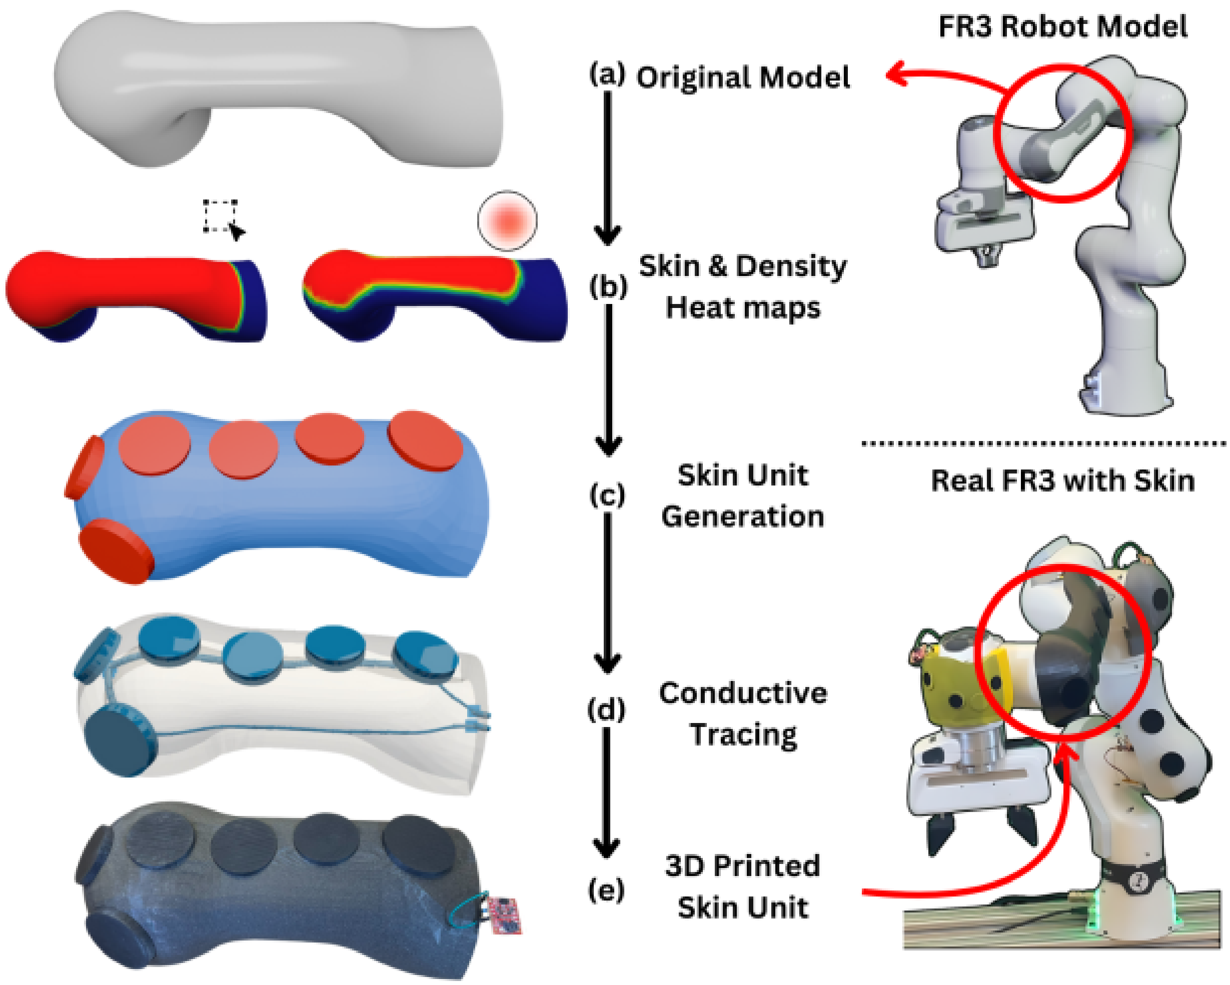

Figure 2.4: Snapshots of link 5 for the FR3 as it goes through the digital skin generation and fabrication stages of the design pipeline.

2.3.2 Printing Functional Conductive and Structural Layers

Multi-material fabrication combines the body and sensing elements in a single process. In GenTact, standard PLA provides structure, while conductive filament forms capacitive sensing nodules and internal connections [40]. The resulting artifact contains both structural and functional regions.

This approach is extended in GenTact-Prox, where conductive PLA is used to print wires and capacitive sensors inside a rigid PLA dermis [43]. The sensors are embedded within the skin rather than exposed directly on the surface. This fabrication choice changes the sensing behavior: the electrodes can be tuned for proximity effects without saturating from direct contact. Fig. 2.5 shows the operations used to produce the dermis, distribute sensors, route connection paths, and generate non-overlapping wires, which are then fabricated as a functional multi-material skin.

The multimodal soft-material work further expands the set of printable functional layers by combining conductive PLA, standard PLA, and TPU [41]. In that system, conductive PLA forms self-capacitance rings and traces, standard PLA forms the structural dermis, and TPU forms a compliant protective covering. Fig. 2.2 illustrates how self-capacitance sensors, ToF mounts, and a compliant soft covering are combined into a single skin architecture. This demonstrates how multi-material printing integrates multiple functional requirements within one skin system.

2.3.3 Maintaining Known Sensor Placement Through Fabrication

Known sensor placement is a recurring requirement across the dissertation. The perception work in Chapter 3 addresses this through calibration; here, fabrication constrains placement directly through printed geometry. This direction grew out of our earlier work on contact localization for a curved, variable-density capacitive skin, where the internal sensor positions were unknown after the skin was applied by hand and had to be recovered with a learned model [42]. That effort showed how much uncertainty comes from not knowing where each sensor ends up, and it motivated the approach here: rather than measure or learn sensor locations after the fact, we fix them in the

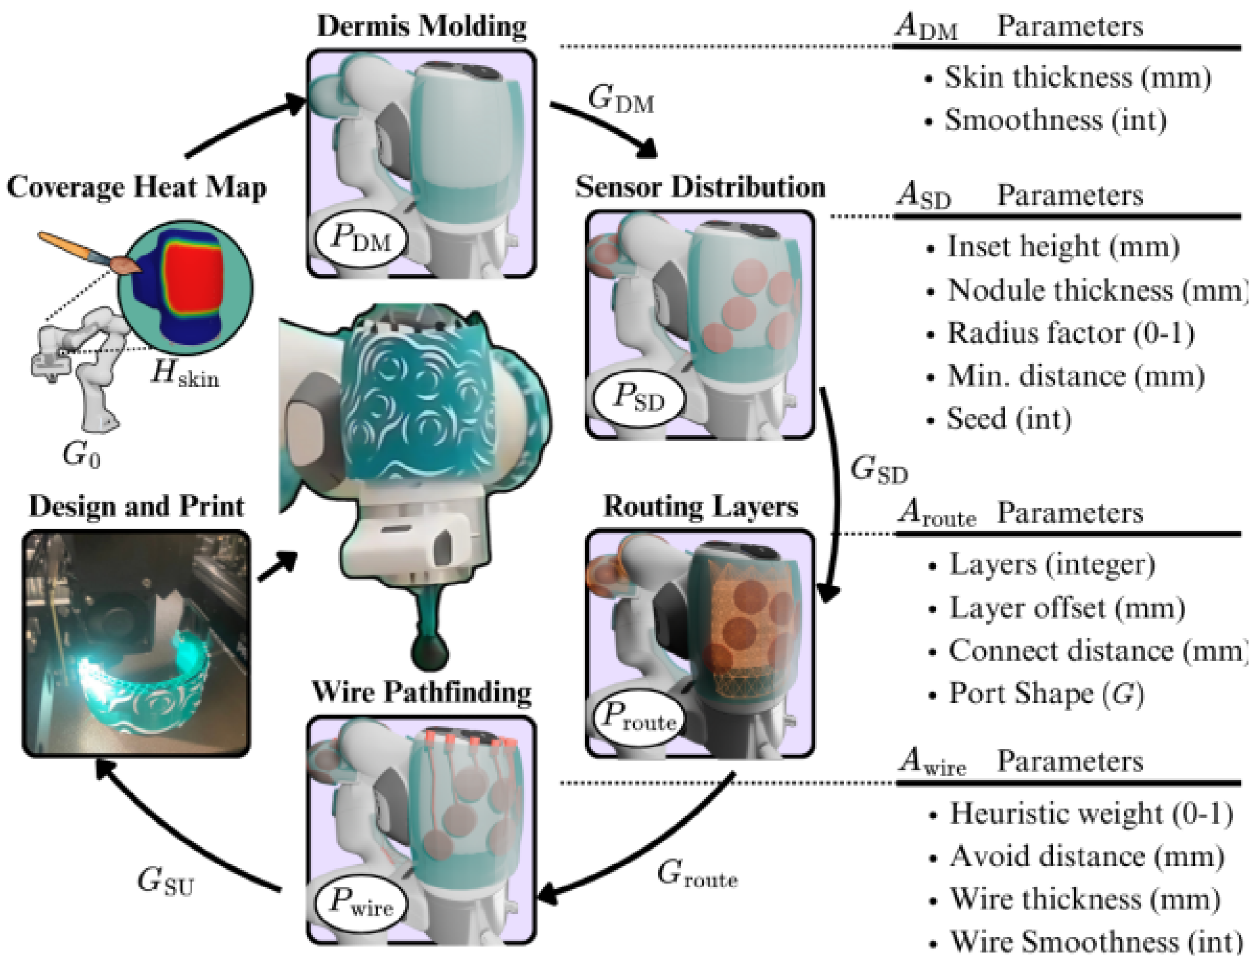

Figure 2.5: The full procedural design stage can be categorized into four main processes: $P_{\mathrm{DM}} \rightarrow P_{\mathrm{SD}} \rightarrow P_{\mathrm{route}} \rightarrow P_{\mathrm{wire}}$. The series of operations molds the dermis layer, distributes the sensors within the dermis, routes a graph of the possible space that wires should occupy, and finally solves for non-overlapping wires between each connection port and sensor. Each process depends on its respective parameters, and the swirl design on the skin serves as a visual reference for the embedded sensors.

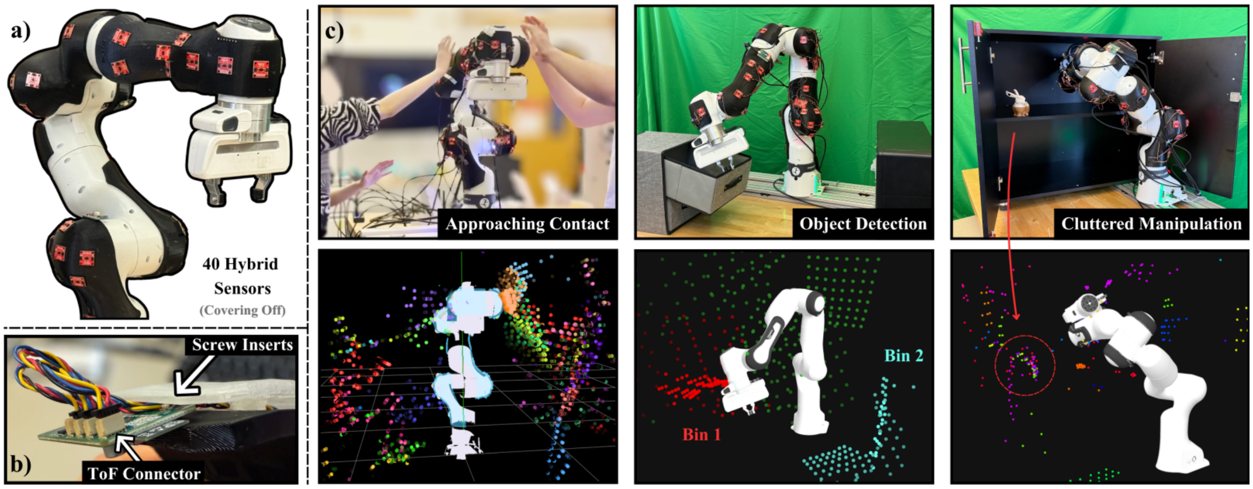

Figure 2.6: a) 40 hybrid sensing nodules distributed over the body of the FR3. b) Screw inserts were used for efficient wire management to the microcontroller. c) The depth data observed by the distributed sensors are demonstrated in various robotics scenarios.

printed geometry.

The form-fitting sensor mounting paper demonstrates this by subtracting a user-specified sensor PCB model from the skin geometry to produce snap-in mounts for ToF imagers [82]. Because the mounts are printed into the skin, the sensors attach at fixed positions and orientations. The point cloud reconstructions in Fig. 2.6 show that the skin can be removed and reapplied while preserving spatial sensing without manual pose recalibration.

GenTact-Prox similarly relies on the fact that the skin units are designed and fabricated for the robot morphology. The printed units snap onto the robot, and the kinematic chain is assumed preserved, so each sensor's pose is defined relative to the link frame [43]. From the manufacturing perspective, this is an advantage because the printed object encodes the intended sensor layout, and the robot-specific fit reduces uncertainty introduced by manual placement.

Fabrication-based placement reduces physical pose uncertainty, although sensor signals may still require calibration for sensitivity, noise, or material-specific response. Preserving sensor geometry through the printed structure makes the skin easier to install, remove, reapply, and reproduce across experiments.

2.4 Scalable Integration with Robot Systems

A robot skin must connect mechanically, electrically, and computationally to the robot system. Multi-material additive manufacturing supports this by incorporating attachment features, signal routing, electrical interfaces, and hardware constraints into the fabricated artifact.

2.4.1 Electrical Interfaces and Embedded Routing

Electrical routing is a practical barrier to scalable robot skin deployment. Whole-body sensing requires many connections across curved, moving surfaces, and exposed wires can shift, bend, snag, introduce signal drift, or vary across builds. GenTact-Prox addresses this by printing conductive paths through the skin geometry [43]. A graph-search strategy connects sensors to ports while avoiding wire overlap, as shown in Fig. 2.5. For this chapter, the key point is that these paths become manufacturable conductive structures.

The multimodal soft-material work uses threaded inserts to align printed conductive terminals with microcontroller pin-outs and manage wiring, as shown in Fig. 2.6 [41]. This reduces ad hoc external wiring and helps stabilize capacitive sensor signals.

2.4.2 Sensor Mounting Without Manual Calibration

The form-fitting ToF mounting work demonstrates how printed skin units can integrate commercial sensors while preserving known placement [82]. Instead of fabricating the sensing element itself, the printed skin provides a mechanically constrained mount for a PCB-based sensor. This allows multi-material additive manufacturing to support sensing modalities that are implemented as commercial boards, such as ToF imagers, by fabricating the robot-specific structures that hold them in known locations.

In the ToF mounting work, each sensor board is snapped into a printed mount. The mount geometry is generated from the sensor PCB model and the robot link geometry, so the sensor pose is defined by the digital model and preserved by the physical print. Fig. 2.7 shows the printed skin unit and the electrical hardware used to aggregate data from the ToF imagers. The skin unit was printed for the fifth link of the FR3, fitted with eight ToF imagers, and used to reconstruct objects in the robot's nearby environment [82].

This example reinforces the chapter's main delineation by showing how additive manufacturing can fabricate robot-specific mounts that preserve sensor placement and make integration easier. This allows robot skins to incorporate sensing modalities that are not themselves printable, while still benefiting from a repeatable digital-to-physical fabrication process.

2.4.3 Deployment on Robot Hardware

The papers discussed in this chapter all move beyond isolated printed samples by deploying skin units on robot hardware. GenTact was evaluated by fabricating six capacitive sensing skins for a Franka Research 3 robot arm in a physical human–robot interaction scenario [40]. The form-fitting ToF mounting work deployed a printed skin unit on the fifth link of the FR3 and used the mounted sensors to reconstruct nearby objects as point clouds [82]. GenTact-Prox deployed five printed tactile/proximity skin units on the FR3 to characterize perisensory space and support online object-aware operation [43]. The multimodal soft-material work deployed six artificial skin units with 40 hybrid sensing elements over the FR3 robot arm [41].

These demonstrations show that multi-material additive manufacturing can produce skins that function after installation on robot hardware. A printed sensing sample on a flat surface is useful for early validation, but whole-body robot skins must also fit around links, avoid joints, route signals, connect to electronics, and provide data during robot operation. The repeated deployment across these works supports the claim that multi-material additive manufacturing is mature enough to serve as a general fabrication method for robot skins.

2.5 Expanding Sensing Modalities Through Additive Manufacturing

Multi-material additive manufacturing supports multiple sensing modalities through different materials and printed structures, including tactile, proximity, pre-touch, hybrid, and magnetic sensing.

2.5.1 Capacitive Tactile Sensing

The GenTact Toolbox demonstrates 3D-printed capacitive tactile sensing as an initial sensing modality [40]. In this system, conductive filament forms sensing nodules that respond to contact. The skin body is printed to fit the robot geometry, while conductive regions are printed into the skin as individually addressable contact points. Fig. 2.4 shows the progression from robot model to skin generation, conductive tracing, and printed skin unit.

Capacitive tactile sensing demonstrates how printed conductive materials can turn a geometric skin into a functional sensing surface. Because the mechanical body and sensing elements are fabricated together, the conductive features do not require separate PCB fabrication or manual electrode placement.

2.5.2 Proximity and Pre-Touch Sensing

GenTact-Prox extends the printed capacitive approach from tactile sensing to proximity and pre-touch sensing [43]. The key fabrication change is that the conductive capacitive electrodes are embedded within the skin rather than exposed on the surface. This reduces direct coupling at contact and allows the sensors to detect weaker proximity effects from nearby objects. The result is a fully 3D-printed skin that senses both tactile and proximity stimuli.

This extension connects the fabrication work to the perception and control chapters by showing that printed skins can support body-centered pre-touch perception. Fig. 2.11 shows the deployed GenTact-Prox skin units and a visualization of proximal receptive fields surrounding the sensors. The skin achieved detection ranges up to 18 cm during evaluation [43], demonstrating that printed capacitive sensors can support nearby-space awareness as well as touch.

2.5.3 Hybrid Self-Capacitance and Time-of-Flight Sensing

My multimodal soft-material work demonstrates that multi-material fabrication can support hybrid sensing by combining self-capacitance and time-of-flight sensing in the same skin unit [41]. Self-capacitance provides close-range tactile and pre-touch sensitivity, while ToF sensing provides longer-range depth measurements. Their complementary ranges expand the observable space around the robot body.

Fig. 2.8 shows this hybrid skin on the FR3, from the internal self-capacitance circuitry and proximity-sensor mounts to the deployed skin reconstructing the nearby scene as a point cloud. The two modalities require different structures, with printed conductive rings and traces for self-capacitance and mounted commercial imagers for ToF sensing. The multimodal skin is created using a layered printed structure: conductive PLA for capacitive rings and wires, standard PLA for the structural dermis, and TPU for the compliant covering. Fig. 2.2 shows the three-layer hybrid sensor structure, and Fig. 2.6 shows 40 hybrid sensing nodules distributed over the FR3.

This example is one demonstration of the chapter's main claim. Multi-material additive manufacturing provides a fabrication framework for combining multiple sensing strategies, commercial sensor mounts, printed conductive features, and compliant mechanical protection in one robot-specific skin.

2.5.4 Magnetic Sensing

Magnetic sensing measures local surface deformation, including normal and shear components, which capacitive and time-of-flight sensing do not. ReSkin [7] showed that a soft magnetized interface over magnetometers can recover contact and force while keeping the electronics separate from a replaceable skin. eFlesh [69] focused on fabrication, 3D printing cut-cell microstructures whose compression response and repeated geometry hold embedded magnets in a printable volume. Both sense contact well, but neither is built to fit an arbitrary robot platform. Our magnetic skin, shown in Fig. 2.9, embeds permanent magnets in a compliant printed epidermis above magnetometers placed at digitally designed locations in the dermis, so local deformation produces a measurable change in magnetic flux.

Magnetic sensing is a useful test of how general the fabrication method is. It is normally fabricated flat and then bent around a surface, an approach that does not transfer to nondevelopable robot geometry. Producing the magnetic skin directly on the robot's conformal geometry, with magnets and magnetometers held at known locations by the printed structure, shows that the additive manufacturing workflow accommodates sensor types that were not originally designed for whole-body robot surfaces. Integrating this new modality reused the existing digital-to-physical pipeline, indicating that additional sensor types can be adopted at low engineering cost.

2.5.5 Soft and Compliant Material Layers

Rigid printed skins can provide structural support and stable sensor placement, but physical human–robot interaction also requires compliance. Hard plastic surfaces may be undesirable during contact because they can increase impact severity or damage the environment. The multimodal soft-material work addresses this by adding a flexible TPU covering over the sensing layer [41].

The compliant layer provides a softer exterior, protects internal wiring and electronics, and allows capacitive sensing to remain responsive to contact and pressure. In the evaluation, the self-capacitance sensors detected contact above the minimum SNR threshold in all tested configurations, including with the flexible covering attached [41]. The signal also increased when the covering was squeezed, indicating that the tactile response was correlated with pressure.

This result shows that material diversity affects both sensing and physical interaction. Multi-material additive manufacturing allows the skin to combine functional sensing materials with mechanical compliance.

2.6 Experimental Demonstrations

The experimental demonstrations are organized around fabrication outcomes, including deployment on robot links, preserved sensor placement, multimodal integration, and support for contact-aware behavior.

2.6.1 Fabrication and Deployment of Whole-Body Skin Units

The GenTact Toolbox demonstrates fabrication and deployment of multiple skin units on a real robot [40]. Six unique sensing units were generated and fabricated for the FR3 robot arm. The resulting skins were tested in a physical human–robot interaction scenario. Fig. 2.10 shows the fabricated skins deployed on the robot and the broader pipeline from digital model to printed skin.

This demonstration supports RQ1 because it shows that multi-material fabrication can produce multiple robot-specific skin units. It also shows that the same digital workflow can generate skins for different robot morphologies, including the Unitree H1-2 humanoid and Go1 quadruped in the GenTact paper [40]. The important idea for this chapter is that the fabrication workflow can translate different generated geometries into physical, functional skin units.

2.6.2 Large-Area Proximity Reconstruction

The form-fitting sensor mounting work demonstrates that printed skin units can support large-area proximity sensing with commercial ToF imagers [82]. A skin unit for the fifth FR3 link was fitted with eight ToF imagers and connected through a multiplexer and microcontroller to reconstruct nearby geometry as point clouds.

Fig. 2.3 shows a visualization of the real obstacles and reconstructed point clouds. The paper reports reconstructions of a person, an arm and fist, and a tape roll, including the hollow center of the tape roll [82]. These reconstructions were obtained after more than ten removals and reapplications without manually calibrating the sensor poses. This result supports the fabrication argument by showing that printed mounts can preserve spatial sensor meaning across removal and reapplication.

Table 2.1: Average contact SNR of the capacitive sensors in Fig. 2.2

	No covering	**Covering (Rest)**	**Covering (Squeeze)**
With ToF	50 ± 20	13 ± 3	22 ± 5
No ToF	120 ± 20	37 ± 4	45 ± 3

2.6.3 Multi-Modal Sensing and Physical Compliance

The multimodal soft-material work demonstrates that additive manufacturing can support multiple sensing modalities and compliant material layers in a deployed robot skin [41]. Six artificial skin units with 40 hybrid sensing elements were deployed over the FR3 robot arm. The ToF sensors reconstructed dynamic scenes as point clouds, while the self-capacitance sensors detected contact through the printed conductive rings.

Fig. 2.6 shows the distributed hybrid sensors over the FR3 and point cloud reconstructions for approaching contact, object detection, and cluttered manipulation. Table 2.1 shows that self-capacitance sensing remained above the minimum SNR threshold across configurations with and without ToF sensors or compliant covering [41]. These results indicate that added proximity sensing and compliance did not prevent capacitive contact detection.

This demonstration supports RQ1 by integrating long-range proximity, close-range capacitive response, and compliant mechanical protection within one skin.

2.6.4 Integration with Contact-Aware Robot Behavior

The final demonstration connects the fabrication work forward to the behavioral goals of the dissertation. GenTact-Prox uses fully 3D-printed tactile and proximity sensors to support contact anticipation and body-centered perisensory space mapping [43]. Five GenTact-Prox units were deployed on the FR3 to characterize the perisensory space where the skin provides information for contact anticipation. Fig. 2.11 shows the deployed skin units and an illustration of the perisensory space around the robot. The important connection is that the fabricated skin is integrated into a robot system that can reason about nearby space. This connects the fabrication foundation to the perception and control chapters that follow (Chapters 3 and 4), by showing that multi-material fabrication can produce the sensing hardware needed for the contact-aware and object-aware robot behavior developed there.

2.7 Whole-Body Deployment Across Robot Platforms

The strongest evidence that multi-material additive manufacturing is a general fabrication method is that the same workflow produces working whole-body skins for robots with very different geometries. We fabricated and deployed skins across three real platforms, the Unitree H1-2 humanoid, the Unitree Go1 quadruped, and the Franka Research 3 robot arm, and evaluated the fabricated skins along three axes: repeatability, scalability, and diversity (Fig. 2.12).

2.7.1 Repeatability

Repeatability is consistency in sensor position and sensor signal across prints. For mounted off-the-shelf sensors, repeatability concerns how accurately each sensor's printed pose matches its digitally designed pose. For fully printed conductive sensors, it concerns the fabrication quality of the conductive traces, their electrical continuity to each sensing node, and the consistency of the resulting signals. Across identically printed copies of a skin design, the ToF mounts reproduce sensor pose closely enough to reconstruct the environment without re-calibration, and the self-capacitance skins reproduce a consistent contact response across prints, with the primary sensing node producing a signal substantially stronger than adjacent-node crosstalk so that contact can be localized reliably [82, 41].

2.7.2 Scalability

Scalability is the range of surface area and geometry the process can cover. The three platforms span very different morphologies, and skin coverage ranges from single patches to full-body deployments. Because the pipeline requires only a surface mesh as input, it extends beyond these platforms to arbitrary conformal surfaces. The ToF skins were integrated into each robot's perception pipeline, providing coverage of near-body regions that each platform's built-in sensors cannot observe. In total, this amounted to 13 printed skin units carrying 64 proximity sensors on the H1-2, 7 units and 21 sensors on the Go1, and 7 units and 40 sensors on the FR3, all produced through the same pipeline.

2.7.3 Diversity

Diversity is the ability to integrate multiple sensing modalities within a single skin while preserving co-designed placement and synchronized measurement. The modalities in this chapter span both off-the-shelf sensors (ToF breakout boards) and fully printed custom sensors (conductive-PLA capacitive nodes and magnetometer-based magnetic sensing). The hybrid skin shows that two modalities can coexist in one unit with complementary coverage of the pre-contact and contact fields, without degrading either signal. Together, these results show that a researcher with access to a multi-material 3D printer can reproduce, modify, or extend these skins across sensing needs and robot platforms.

2.8 Discussion & Conclusion

This chapter presented multi-material additive manufacturing as a method for turning robot-specific digital skin geometries into physical, functional, and deployable sensing systems. In response to RQ1, the works discussed in this chapter show that multi-material additive manufacturing improves reproducibility and scalability by moving fabrication into a digital workflow: form-fitting skins preserve the relationship between the robot model and the physical artifact, printed mounts and embedded sensors preserve known placement off the print bed, and multi-material structures integrate bodies, conductive elements, routing paths, and connection interfaces. Across the demonstrated systems, skins were fabricated for the Unitree H1-2 humanoid, the Unitree Go1 quadruped, and the Franka Research 3 robot arm, deployed on hardware, removed and reapplied, and used to collect sensing data.

The same fabrication method also supports sensing diversity, through printed capacitive elements, commercial ToF mounts, TPU coverings, hybrid self-capacitance/ToF layouts, and a new magnetic skin. Together, these examples show how additive manufacturing can combine multiple sensing modalities, material properties, and functional layers within robot skins.

This chapter treats fabrication as a core technology of whole-body robot skin research. Reproducible fabrication methods make skins easier to share and extend, and multi-material additive manufacturing supports this goal by making the process more explicit and repeatable across robots and tasks. Fabrication is the foundation; the chapters that follow show how the distributed sensing on these skins is calibrated and integrated (Chapter 3), and how that perception is used for safe physical interaction (Chapter 4).

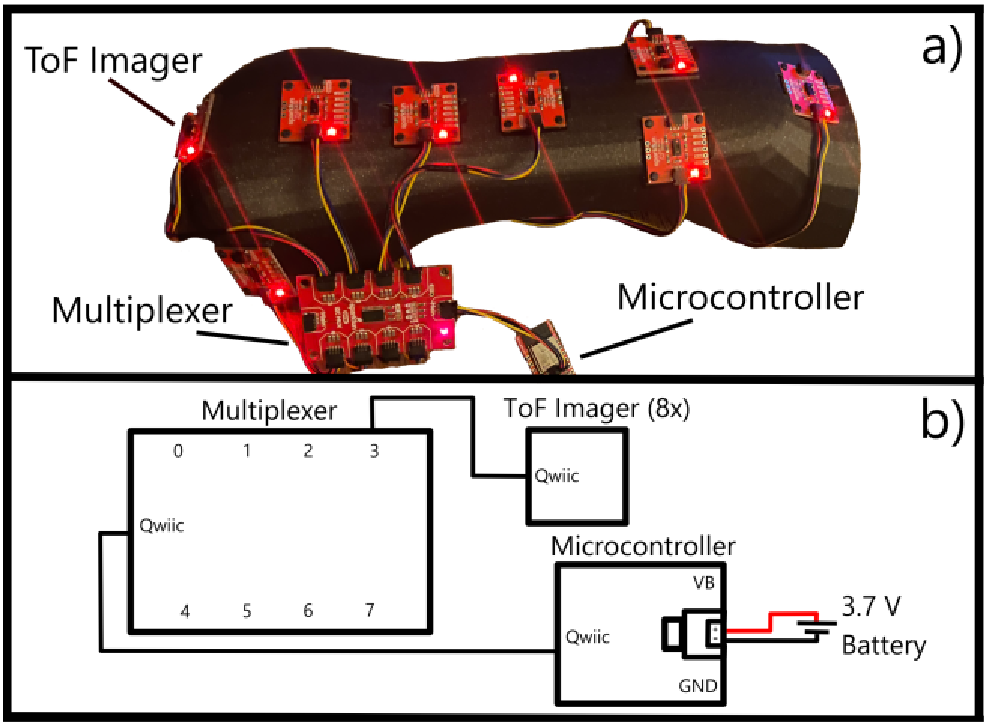

Figure 2.7: (a) Photo of the skin unit. For proximity sensing, we chose to mount time of flight (ToF) imagers. Each ToF imager is attached to an index (0-7) on the multiplexer. (b) General schematic of the electrical hardware.

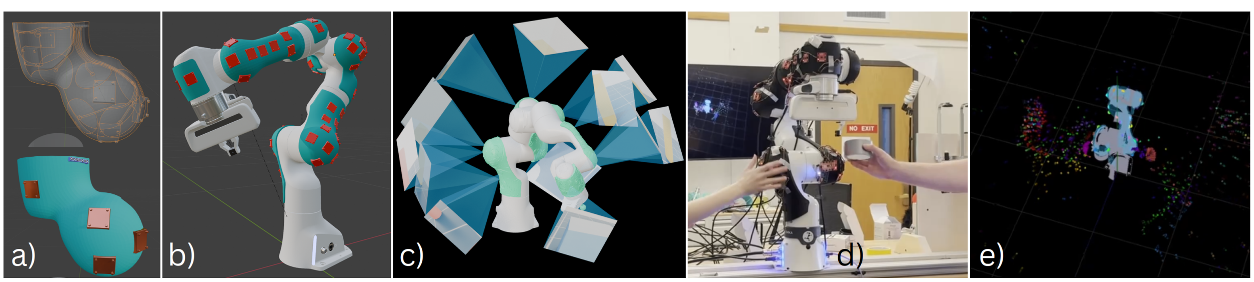

Figure 2.8: The hybrid self-capacitance and time-of-flight skin fabricated for the Franka Research 3. (a, top) The self-capacitance circuitry printed inside the skin body, together with the mounts that hold the off-the-shelf time-of-flight proximity sensors; (a, bottom) the custom screw-in PCB mounting port on the top of the skin and the proximity sensor mounts shown in bronze. (b) Digital model of form-fitting skin units with sensor mounting locations marked in red. (c) Representative fields of view of the time-of-flight sensors; each proximity sensor observes a pyramidal volume extending outward from the robot surface. (d) The printed skin units deployed on the physical FR3, with proximity sensors across six links providing real-time perception of two people moving around the arm. (e) The corresponding point-cloud reconstruction of the scene in (d), rendered in OpenGL from the onboard sensor data.

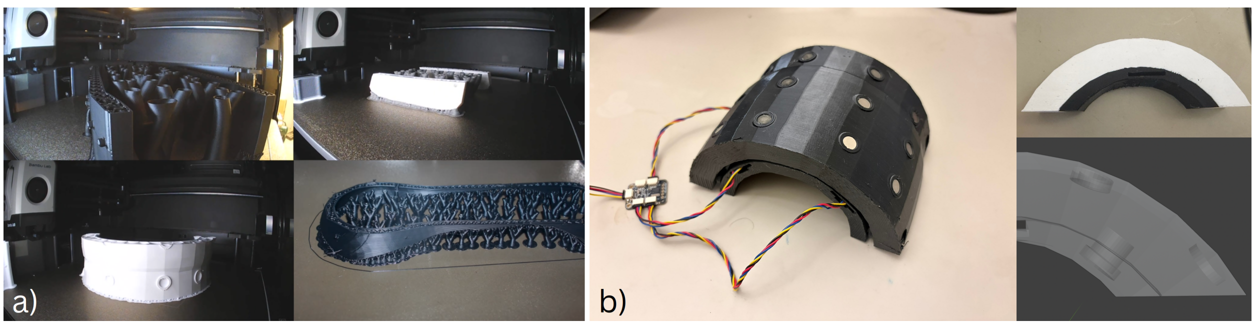

Figure 2.9: Multi-material fabrication and a printed magnetic skin. (a) Rigid (black PLA) and flexible (white TPU) skin units partway through fabrication on the multi-material printer. (b) A magnetic skin instantiation built around the same magnetometer used in eFlesh [69], printed to fit the end-effector of the Franka Research 3 and sense contact forces; (b, top right) a skin unit combining compliant and rigid components in a single multi-material print, (b, bottom right) the internal peg-in-hole structure that mechanically bonds the rigid and flexible layers.

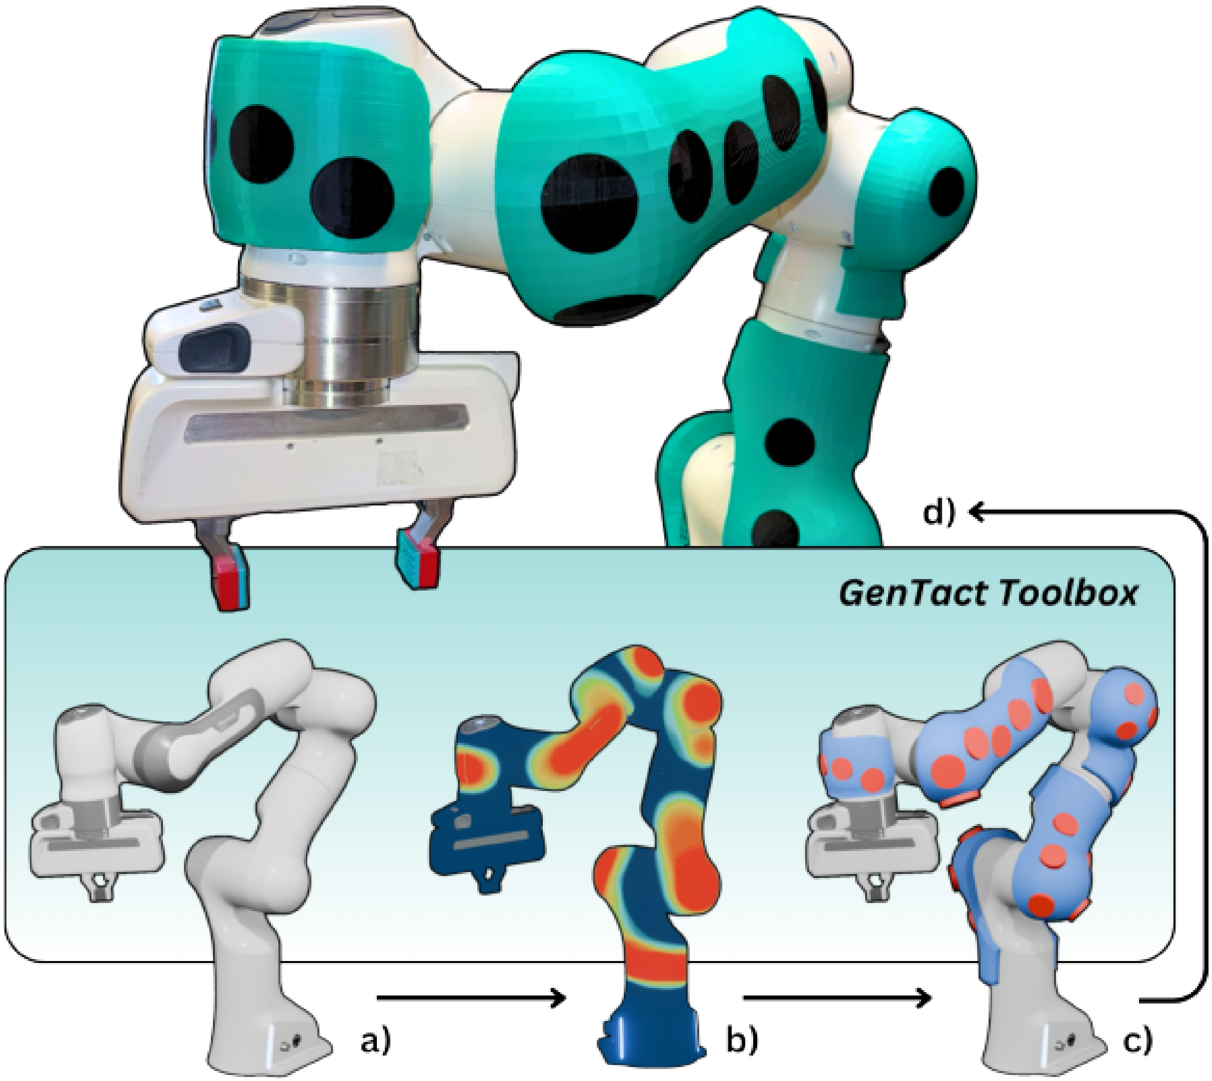

Figure 2.10: The computational pipeline presented in this work, GenTact Toolbox, generates form-fitting and adaptable whole-body tactile sensors. GenTact Toolbox uses the 3D model of a given robot (a) and a user-generated heat map (b) to create digital meshes of sensor arrays (c) that can be 3D printed as functional tactile sensors (d).

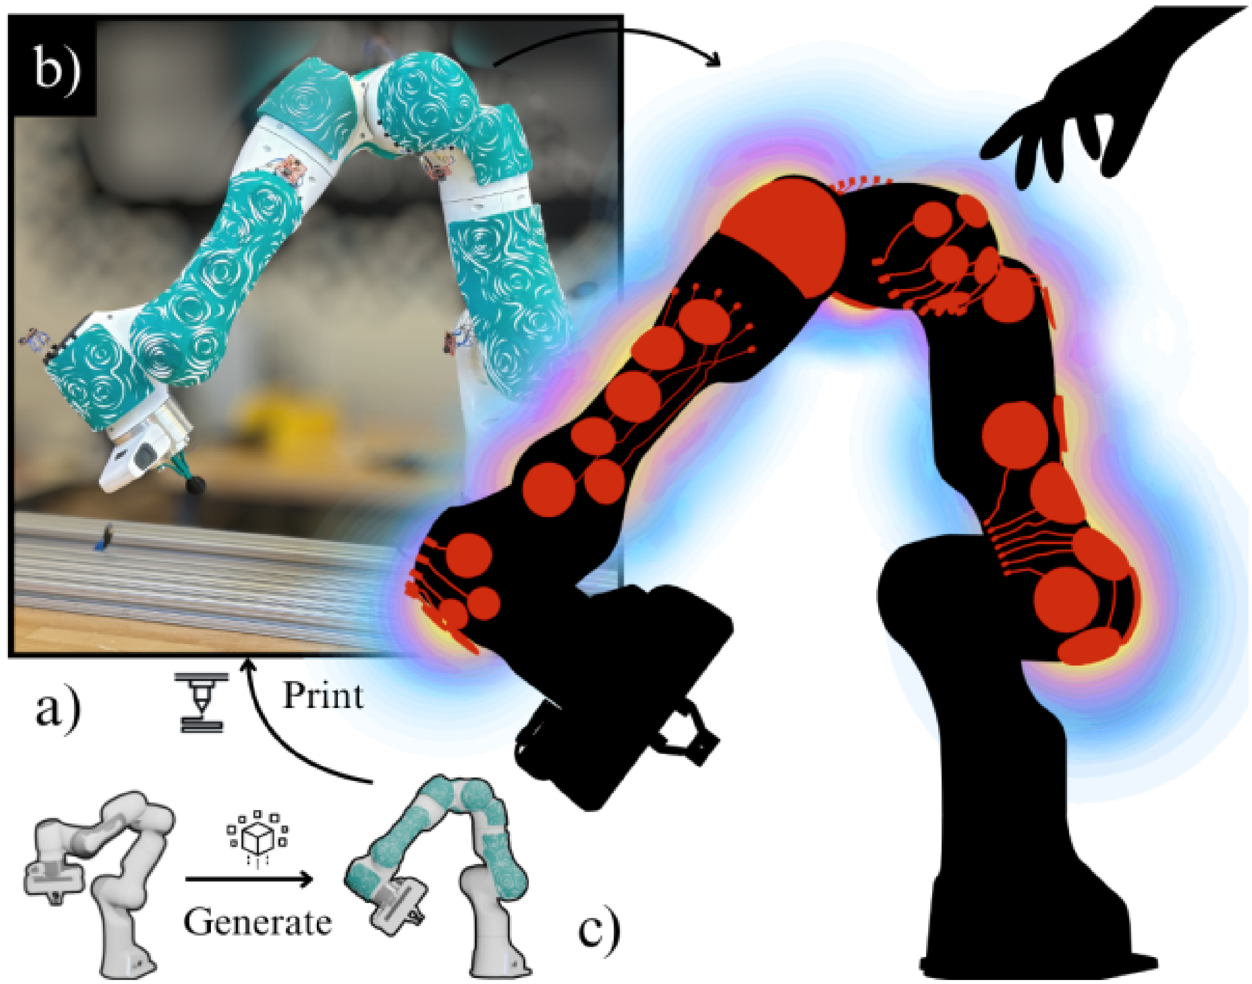

Figure 2.11: a) GenTact-Prox whole-body artificial skins are digitally generated to fit a robot's morphology and 3D-printed with functional tactile and proximity sensors. b) Five unique skin units were deployed on the FR3 robot for evaluation. c) An "x-ray" view of the individual sensors nested within the GenTact-Prox skin units surrounded by the sensor's proximal receptive field (PRF) in the nearby surrounding space.

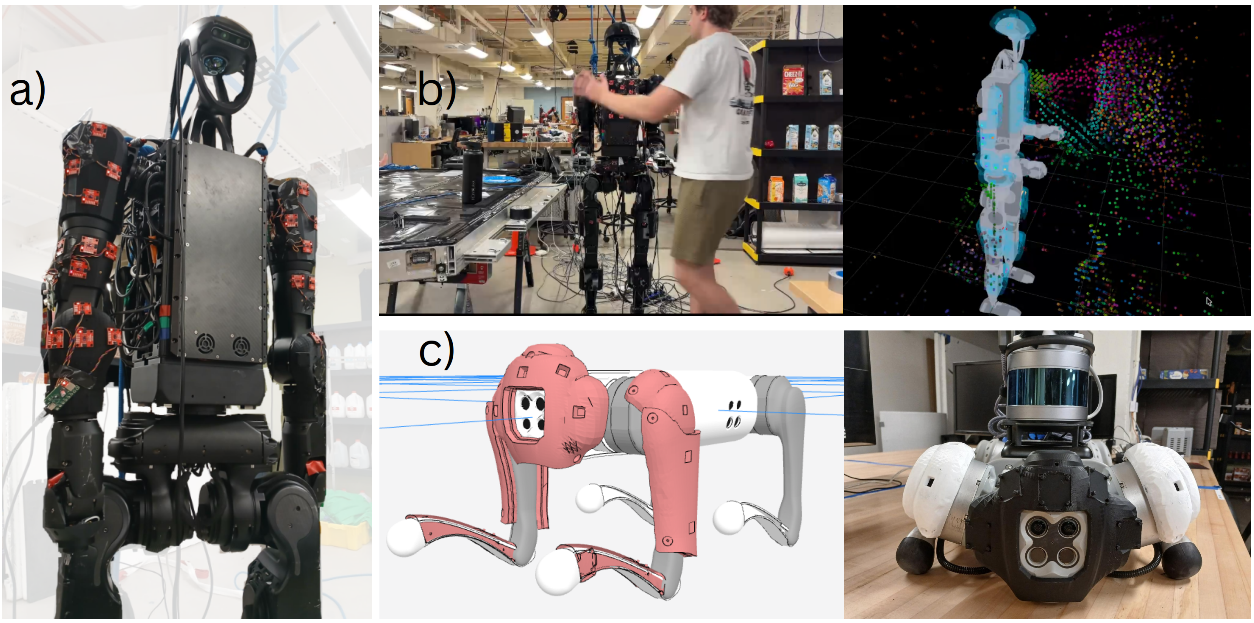

Figure 2.12: Extending the SPAD-based proximity skin across robot morphologies. The same digital-to-physical workflow produces form-fitting skin units for a humanoid and a quadruped, each unit combining a rigid printed dermis with a flexible TPU epidermis that shields the sensors and provides a clean form factor for wire management. (a) The Unitree H1-2 humanoid with its arms covered in printed proximity-sensing skin units. (b, left) A person interacting with the sensorized humanoid and (b, right) the corresponding point-cloud reconstruction from the onboard SPAD proximity sensors. (c, left) The seven skin units designed for the Unitree Go1 quadruped and (c, right) the units mounted on the quadruped prior to attaching the SPAD sensors. Across these two platforms, the H1-2 humanoid is covered by 64 sensors across 13 skin units and 16 microcontrollers, and the Go1 quadruped by 21 sensors across 7 skin units and 7 microcontrollers.

Chapter 3

Nearby-Space Perception from Distributed Onboard Sensing

3.1 Overview

A skin covered in sensors is worth little until its readings can be trusted and located on the body; that step is what turns distributed hardware into perception a controller or a learning system can build on. Chapter 2 showed how whole-body skins can be fabricated so that sensors sit at known locations on the robot body. Fabrication puts the sensors in place; it does not, by itself, make their readings usable. A distributed skin produces many local measurements, and those measurements only inform a controller when the robot knows where each reading originates on its body and how onboard sensing relates to what external sensors already observe. This chapter is about turning distributed onboard sensing into reliable perception of the space immediately around the robot.

Robots operating around people are commonly equipped with vision, depth, and force/torque sensors. Each of these is limited near the robot body: cameras are occluded by the robot's own links, external depth is blocked as objects and people move closer, and force/torque sensors report contact only after it has happened. The region within roughly 20 cm of the robot surface—where close interaction is decided—is largely unobserved. Sensing this region directly from the robot's own body is what lets a robot reason about nearby objects before contact.

This chapter draws on two contributions. A kinematic calibration method estimates the pose of each sensor unit mounted on a robot arm, so that proximity readings can be interpreted in the robot's frame of reference [91]. A volumetric fusion method combines external depth and onboard proximity sensing, showing that the two are complementary and together reduce the amount of unobserved space around the robot [83]. Calibration grounds the sensing on the body; fusion integrates it with external perception. The perception this produces is what the object-aware controllers in Chapter 4 act on.

3.1.1 Research Question

> **RQ2:** How can distributed onboard sensing be calibrated and integrated to support reliable nearby-space perception?

RQ2 has two parts. The first is spatial grounding: onboard proximity sensing must be located on the robot body before a distance reading can be turned into an estimate of where an object is in the robot's workspace. The second is integration: onboard sensing should complement external perception rather than duplicate it, reducing uncertainty in the occluded regions near the robot that external sensors cannot see. The calibration work addresses the first part [91], and the volumetric fusion work addresses the second [83].

3.2 Sensing Requirements for Nearby-Space Perception

When deploying robot skin for nearby-space perception, the robot must know where each sensor reading originates on its body so that local distance measurements can be converted to the robot base frame. Current human–robot interaction is facilitated by vision, depth, and force/torque sensors, which are limited in perception near the robot body due to occlusions, viewpoint, and latency. Onboard sensors mounted directly on the robot body provide an additional layer of perception by observing the space surrounding the robot itself. Two requirements make this sensing usable: each sensor must be spatially grounded on the robot body, and onboard sensing must be integrated with external perception to reduce unobserved space.

3.2.1 Calibrating Distributed Onboard Sensors

Distributed proximity sensing requires more than placing sensors on the robot. A proximity sensor reports a distance measurement along its sensing direction, but that measurement is only useful for a robot controller if the system knows where the sensor is located and how it is oriented relative to the robot. Without this information, the robot cannot transform a local distance reading into an estimate of where an object is in the robot's workspace.

Our calibration work [91] addresses this by treating the pose of each sensor unit as a kinematic calibration problem. Each sensor unit contains an inertial measurement unit and a time-of-flight proximity sensor. The goal is to estimate the six-dimensional pose of each unit mounted on the robot body so that proximity readings can be interpreted in the robot's frame of reference [91]. As shown in Fig. 3.1, the calibration procedure collects accelerometer data from sensor units in static and dynamic robot poses, defines a kinematic model from the robot joint to each sensor unit, and optimizes the sensor pose parameters by minimizing the difference between measured and predicted accelerations.

If sensor unit k reports a distance measurement $d_{obs,k}$ and its pose is known with respect to the robot base, then the detected object can be estimated in the robot's workspace. The visual example in Fig. 3.2 shows how calibrated skin-mounted sensor units support close-proximity collision avoidance. As a human hand approaches the robot, proximity measurements from the robot arm are interpreted relative to the robot's base, allowing the controller to generate a collision-free motion [91].

Calibration and fabrication address the same requirement—known sensor placement—from two directions. The printed skins in Chapter 2 preserve sensor geometry through the print, reducing physical pose uncertainty. Calibration recovers sensor pose for units that are mounted rather than printed in place, and can further correct residual uncertainty in either case. Together they mean a robot can trust where its sensing data originates on its body.

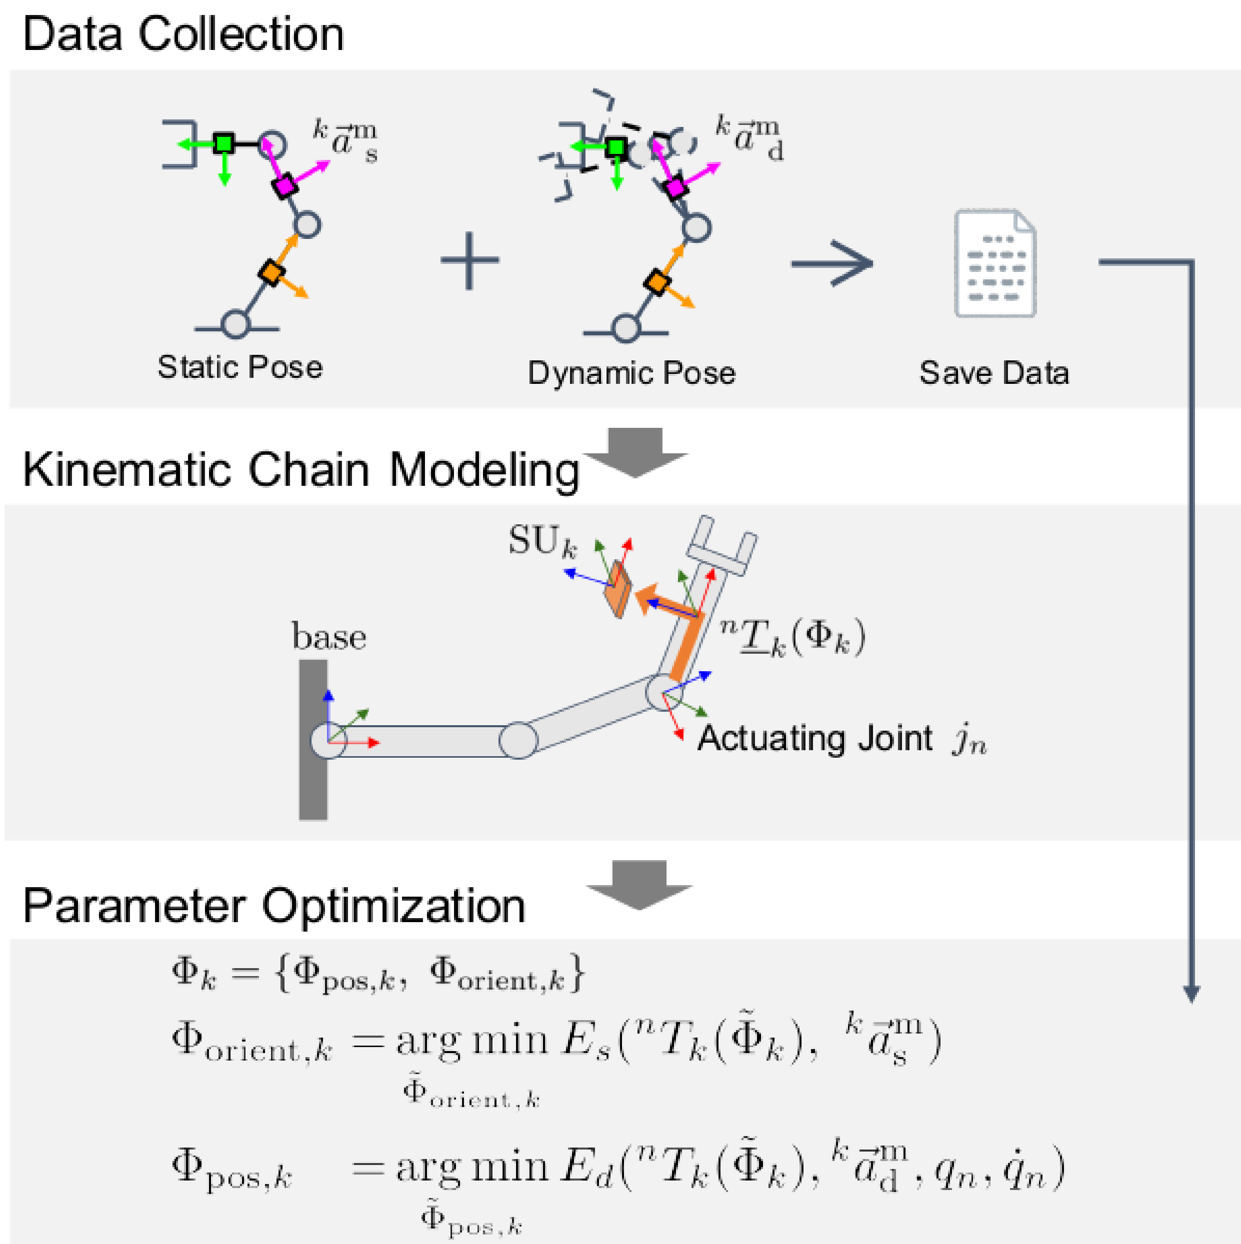

Figure 3.1: Overview of the kinematic calibration algorithm: 1) collect IMU data $\tilde{\mathbf{a}}_s^k$ and $\tilde{\mathbf{a}}_d^k$ at static and dynamic poses for all sensor units SU_k, $\forall k \in \{1, \ldots, K\}$; 2) define a kinematic model from joint j_n to SU_k, expressed as ${}^{n}\mathbf{T}_k$ and parameterized by $\boldsymbol{\Phi}_k$; 3) optimize the parameters $\boldsymbol{\Phi}_k$ by minimizing the static and dynamic error functions E_s and E_d for each SU_k. The variables q_n and $\dot{q}_n$ represent the joint angle and joint velocity of j_n, respectively.

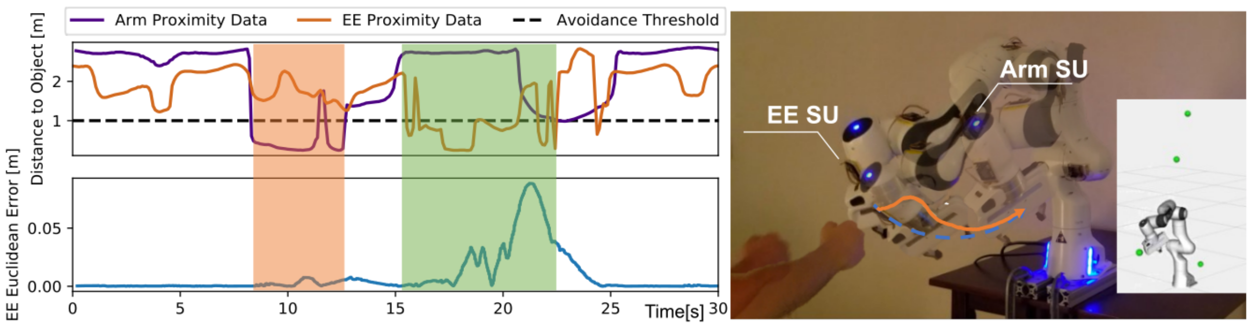

Figure 3.2: Trajectory redirection for obstacle avoidance. Right: a human approaches a robot with calibrated skin units; their accurate pose estimation allows the robot to precisely perceive the person from a distance. Bottom-Right: real-time visualization of the robot and obstacles (shown as green spheres) as sensed by the proximity sensors placed on the calibrated SUs; only below-threshold signals (arm and end-effector SU) affect the robot's motion. Left: combined graph of proximity data over time vs. the robot's end effector error along a circular trajectory. The two shaded areas indicate intervals when the human approaches a skin unit (below the avoidance threshold), causing only a temporary perturbation in the error. The perturbation stays low in the orange colored area (arm proximity activation) compared to the green area (end effector proximity activation) because the robot leverages the redundancy afforded by its kinematic chain to avoid colliding with the person while concurrently satisfying its main task objective.

3.2.2 Reducing Occluded-Space Uncertainty

A second motivation for onboard proximity sensing is the limitation of external perception. Depth cameras are commonly used to sense the robot's workspace, but they are susceptible to occlusion. When an object is placed between the camera and the robot, or when the robot body blocks the camera view, the resulting map may contain incomplete or missing information [63]. This is problematic in close-proximity interaction because occlusions become more likely as people and objects move closer to the robot. My volumetric fusion work addresses this by combining external depth with onboard proximity sensing to produce a more complete 3D representation of the robot's environment [83]. The method extends an Octomap-style occupancy mapping framework by incorporating external depth camera measurements and onboard proximity measurements into a single voxel representation. Each sensor contributes to the occupancy estimate according to a log-odds update based on its optimal sensing range. As shown in Fig. 3.3, the depth camera provides a broad external view and onboard proximity sensors provide measurements of nearby or otherwise occluded regions.

The fusion result reframes the robot's own body as an active perceptual surface. External depth answers where objects are across the wider workspace; onboard proximity answers what is happening in the near-body region the camera cannot see. Combining the two produces a representation more complete than either alone, and it is this combined, body-grounded perception that the controllers in Chapter 4 rely on.

3.3 Experimental Evaluation

The evaluation validates the two requirements from this chapter: that distributed sensors can be located on the robot body, and that onboard proximity sensing complements external depth in occluded nearby space [91, 83].

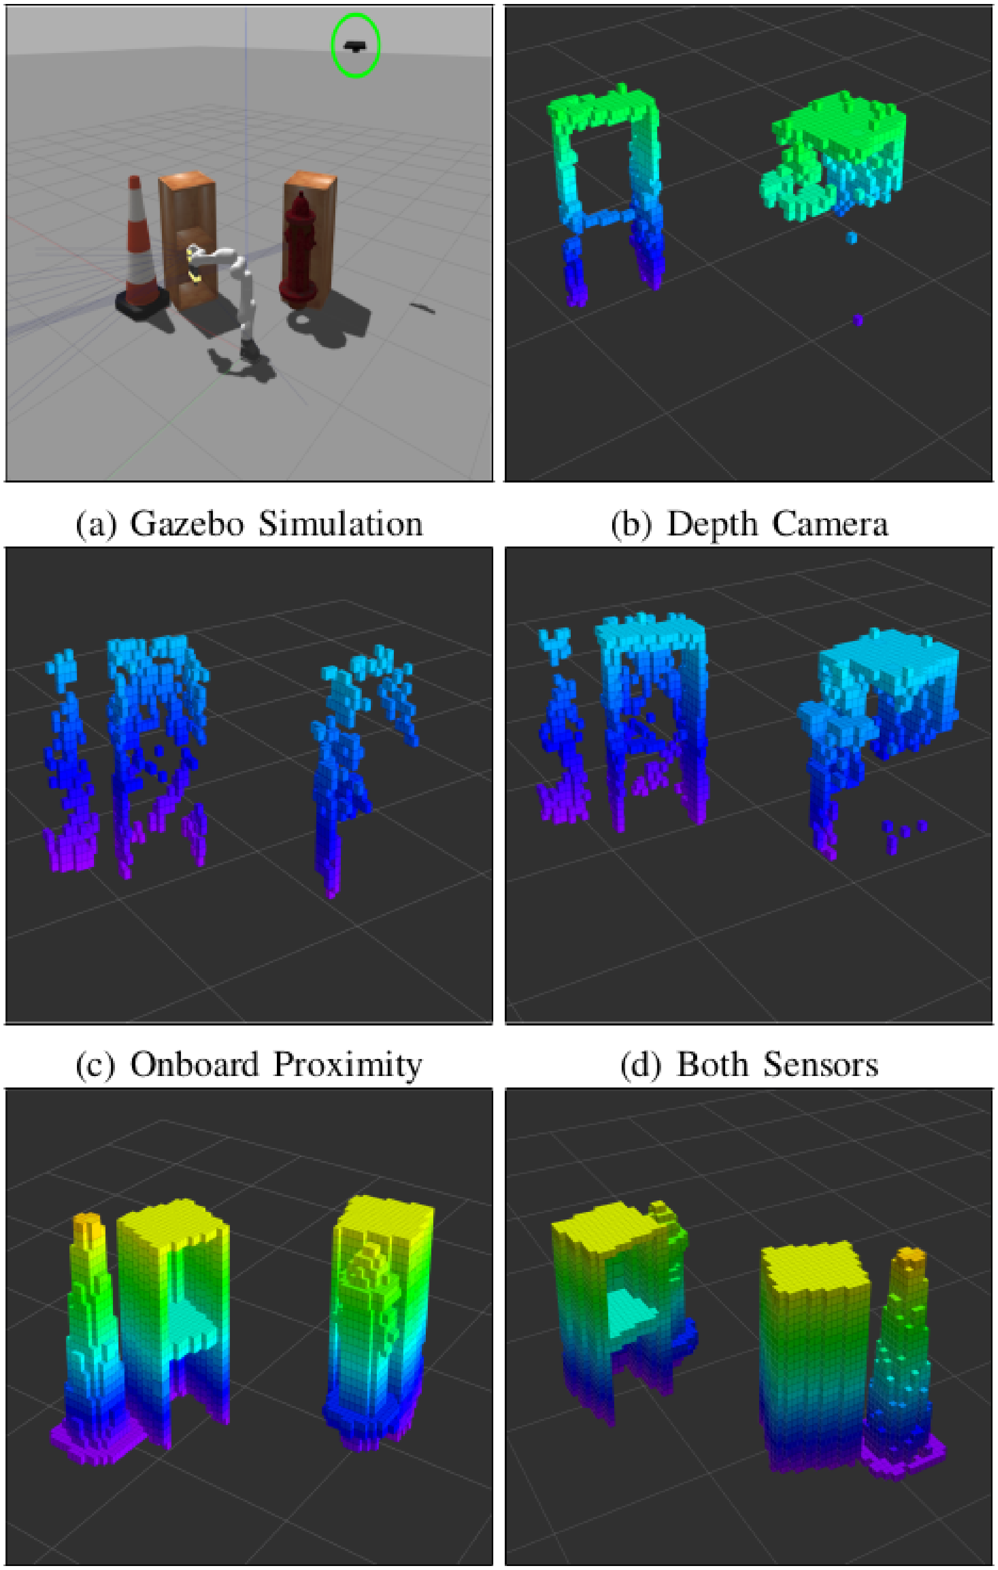

(a) Gazebo Simulation
(b) Depth Camera
(c) Onboard Proximity
(d) Both Sensors
(e) Ground Truth Front
(f) Ground Truth Back

Figure 3.3: Image (a) shows our simulated Gazebo environment with multiple objects occluded from the view of a depth camera, circled in green. Images (b), (c), and (d) show an Octomap representation when using specific sensor data. The full scene used to evaluate our method is shown in (e) and (f).

3.3.1 Validation of Nearby-Space Sensing

The kinematic calibration work validates the ability to locate sensor units on the robot body. In simulation, the method achieved sub-centimeter positional error, and in the real-world evaluation it reduced average position error compared to the baseline method [91]. The estimated sensor poses were then used in a control example where proximity readings caused the robot to redirect its trajectory around a nearby human. As shown in Fig. 3.2, calibrated sensor readings enabled real-time obstacle detection and motion adjustment, supporting their use as input for nearby-space perception.

The volumetric fusion work [83] confirms the complementary role of onboard proximity sensors and external depth cameras. The simulated evaluation used a Franka Emika Panda robot, an external depth camera, and multiple onboard proximity sensors. As shown in Fig. 3.3, depth-only sensing was limited by occlusions, while proximity-only sensing captured local information near the robot. The combined map included more occupied and free space than either source alone, showing that close-proximity interaction benefits from both external and body-mounted sensing [83].

3.4 Discussion & Conclusion

This chapter addressed how distributed onboard sensing becomes reliable nearby-space perception. In response to RQ2, two requirements were established and validated. The calibration work shows that distributed sensor readings must be grounded in the robot body before they can inform control, recovering sensor pose accurately enough to redirect motion around a nearby person [91]. The volumetric fusion work shows that onboard sensors provide information external depth cameras miss, particularly in occluded nearby regions, and that fusing the two yields a more complete representation of the robot's surroundings [83].

Together with Chapter 2, this chapter closes the gap between hardware and use: fabrication places sensors at known locations on the body, and calibration and fusion turn their readings into perception the robot can act on. That perception is the input to the object-aware controllers developed in Chapter 4, where nearby-space sensing is used to anticipate, avoid, and admit contact during physical interaction.

Chapter 4

Object-Aware Control for Safe Physical Interaction

4.1 Overview

Sensing across the body matters only if it changes how a robot acts around people; this chapter is where whole-body perception has to earn its value, by making physical interaction safer rather than only better observed. Chapter 2 fabricated whole-body skins with known sensor placement, and Chapter 3 turned their distributed readings into reliable perception of the space immediately around the robot. This chapter is about what that perception is *for*: using nearby-space awareness to control how a robot moves during close physical interaction.

Safe robot operation in human environments requires more than detecting contact after it occurs. Methods for safe human-robot interaction span low-level control, motion planning, prediction of human motion, and consideration of human factors [50]. In dynamic, shared spaces, robots must perceive nearby objects and adjust their motion in real time. This chapter treats that adjustment as three related behaviors: anticipating likely contact and softening motion before it happens, avoiding obstacles in a way that is smooth and kinematically feasible, and admitting contact deliberately when it is useful for the task.

The chapter draws on three contributions. First, a contact anticipation framework uses onboard proximity sensing to detect nearby objects, modify robot motion before contact, dynamically adjust contact thresholds, and react after contact occurs [23]. Second, a systematic evaluation framework analyzes object-aware controllers (OACs) beyond collision-avoidance success, examining the kinematic feasibility, smoothness, and constraint behavior of the motion they produce [22].

Third, a contact-admitting motion planner shows that contact can be selectively permitted rather than always avoided [64]. Together, these move a robot along a continuum from avoiding contact, to anticipating it, to admitting it on purpose.

4.1.1 Research Question

> **RQ3:** How can this perception be utilized by object-aware controllers to enable safe, kinematically feasible contact anticipation and avoidance?

RQ3 has two sides. The first is how object-aware controllers turn nearby-object information into motion—through repulsive vectors, velocity scaling, movement restrictions, control-point selection, and optimization constraints—and how they anticipate, detect, and react to contact. The second is how that motion should be judged: avoiding a collision is not sufficient if the controller produces high jerk, discontinuous control-point switching, poor manipulability, or constraints the robot cannot physically satisfy. The contact anticipation framework [23] and contact-admitting planner [64] address the first side; the evaluation framework [22] addresses the second.

4.2 From Perception to Reactive Control

Object-aware control assumes that some perception system can provide information about nearby obstacles. In this dissertation, that information comes from the fabricated skins of Chapter 2 and the calibrated, fused nearby-space perception of Chapter 3: distributed onboard proximity readings, grounded on the robot body and complemented by external depth. The controllers in this chapter abstract over the specific modality and operate on the resulting set of nearby-object detections.

Body-centered sensing for reactive control is not limited to the modalities used here. AmbiSense, for example, uses an acoustic field around a robot link to estimate proximity and bearing and feeds those estimates into a collision-avoidance inverse-kinematic solver [75], another demonstration that on-body sensing can produce actionable input for reactive control. The question this chapter takes up begins once such detections are available: given the same obstacle information, different controllers produce very different motion.

4.3 Contact Anticipation Framework

The first contribution is a contact anticipation framework that uses onboard proximity sensors to unite obstacle avoidance and deliberate physical interaction [23]. Rather than treating contact as a failure condition detected only after impact, the robot uses nearby-space perception to anticipate likely contact and change its behavior before interaction occurs. As shown in Fig. 4.2, the system moves through object detection, prior-to-contact behavior, contact detection, and post-contact behavior.

4.3.1 System Overview

The system is implemented on a Franka Emika Panda robot arm equipped with custom sensor units. Each sensor unit contains a VL53L1X time-of-flight sensor for distance measurement and an IMU for calibration. The distance sensor reports proximity data at 50 Hz and can measure up to four meters from the sensor [23]. The sensor units are distributed along the robot body so that they can detect objects in the robot's immediate surroundings, using a custom proximity PCB for nearby-space perception.

The first step in the framework is object position estimation. Each proximity sensor is oriented so that its distance measurement is aligned with the sensor unit's local z-axis. Given the sensor pose and a proximity reading, the system estimates the position of the detected object in the robot base frame. For sensor unit k, the detected object position $h_k \in \mathbb{R}^3$ is computed as:

$$h_k = {}^{O}\mathbf{r}_{SU_k} + {}^{O}\mathbf{R}_{SU_k} \begin{bmatrix} 0 \\ 0 \\ d_{obs,k} \end{bmatrix}, \tag{4.1}$$

where ${}^{O}\mathbf{r}_{SU_k}$ is the position of sensor unit k with respect to the robot base frame, ${}^{O}\mathbf{R}_{SU_k}$ is the rotation matrix of the sensor unit with respect to the robot base frame, and $d_{obs,k}$ is the

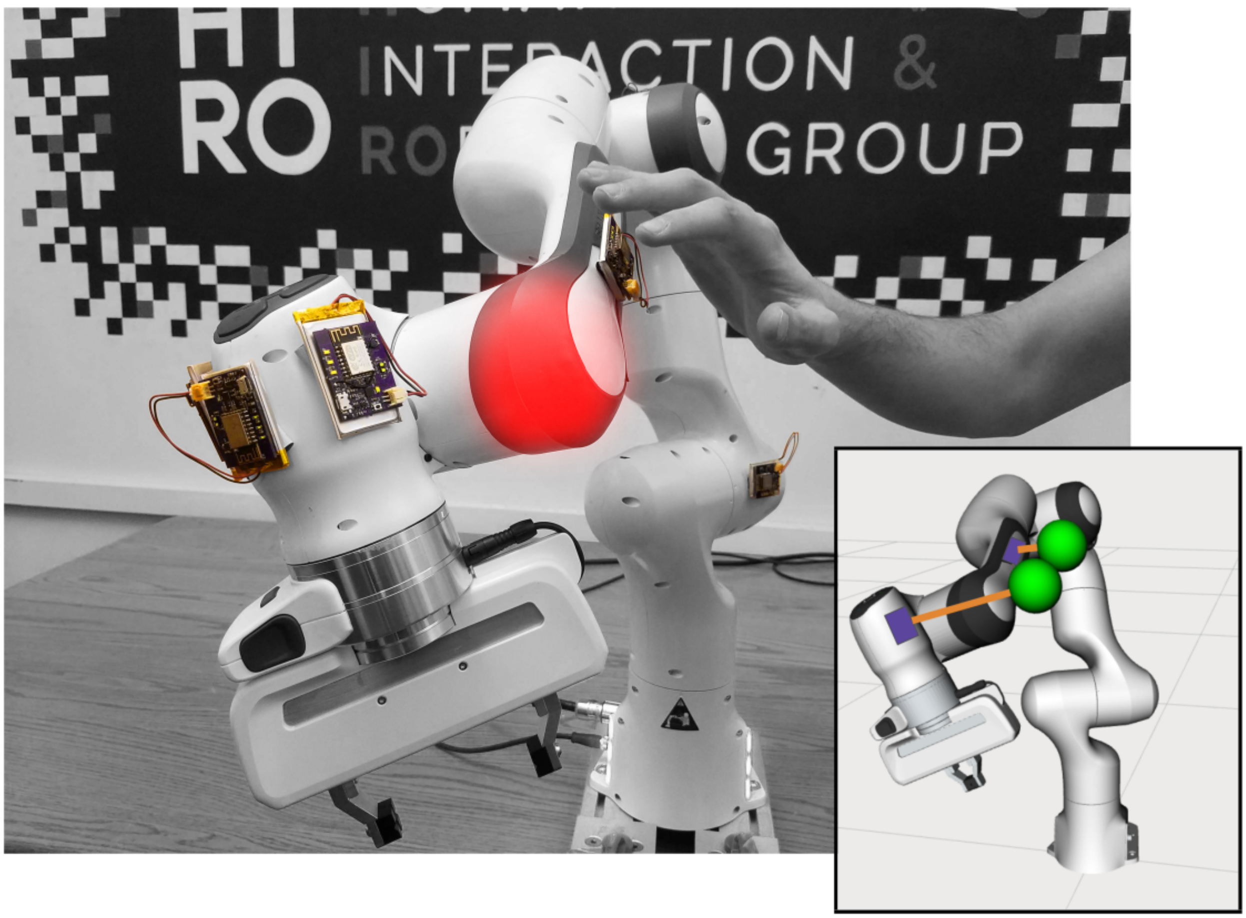

Figure 4.1: Four sensor units (shown in color) detecting a human hand near the robot's surface. The bottom right image shows a visualization of sensed objects (green spheres) based on values from proximity sensors. The red portion of the robot shows where contact is expected to be made (an interpolation of the data from the proximity sensors).

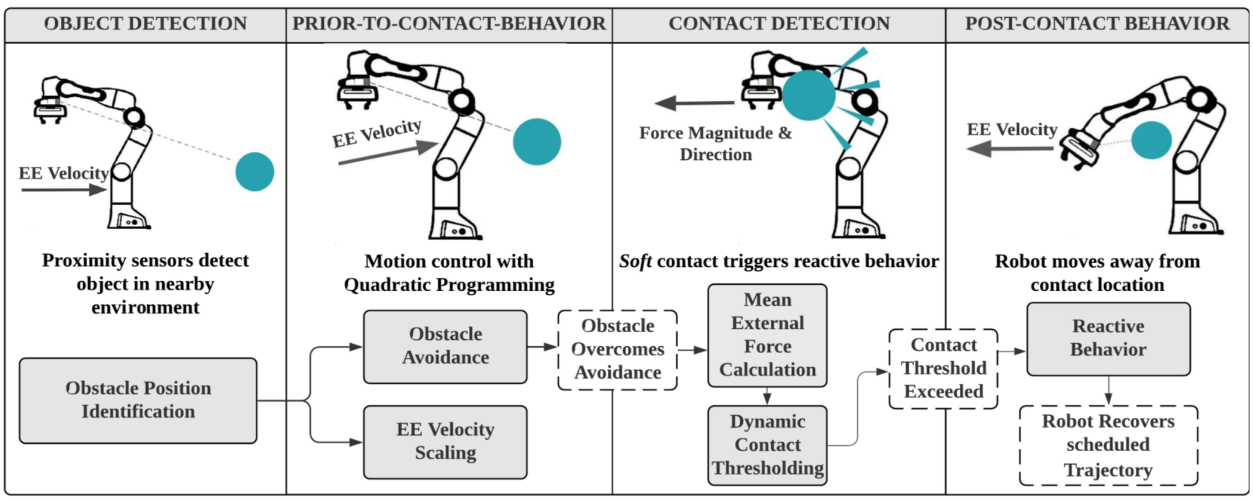

Figure 4.2: Diagram of framework for contact anticipation, detection, and reaction. Each module represents a distinct interaction from when an object enters the robot's environment to when the repulsive velocity applied to the robot arm becomes zero, and the initial trajectory is resumed.

proximity reading. This estimated object position becomes the input to the robot's avoidance behavior, velocity scaling, and contact thresholding.

The robot's main task and avoidance behaviors are implemented through a Cartesian velocity controller formulated as a quadratic programming (QP) problem. The controller solves for joint velocities while balancing task execution, singularity avoidance, and a preference for joint configurations near the middle of the robot's joint limits. This allows the robot to continue pursuing its original task while modifying its motion in response to nearby objects. This QP structure is shared with the controllers analyzed in Section 4.4, where it is stated explicitly in Eq. 4.4.

4.3.2 Prior-to-Contact Motion Modification

Prior-to-contact behavior is the first point where perception modifies robot action. The framework reduces the robot's end-effector velocity and constrains motion at selected body points based on the distance to detected obstacles.

The first mechanism is end-effector velocity reduction. As an obstacle approaches, the robot slows its task execution based on the closest measured distance. If d_{lowest} is the smallest distance between a detected object and the robot end-effector, and d_{max} is a user-defined maximum sensing distance, then the desired end-effector velocity is scaled as:

$$\dot{x}_d = \frac{\|d_{lowest}\|}{d_{max}} \dot{x}_d. \tag{4.2}$$

This scaling causes the robot to move more slowly as an object approaches. Slower motion reduces the severity of potential contact and makes interaction more natural. This behavior prepares the robot for contact without necessarily preventing contact altogether.

The second mechanism is QP-based movement restriction. The controller identifies points along the robot body that are closest to nearby obstacles and constrains their approach velocity. This use of body-level control points allows the robot to alter motion locally through its kinematic redundancy while continuing to perform the task.

Together, velocity scaling and movement restrictions create a controlled prior-to-contact behavior. The robot responds to nearby objects by slowing down, reducing motor noise, and allowing limited deviation from the original trajectory. The goal is to produce robot behavior that remains safe and controlled as it moves from nearby-object avoidance toward possible physical interaction.

4.3.3 Dynamic Contact Thresholding

Dynamic contact thresholding is the main contribution of the contact anticipation framework. The purpose of this component is to detect when contact has occurred, in which direction, and with what magnitude. A fixed threshold on external force is not sufficient for this task because the estimated external force signal is noisy and depends on robot motion. When the robot moves quickly, external force estimates may fluctuate even without contact. When the robot slows down near an object, the signal becomes less noisy and smaller forces can be interpreted more reliably.

The framework addresses this issue by computing contact thresholds dynamically from both internal force statistics and proximity readings. A sliding window is used to calculate a running average and standard deviation of the estimated external Cartesian force. The running mean provides a local baseline for the force signal, while the standard deviation captures the current level of noise. The threshold is then expanded when the signal is noisy and tightened when the signal is more stable.

For each axis, upper and lower contact thresholds are computed around the running mean:

$$\text{Contact Threshold} = \begin{cases} T_u = \mu + F_b + F_\sigma - F_{obs}^-, \\ T_l = \mu - F_b - F_\sigma + F_{obs}^+, \end{cases} \tag{4.3}$$

where μ is the running mean of the external force window, F_b is the base additional force required to trigger contact behavior, F_σ increases the threshold according to force-signal variability, and F_{obs}^- and F_{obs}^+ reduce the upper and lower thresholds based on nearby obstacle readings.

The takeaway is that proximity sensing changes the robot's interpretation of force. As objects approach, dynamic contact thresholding increases sensitivity to physical interaction by linking nearby-space perception to contact detection. Proximity sensing therefore extends beyond obstacle avoidance by modifying the robot's response to external force [23].

This mechanism is graphed in a contact interaction sequence shown in Fig. 4.3. Before contact, a human enters the robot workspace and is sensed by onboard sensor units. During contact, the external force is evaluated against the dynamically computed threshold. After contact is detected, the robot moves away from the contact direction. This sequence captures the main conceptual contribution of the framework by treating contact as a process that can be anticipated, detected, and acted upon.

4.3.4 Post-Contact Reaction

Once contact is detected, the robot enters a reactive behavior. The direction and magnitude of the estimated external force determine how the robot moves away from the interaction. This response creates space near the affected area, reduces continued force, and allows the robot to recover.

The post-contact response is connected to the prior-to-contact motion modification. The robot first slows and adjusts its trajectory as an object approaches. If physical interaction occurs, the dynamic threshold identifies the event and the robot moves away from the applied force. Once the reactive motion is complete and the repulsive velocity becomes zero, the robot resumes its original trajectory. The result is a continuous interaction loop in which the system perceives nearby objects, softens motion, detects contact, reacts, and returns to task execution.

This behavior is relevant to physical human–robot interaction because it avoids treating all contact as an emergency stop. Instead, the robot responds in a way that is proportional to the interaction and compatible with continued task execution.

4.4 Object-Aware Controller Evaluation Framework

The contact anticipation framework is one object-aware controller among many. The second contribution of this chapter is a systematic framework for analyzing object-aware controllers in general, defined here as methods that modify robot motion online to avoid or anticipate possible collisions with nearby obstacles [22]. These controllers are useful when dynamic obstacles are absent from the original motion plan, due to perception errors, occlusions, dynamic motion, or clutter.

The analysis moves beyond collision-avoidance success by examining the quality and feasibility of the resulting motion. As illustrated in Fig. 4.4, the framework is organized around three design considerations: kinematic feasibility, motion smoothness, and reactive constraint behavior. Effective object-aware control requires obstacle avoidance, kinematic feasibility, smooth motion, and meaningful constraints that the robot can realistically satisfy.

4.4.1 Controller Selection and Assumptions

The evaluation compares three representative OACs: Flacco, Ding, and Escobedo. They share similar optimization-based formulations and avoidance goals, while differing in how they impose avoidance behavior. The evaluation assumes that a perception system can reduce relevant sensor data to a finite set of rigid obstacles O. Each obstacle $o \in O$ contains the information required by the controller, such as position relative to the robot. This abstraction allows the framework to compare control behavior without depending on a specific sensing modality.

Each evaluated controller is expressed using a quadratic programming formulation:

$$\min_{\dot{q}} \frac{1}{2}\dot{q}^\top H \dot{q} + f^\top \dot{q} \quad \text{s.t.} \begin{cases} A\dot{q} \leq b, \\ b_l \leq \dot{q} \leq b_u. \end{cases} \tag{4.4}$$

Here, $\dot{q}$ represents the robot joint velocities, H and f define the quadratic and linear objective terms, and $A\dot{q} \leq b$ represents linear inequality constraints. All controllers share a Cartesian velocity control objective that minimizes the error between the desired end-effector velocity and the actual velocity induced by the robot Jacobian:

$$g(\dot{q}) = \frac{1}{2}(\dot{x}_d - J\dot{q})^\top (\dot{x}_d - J\dot{q}). \tag{4.5}$$

The controllers differ in how they modify this objective or impose avoidance constraints. The Flacco controller uses repulsive vectors and statically defined control points, as shown in Fig. 4.5. The Ding controller imposes constraints on dynamically selected control points along the robot body and includes a manipulability-related damping term, as shown in Fig. 4.6. The Escobedo controller also uses dynamically selected control points, but adds velocity scaling and a middle-joint preference term to reduce undesirable joint configurations, as shown in Fig. 4.7.

4.4.2 Kinematic Feasibility

Kinematic feasibility captures whether a controller commands behavior that is compatible with the robot's current configuration. Because a manipulator is a kinematic chain, points along the body vary in their ability to move in different directions. This matters for object-aware control because avoidance constraints may be applied to locations with limited mobility. Kinematic feasibility is evaluated using manipulability, which describes how freely the robot can move from a given joint configuration. Although this measure is often applied to the end-effector, the framework extends it to control points along the body to determine whether avoidance constraints are being placed on regions that can move effectively.

One metric is the manipulability scalar:

$$w = \sqrt{\det(JJ^{\top})}, \tag{4.6}$$

where J is the Jacobian associated with the point being evaluated. Larger values indicate greater freedom of motion in Cartesian space, while smaller values indicate that the robot is closer to a singular or constrained configuration. The framework also evaluates the alignment between a repulsive vector or equivalent movement restriction and the direction of lowest manipulability. This is measured by projecting the repulsive force $V(p)$ onto the minimum-operability eigenvector v_{min} of the manipulability ellipsoid:

$$\text{proj}_{v_{min}} V(p) = \frac{V(p) \cdot v_{min}}{\|v_{min}\|^2}. \tag{4.7}$$

If a controller frequently directs a control point to move along the axis of lowest manipulability, then it is asking the robot to move in a direction that is difficult to realize. Fig. 4.4 illustrates this by contrasting high- and low-manipulability configurations. The evaluation results later show that this metric reveals distinct differences among controllers, especially when constraints are applied near the base of the robot.

4.4.3 Motion Smoothness

Motion smoothness evaluates how the robot's trajectory evolves over time. Although object-aware control may require quick reactions to nearby obstacles, those responses should avoid abrupt velocity or acceleration changes. Sudden commands can increase mechanical wear, introduce vibration, cause servo errors, and make human–robot interaction less predictable. The framework analyzes motion profiles, including position, velocity, acceleration, and jerk. Emphasis is placed on jerk because spikes in jerk indicate that a controller is producing abrupt changes in motion. Fig. 4.4 illustrates the difference between low-jerk and high-jerk behavior: in the low-jerk case, the end-effector moves smoothly away from the obstacle, while in the high-jerk case, the path becomes jagged.

This criterion is meaningful because a controller can technically avoid an obstacle while still producing undesirable robot behavior. A high-jerk response may indicate that the controller is switching constraints abruptly, changing control points discontinuously, or generating unstable movement profiles. Jerk provides a measure of whether the controller's motion is safe, mechanically reasonable, and predictable during interaction. Predictable, low-jerk motion also makes the robot's behavior easier for nearby people to anticipate, echoing the distinction between predictable and legible motion in human-robot interaction [20].

4.4.4 Reactive Constraint Behavior

Reactive constraint behavior describes how motion requirements are generated in response to nearby objects. OACs often use virtual constraints or repulsive forces to modify robot motion, encoding the relationship between the robot and surrounding obstacles.

The framework analyzes the location of control points, the magnitude and direction of repulsive forces, and how these quantities change as the obstacle moves. Control point selection is especially consequential. A controller that uses statically defined control points may produce smoother control point behavior, but may not always constrain the point on the robot closest to the obstacle. A controller that uses dynamically selected control points may better track the closest point on the body, but may introduce discontinuities when the selected control point changes suddenly.

This distinction is shown in Fig. 4.4, which contrasts static and dynamic control points. It is also reflected in the controller diagrams in Figs. 4.5–4.7. The Flacco controller uses statically defined control points, while Ding and Escobedo select control points dynamically based on proximity to the obstacle. These differences affect manipulability, smoothness, and the stability of the resulting avoidance behavior.

4.5 Contact-Admitting Control

Anticipation and avoidance both aim to reduce contact. But in cluttered, shared spaces, a collision-free path may not exist, and some contact is not only unavoidable but useful. The third contribution of this chapter treats contact as something a robot can deliberately admit rather than always prevent [64].

CAT-RRT is a motion planner that admits contact one link at a time. Instead of rejecting any path that touches an obstacle, it reasons about contact on a per-link basis, selectively permitting contact where it is acceptable while still constraining it elsewhere. This lets the robot plan through spaces where a strictly collision-free path would fail, complementing the reactive behaviors above:

anticipation softens contact before it happens, avoidance steers around obstacles when feasible, and contact-admitting planning allows deliberate contact when the task requires it.

4.6 Experimental Evaluation

The evaluation is in two parts. The first validates the contact anticipation framework in physical human–robot interaction [23]. The second compares object-aware controllers under matched obstacle conditions [22].

4.6.1 Contact Anticipation Evaluation

The contact anticipation framework is evaluated on a real Franka Emika Panda robot using static obstacle collision, obstacle avoidance, and dynamic obstacle collision scenarios. These scenarios test whether proximity sensing reduces contact forces, whether the robot can avoid or deviate around nearby obstacles, and whether the full framework supports repeated physical interaction with a human.

In the static obstacle collision scenario, the robot follows predefined trajectories while making contact with a static obstacle. The same movements are tested with and without onboard sensor information. The results demonstrate that contact forces are approximately two times smaller when proximity sensing is used to anticipate the collision [23]. This reduction occurs because the robot slows down as the object enters nearby space and lowers the contact threshold in the direction of the detected object. The reduction shows that proximity sensing can reduce contact severity before physical interaction occurs.

In the obstacle avoidance scenario, the robot follows a circular trajectory while a human enters its path. When the human enters the robot's nearby space, the controller scales the robot's velocity and imposes movement restrictions that cause a slight deviation from the original trajectory. Once the object is no longer close enough to influence the robot, the manipulator returns to its original path. This scenario shows the proactive trajectory adjustment, where the robot adapts its path before impact rather than responding only after contact [23].

The dynamic obstacle collision scenario demonstrates the full interaction cycle. A human makes multiple contacts with the robot from different directions while the robot is moving. The system detects nearby objects through multiple sensor units, dynamically adjusts contact thresholds, identifies contact, moves away from the direction of applied force, and then resumes the original trajectory. This interaction is summarized in forces in Fig. 4.3 and Cartesian position in Fig. 4.8. This scenario tests the full interaction loop by combining object detection, threshold adaptation, contact response, and task recovery.

Across the three scenarios, the results demonstrate that onboard proximity sensing and dynamic contact thresholding allow the robot to operate in a space between complete avoidance and passive collision response.

4.6.2 Object-Aware Controller Evaluation

The OAC framework compares Flacco, Ding, and Escobedo using two fundamental robot-obstacle interaction scenarios. In the *Static Robot, Dynamic Obstacle* condition, the manipulator maintains a fixed end-effector position as an obstacle moves toward the body. In the *Dynamic Robot, Dynamic Obstacle* condition, the manipulator follows a circular Cartesian trajectory while an obstacle approaches the end-effector. Both experiments isolate controller response while keeping obstacle information consistent across methods.

The first metric concerns manipulability. Fig. 4.9 shows the manipulability scalar values for control points and the end-effector during the dynamic robot experiment. Flacco's statically defined control points produce smoother manipulability profiles, while Ding and Escobedo exhibit sudden changes as the dynamically selected control point switches to a different part of the robot body. These discontinuities matter because the new control point may have a substantially different manipulability value. A point with higher manipulability can often avoid an obstacle without strongly affecting the main task, while a point with lower manipulability may be physically constrained and less able to satisfy the imposed movement restriction.

The second kinematic metric uses the manipulability ellipsoid. Fig. 4.10 shows the proportion of each controller's repulsive vector or equivalent restriction projected onto the direction of lowest axial operability, as per Eq. 4.7. Values near one indicate that the controller is commanding motion primarily along the axis of lowest manipulability. This is undesirable because the robot is being asked to move in a direction that is difficult or infeasible from its current configuration. The results show that constraints imposed near proximal links can be especially problematic because points near the base of the robot are naturally more kinematically constrained.

Motion smoothness is evaluated through joint jerk. Fig. 4.11 shows the jerk profiles of the robot joints during obstacle avoidance. Spikes in jerk indicate abrupt changes in commanded motion. These spikes reveal that certain controller behaviors, such as sudden changes in control point selection or constraint activation, can generate unstable motion profiles. This result supports evaluating object-aware controllers based on both obstacle avoidance and the smoothness and mechanical reasonableness of the resulting motion.

Reactive constraint behavior is examined through the end-effector paths generated during the dynamic robot experiment. Fig. 4.12 shows the end-effector paths generated by the different controllers. The paths reveal how controller design choices affect the resulting avoidance trajectory. Some controllers produce smooth deviations from the nominal path, while others generate less stable or more discontinuous behavior. These differences reflect the underlying design choices identified earlier, including static versus dynamic control points, repulsive vectors versus hard constraints, velocity scaling, and the treatment of the robot's kinematic structure.

Together, the experiments show substantial differences among OACs even when they receive the same obstacle information. The metrics make these differences visible through analysis of kinematic feasibility, motion smoothness, and reactive constraints. These results suggest that future designs should account for manipulability when selecting control points and scaling repulsive behavior, avoid unrealistic constraints near kinematically limited regions, and reduce discontinuities in movement profiles.

4.7 Discussion & Conclusion

This chapter turned nearby-space perception into control for safe physical interaction. In response to RQ3, it showed how object-aware controllers use nearby-object information to modify robot motion, and how that motion should be judged beyond collision-avoidance success. The contact anticipation framework uses proximity readings to estimate nearby object positions, reduce velocity, impose movement restrictions, dynamically adjust contact thresholds, and trigger post-contact reactions, reducing contact force and supporting repeated transitions between avoidance and contact [23]. The evaluation framework showed that controllers receiving identical obstacle information can produce very different motion, and that manipulability, jerk, and reactive constraint behavior are necessary criteria for judging whether that motion is safe and kinematically feasible [22]. Contact-admitting planning extended the range of behavior further, allowing deliberate contact when a collision-free path is not available [64].

Taken together with the previous chapters, this completes the arc from hardware to application. Fabrication (Chapter 2) places sensors on the body with known locations; perception (Chapter 3) grounds and integrates their readings into nearby-space awareness; and this chapter uses that awareness to anticipate, avoid, and admit contact during physical interaction. Whole-body sensing matters most when it supports action, and the value of that action depends on control that is safe, interpretable, and physically feasible.

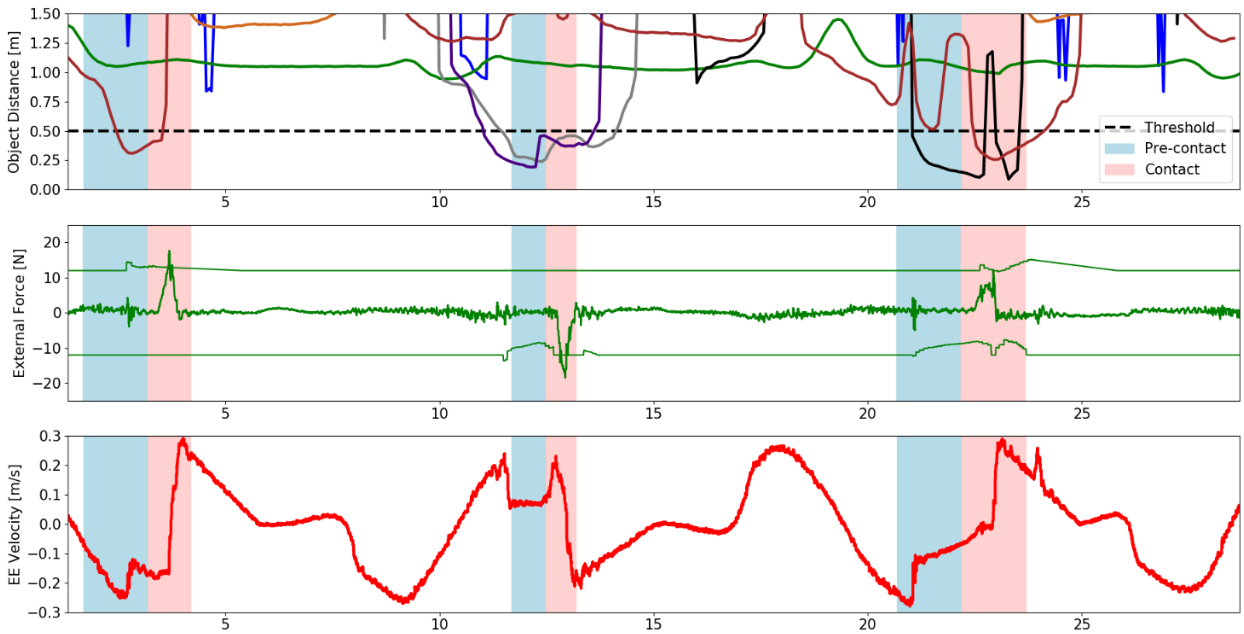

Figure 4.3: Three components of a continuous human–robot interaction over 30 seconds: during this time, three distinct contacts are made from multiple directions. Top graph: depth values of seven proximity sensors placed along the robot body and EE, with the activation threshold as a dotted black line. Middle graph: external force exerted on the robot EE in the y-axis with the dynamic thresholds shown above and below. Bottom graph: The manipulator's EE velocity along the y-axis.

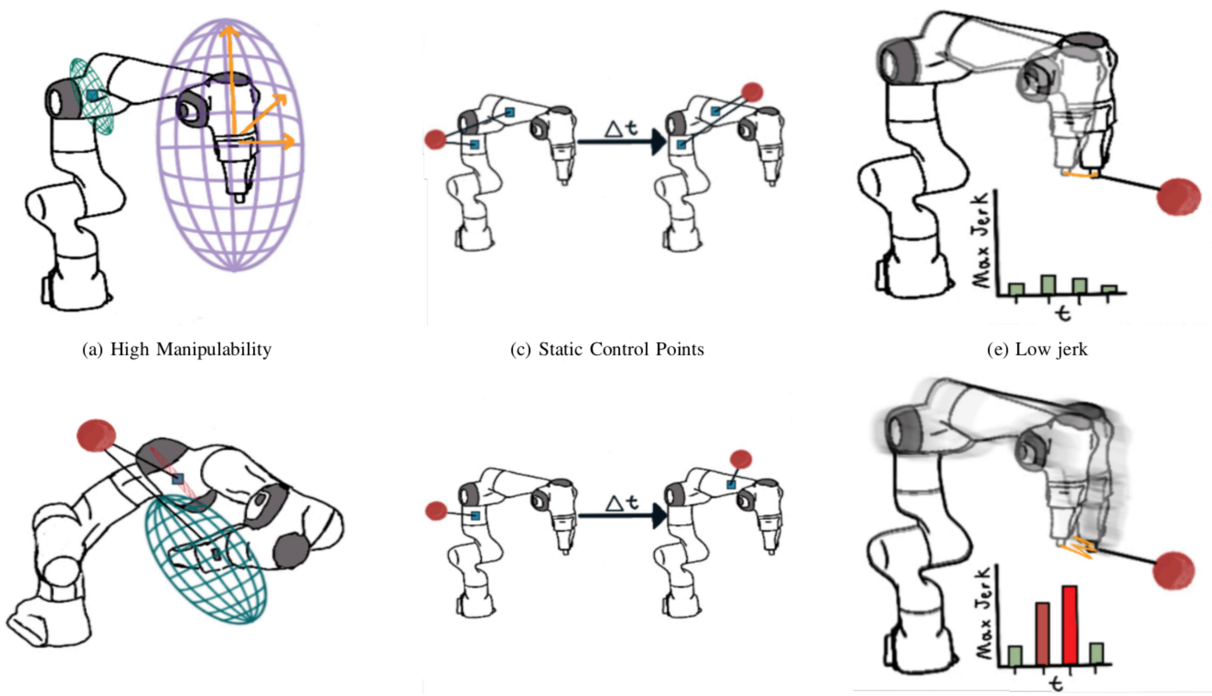

(a) High Manipulability (c) Static Control Points (e) Low jerk

(b) Low Manipulability (d) Dynamic Control Points (f) High jerk

Figure 4.4: Each column represents one of the three design considerations used to evaluate controllers in this work. **(a)** High-manipulability configurations, shown from the large manipulability ellipsoids. Eigenvectors of $\mathbf{J}\mathbf{J}^{\top}$, where $\mathbf{J}$ is the robot Jacobian, are shown as orange arrows. **(b)** An undesirable configuration caused by movement restrictions from the red object, where the smaller ellipsoids indicate that the arm has less ability to move in an arbitrary direction freely. **(c)** Static control points are determined by the controller designer and placed in predefined locations. **(d)** Dynamic control points are selected by determining the closest point on the robot body to an obstacle. **(e)** Low jerk is exhibited when the robot smoothly moves away from the obstacle, as shown by the orange path. **(f)** Movement restrictions imposed on the end-effector cause jittery motion, as seen in the jagged path to its final location.

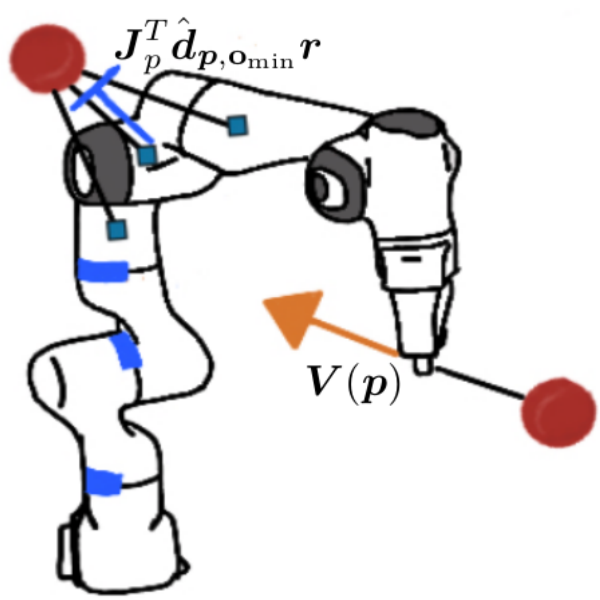

Figure 4.5: Flacco [27] controller diagram with obstacles represented as red circles. The orange arrow originating from the end-effector is a repulsive force. Blue rectangles show joint restrictions caused by constraints added for this control formulation.

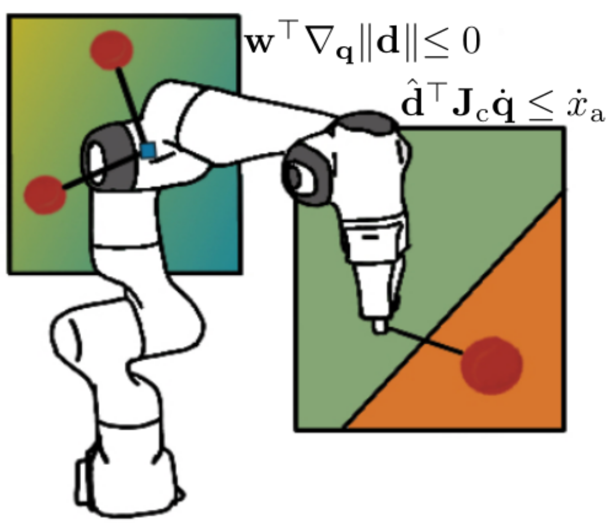

Figure 4.6: Ding [14] controller diagram with obstacles represented as red circles. The rectangle encapsulating the EE shows movement restrictions, the orange portion shows restricted movement while the green shows where the robot can still move. This first restriction is also added to the body control points, when obstacles are nearby. The rectangle near the robot body shows the weighted sum distance gradient restriction on the robot motion.

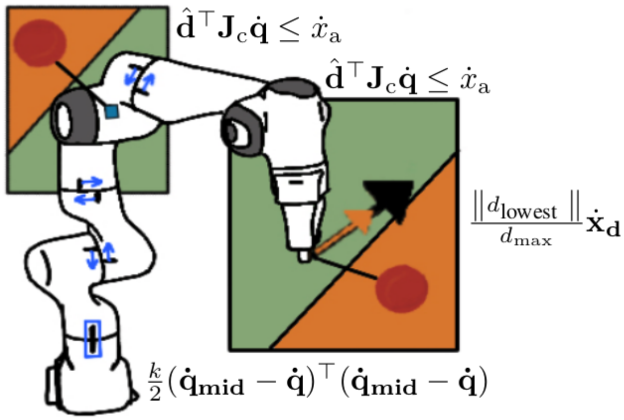

Figure 4.7: Escobedo [23] controller diagram with obstacles represented as red circles. The rectangles show movement restrictions from the prior-to-contact behavior in Section 4.3.2; the orange portion shows restricted movement while the green shows where the robot can still move. The orange arrow shows the scaled EE velocity from Eq. 4.2, and the black arrow is the initial velocity.

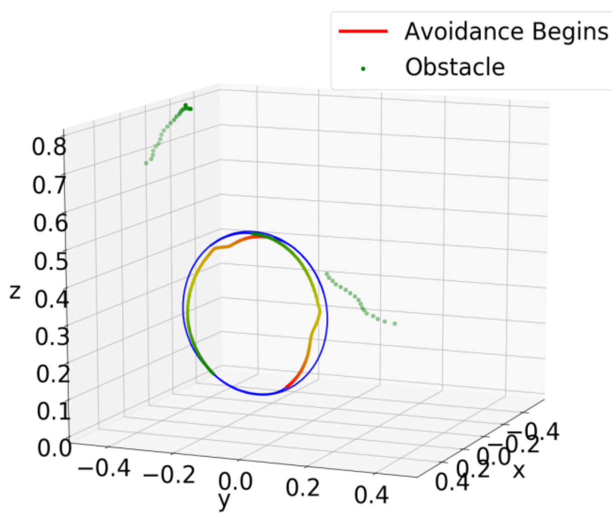

Figure 4.8: Robot EE moving in a circular path (blue circle); when an object is detected (green dots), the end-effector trajectory is modified to avoid collision (red) then gradually recovers (green). We demonstrate this behavior in two interactions on opposite sides of the robot; the objects are detected by independent proximity sensors.

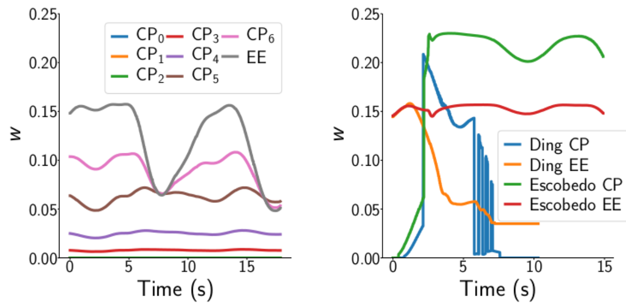

Figure 4.9: The manipulability at the control points (CP) and the end-effector (EE) tracked across time for the five control points of Flacco's work (left), and the control point and end-effector of Ding and Escobedo (right). Both graphs are constructed from the dynamic robot and dynamic object experimental scenario.

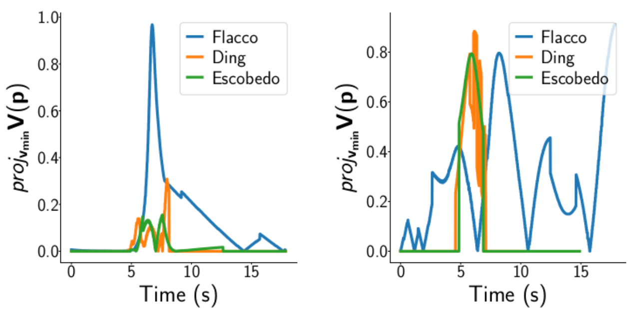

Figure 4.10: The proportion of the repulsive vector's magnitude projected onto the minimum dimension of the manipulability ellipsoid as per Eq. 4.7. The left graph is from the static robot dynamic object scenario and the right graph is from the dynamic robot dynamic object scenario.

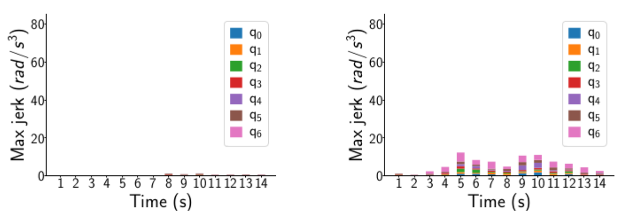

(a) No obstacle

(b) Flacco

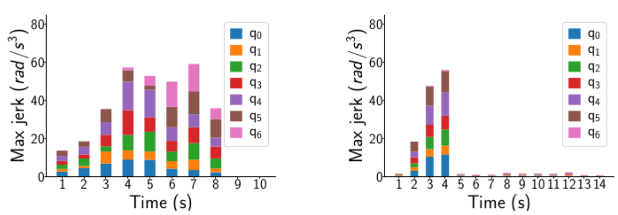

Figure 4.11: Maximum jerk per second of experimental time plotted for each joint of the robot arm. The spikes in jerk correspond to avoidance behaviors in the dynamic robot dynamic object scenario.

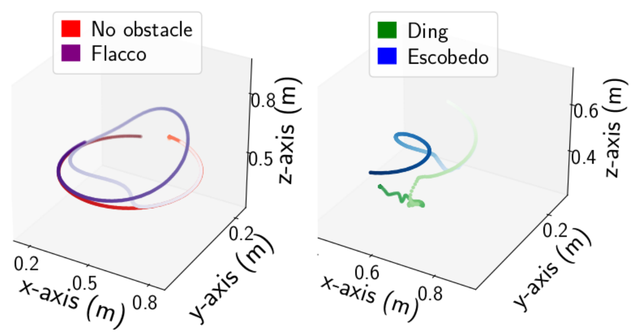

Figure 4.12: The end-effector path of the robot as it is commanded by each controller to travel in a circular path while being approached by an obstacle. Both graphs are constructed from the dynamic robot dynamic object experimental scenario.

Chapter 5

Conclusion

Robots are beginning to move out of the isolated cells they were built for and into homes, hospitals, and laboratories, spaces defined by people, clutter, and unplanned contact. This shift is still tentative and narrow, because a robot cannot safely share space with people until it can sense across its entire body, and whole-body sensing has so far remained the privilege of a few specialized labs. This dissertation argued that the central challenge in whole-body robot sensing is not inventing a new sensor, but building skins that are practical to fabricate, reliable to use, and adaptable across robots, tasks, and sensing modalities. It made the case that multi-material additive manufacturing is a foundation for whole-body robot skins, and that skins built on it are repeatable, scalable, and compatible with diverse sensing modalities across robot platforms. The work was organized around three research questions that move from hardware to application: fabrication establishes the foundation, perception grounds distributed sensing on the robot body, and object-aware control turns that sensing into safe physical interaction.

RQ1 asked: How can multi-material additive manufacturing enable the reproducible fabrication of whole-body robotic skins with diverse sensing modalities? This was addressed by treating fabrication as a first-class research problem rather than a prototyping detail. Whole-body skins were generated from a robot's own 3D model and printed as form-fitting units that combine structure, conductive sensing regions, embedded routing, sensor mounts, and compliant coverings in a single artifact [40, 43, 41, 82]. The concrete lesson is that sensor placement should be fixed in the digital model and preserved by the print rather than recovered per build:

printed snap-in mounts hold off-the-shelf sensors at known poses through repeated removal and reapplication, and co-printing standard PLA, conductive PLA, and TPU lets structure, sensing, and compliance be fabricated together. The same workflow produced working skins spanning capacitive, time-of-flight, hybrid, and magnetic sensing, deployed across the Unitree H1-2 humanoid, the Unitree Go1 quadruped, and the Franka Research 3 arm. The actionable outcome is a path any researcher with a multi-material printer and a surface mesh can reproduce, modify, or extend.

RQ2 asked: How can distributed onboard sensing be calibrated and integrated to support reliable nearby-space perception? This was addressed by showing that a distributed skin is only useful once the robot knows where each reading originates on its body and how onboard sensing relates to external perception. Kinematic calibration recovered the pose of each sensor unit accurately enough to redirect motion around a nearby person, treating sensor placement as a calibration problem grounded in the robot's kinematic chain [91]. Volumetric fusion then combined onboard proximity with external depth in a single occupancy representation, weighting each sensor by its optimal sensing range and filling the occluded near-body region that external cameras cannot observe [83]. The concrete lesson is that distance readings must be transformed into the body frame before they can inform control, and that the robot's own body should be treated as an active perceptual surface for the roughly 20 cm shell where close interaction is decided.

RQ3 asked: How can this perception be utilized by object-aware controllers to enable safe, kinematically feasible contact anticipation and avoidance? This was addressed by turning nearby-space perception into motion and by showing how that motion should be judged. A contact anticipation framework used onboard proximity to slow the robot, impose movement restrictions, and dynamically adjust contact thresholds, roughly halving contact forces and supporting smooth transitions between avoidance and contact [23]. A systematic evaluation of object-aware controllers showed that controllers receiving identical obstacle information can produce very different motion, and that collision-avoidance success is not enough: manipulability, jerk, and reactive constraint behavior are necessary criteria, and avoidance constraints should not be placed on low-manipulability proximal links [22]. Contact-admitting planning extended the range of behavior further, letting a robot deliberately permit contact one link at a time when a collision-free path is not available [64]. The concrete lesson is that whole-body sensing becomes valuable when control is designed to be manipulability-aware, smooth, and able to move along a continuum from avoiding contact, to anticipating it, to admitting it on purpose.

Taken together, the three research questions support a single conclusion: progress in whole-body robot skins depends on more than a new sensor. It depends on placing sensors where they matter, fabricating them with consistency, grounding their readings on the body, and using them in control that is safe and physically feasible. Fabrication makes the hardware reproducible, perception makes the readings trustworthy, and object-aware control makes them useful. This dissertation connects the behavioral value of whole-body sensing to the manufacturing methods that make such systems practical to build and share. The larger contribution is a change in what whole-body sensing demands of the people who want to use it. When a capability moves from a one-off build that requires specialized hardware knowledge, custom wiring, and months of integration to something reproducible from a shared digital file, it becomes a tool others can simply assume rather than a research problem of its own. A human-robot interaction researcher studying physical collaboration, or a team building learning systems that need contact-rich data, can then treat whole-body perception as a starting point rather than a hardware project that must come first.

5.1 Future Work

Several directions follow directly from this work. On the fabrication side, the open question is which sensing modalities are most useful for which applications; the magnetic and hybrid skins here are early evidence that new modalities can be added at low engineering cost, and the more the community explores this space, the faster the shared architecture will mature. On the perception side, onboard sensing should be integrated more tightly with external perception and with learning-based representations, so that the near-body region becomes a routine input to modern perception pipelines rather than a special case. Closely related is the robot's representation of its own body: adapting a geometric model of the robot's morphology to the task at hand, rather than holding it fixed, lets a system balance computational efficiency against the fidelity needed for safe collision checking and contact-rich manipulation, and the onboard sensing developed here is a natural signal for driving that adaptation [65]. On the control side, contact anticipation and contact-admitting behavior should be unified into controllers that reason continuously across the avoid, anticipate, admit continuum, and evaluated with the kinematic and smoothness criteria developed here. One direction is deployment in human-centric environments. We showed a completed example in a robot-assisted chemical dialysis system, where a manipulator worked alongside scientists to carry out a real, multi-day laboratory procedure [38]. The robot fit into the scientists' existing space and workflow rather than a purpose-built cell. Related manipulation work targets these same environments, learning container dynamics to move open liquids through cluttered laboratory spaces such as fume hoods without spilling [1], exactly the kind of dynamic, contact-sensitive task where whole-body sensing improves safety. That is the setting whole-body sensing is meant to support: robots placed in labs, clinics, and homes to help people without redesigning those spaces or making the people working near them uncomfortable. Sensing proximity and contact across the whole body is what lets a robot share that space safely, keeping clear when it should and admitting contact when a task calls for it. Finally, the field needs better ways to share designs, compare results, and build on prior systems; because the skins in this dissertation are defined by shareable digital files, they offer a concrete starting point for that kind of community reuse.

This dissertation showed that whole-body robot skins require advances in both perception and hardware, and that multi-material additive manufacturing offers a clear path forward. Its deeper contribution is to make whole-body perception ordinary: something any lab can fabricate, any robot can wear, and any subfield can build on. For robots to earn a place in the spaces where people live and work, they must be able to feel the world across their entire bodies; this dissertation turns that ability from a barrier only a few labs can clear into a foundation the whole field can build the next generation of human-centered robots on.

Bibliography

[1] Ava Abderezaei, Anuj Pasricha, Alex Klausenstock, and Alessandro Roncone. Clutter-aware spill-free liquid transport via learned dynamics. In 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 9908–9915. IEEE, 2024.

[2] Hosam Alagi, Alexander Heiligl, Stefan Escaida Navarro, Torsten Kroegerl, and Björn Hein. Material recognition using a capacitive proximity sensor with flexible spatial resolution. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 6284–6290. IEEE, 2018.

[3] Alessandro Albini, Mohsen Kaboli, Giorgio Cannata, and Perla Maiolino. Representing data in robotic tactile perception–a review. arXiv preprint arXiv:2510.10804, 2025.

[4] Giovanni Buizza Avanzini, Nicola Maria Ceriani, Andrea Maria Zanchettin, Paolo Rocco, and Luca Bascetta. Safety control of industrial robots based on a distributed distance sensor. IEEE Transactions on Control Systems Technology, 22(6):2127–2140, 2014.

[5] S Sandra Bae, Takanori Fujiwara, Danielle Szafir, Ellen Yi-Luen Do, and Michael L Rivera. Computational design and single-wire sensing of 3d printed objects with integrated capacitive touchpoints. In Proceedings of the ACM Symposium on Computational Fabrication, pages 1–19, 2025.

[6] S Sandra Bae, Takanori Fujiwara, Anders Ynnerman, Ellen Yi-Luen Do, Michael L Rivera, and Danielle Albers Szafir. A computational design pipeline to fabricate sensing network physicalizations. IEEE Transactions on Visualization and Computer Graphics, 30(1):913–923, 2023.

[7] Raunaq Bhirangi, Tess Hellebrekers, Carmel Majidi, and Abhinav Gupta. Reskin: versatile, replaceable, lasting tactile skins. In 5th Annual Conference on Robot Learning, 2021.

[8] Giorgio Cannata, Marco Maggiali, Giorgio Metta, and Giulio Sandini. An embedded artificial skin for humanoid robots. In 2008 IEEE International conference on multisensor fusion and integration for intelligent systems, pages 434–438. IEEE, 2008.

[9] Haohua Chen, Yixuan Zhou, Yifan Zhou, and Hesheng Wang. Cssdf-net: Safe motion planning based on neural implicit representations of configuration space distance field. arXiv preprint arXiv:2603.18669, 2026.

[10] Gordon Cheng, Emmanuel Dean-Leon, Florian Bergner, Julio Rogelio Guadarrama Olvera, Quentin Leboutet, and Philipp Mittendorfer. A comprehensive realization of robot skin: Sensors, sensing, control, and applications. Proceedings of the IEEE, 107(10):2034–2051, 2019.

[11] Edward Cheung and Vladimir Lumelsky. A sensitive skin system for motion control of robot arm manipulators. Robotics and Autonomous Systems, 10(1):9–32, 1992.

[12] 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.

[13] Ravinder S Dahiya, Philipp Mittendorfer, Maurizio Valle, Gordon Cheng, and Vladimir J Lumelsky. Directions toward effective utilization of tactile skin: A review. IEEE Sensors Journal, 13(11):4121–4138, 2013.

[14] Yitao Ding and Ulrike Thomas. Collision avoidance with proximity servoing for redundant serial robot manipulators. In 2020 IEEE International Conference on Robotics and Automation (ICRA), pages 10249–10255. IEEE, 2020.

[15] Yitao Ding and Ulrike Thomas. Improving safety and accuracy of impedance controlled robot manipulators with proximity perception and proactive impact reactions. In 2021 IEEE International Conference on Robotics and Automation (ICRA), pages 3816–3821, 2021.

[16] Yitao Ding, Hongyin Zhang, and Ulrike Thomas. Capacitive proximity sensor skin for contactless material detection. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 7179–7184. IEEE, 2018.

[17] Won Kyung Do, Ankush Kundan Dhawan, Mathilda Kitzmann, and Monroe Kennedy. Densetact-mini: an optical tactile sensor for grasping multi-scale objects from flat surfaces. In 2024 IEEE international conference on robotics and automation (ICRA), pages 6928–6934. IEEE, 2024.

[18] Won Kyung Do, Bianca Jurewicz, and Monroe Kennedy III. Densetact 2.0: Optical tactile sensor for shape and force reconstruction, 2023.

[19] Won Kyung Do and Monroe Kennedy. Densetact: Optical tactile sensor for dense shape reconstruction. In 2022 International Conference on Robotics and Automation (ICRA), pages 6188–6194. IEEE, 2022.

[20] Anca D Dragan, Kenton CT Lee, and Siddhartha S Srinivasa. Legibility and predictability of robot motion. In 2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI), pages 301–308. IEEE, 2013.

[21] Zackory Erickson, Henry M. Clever, Vamsee Gangaram, Greg Turk, C. Karen Liu, and Charles C. Kemp. Multidimensional capacitive sensing for robot-assisted dressing and bathing, 2019.

[22] Caleb Escobedo, Nataliya Nechyporenko, Shreyas Kadekodi, and Alessandro Roncone. A framework for the systematic evaluation of obstacle avoidance and object-aware controllers. In 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 8117–8124. IEEE, 2022.

[23] Caleb Escobedo, Matthew Strong, Mary West, Ander Aramburu, and Alessandro Roncone. Contact anticipation for physical human–robot interaction with robotic manipulators using onboard proximity sensors. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 7255–7262. IEEE, 2021.

[24] Xiaoran Fan, Riley Simmons-Edler, Daewon Lee, Larry Jackel, Richard Howard, and Daniel Lee. Aurasense: Robot collision avoidance by full surface proximity detection. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 1763–1770. IEEE, 2021.

[25] Faraz Faruqi, Ahmed Katary, Tarik Hasic, Amira Abdel-Rahman, Nayeemur Rahman, Leandra Tejedor, Mackenzie Leake, Megan Hofmann, and Stefanie Mueller. Style2fab: Functionality-aware segmentation for fabricating personalized 3d models with generative ai. In Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology, UIST '23, page 1–13. ACM, October 2023.

[26] Jeremy A Fishel and Gerald E Loeb. Sensing tactile microvibrations with the biotac—comparison with human sensitivity. In 2012 4th IEEE RAS & EMBS international conference on biomedical robotics and biomechatronics (BioRob), pages 1122–1127. IEEE, 2012.

[27] Fabrizio Flacco, Torsten Kröger, Alessandro De Luca, and Oussama Khatib. A depth space approach to human-robot collision avoidance. In 2012 IEEE international conference on robotics and automation, pages 338–345. IEEE, 2012.

[28] Diogo Fonseca, Mohammad Safeea, and Pedro Neto. A flexible piezoresistive/self-capacitive hybrid force and proximity sensor to interface collaborative robots. IEEE Transactions on Industrial Informatics, 19(3):2485–2495, 2022.

[29] Francesco Grella, Alessandro Albini, and Giorgio Cannata. Voluntary interaction detection for safe human-robot collaboration. In 2022 sixth IEEE international conference on robotic computing (IRC), pages 353–359. IEEE, 2022.

[30] Sami Haddadin, Alessandro De Luca, and Alin Albu-Schäffer. Robot collisions: A survey on detection, isolation, and identification. IEEE Transactions on Robotics, 33(6):1292–1312, 2017.

[31] Md Sahid Hassan, Saqlain Zaman, Joshua Z. R. Dantzler, Diana Hazel Leyva, Md Shahjahan Mahmud, Jean Montes Ramirez, Sofia Gabriela Gomez, and Yirong Lin. 3d printed integrated sensors: From fabrication to applications—a review. Nanomaterials, 13(24), 2023.

[32] Jesse Haviland and Peter Corke. A purely-reactive manipulability-maximising motion controller. arXiv preprint arXiv:2002.11901, 2020.

[33] Jesse Haviland and Peter Corke. Neo: A novel expeditious optimisation algorithm for reactive motion control of manipulators. IEEE Robotics and Automation Letters, 6(2):1043–1050, 2021.

[34] Van Anh Ho, Masaaki Makikawa, and Shinichi Hirai. Flexible fabric sensor toward a humanoid robot's skin: Fabrication, characterization, and perceptions. IEEE Sensors Journal, 13(10):4065–4080, 2013.

[35] Dana Hughes and Nikolaus Correll. A soft, amorphous skin that can sense and localize textures. In 2014 IEEE International Conference on Robotics and Automation (ICRA), pages 1844–1851. IEEE, 2014.

[36] Dana Hughes, John Lammie, and Nikolaus Correll. A robotic skin for collision avoidance and affective touch recognition. IEEE Robotics and Automation Letters, 3(3):1386–1393, 2018.

[37] Fumiya Iida, Chapa Sirithunge, Perla Maiolino, and Josie Hughes. Foundations of embodied intelligence for robotic systems. Foundations and Trends® in Robotics, 12(4):350–425, 2026.

[38] Diane N Jung, Caleb Escobedo, Noah Liska, Maitrey Gramopadhye, Daniel Szafir, Alessandro Roncone, and Carson J Bruns. Design of a robot-assisted chemical dialysis system. In Companion Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction, pages 416–420, 2026.

[39] Daehwa Kim, Mario Srouji, Chen Chen, and Jian Zhang. Armor: Egocentric perception for humanoid robot collision avoidance and motion planning, 2024.

[40] Carson Kohlbrenner, Caleb Escobedo, S Sandra Bae, Alexander Dickhans, and Alessandro Roncone. Gentact toolbox: A computational design pipeline to procedurally generate context-driven 3d printed whole-body artificial skins. In 2025 IEEE International Conference on Robotics and Automation (ICRA), pages 4716–4722. IEEE, 2025.

[41] Carson Kohlbrenner, Caleb Escobedo, Sayak Ray, Alexander Dickhans, Anna Soukhovei, Nickolaus Jackoski, Lyle Antieau, and Alessandro Roncone. Improving sensing coverage and compliance of 3d-printed artificial skins through multi-modal sensing and soft materials. arXiv preprint arXiv:2604.25563, 2026.

[42] Carson Kohlbrenner, Mitchell Murray, Yutong Zhang, Caleb Escobedo, Thomas Dunnington, Nolan Stevenson, Nikolaus Correll, and Alessandro Roncone. A machine learning approach to contact localization in variable density three-dimensional tactile artificial skin. In 2nd NeurIPS Workshop on Touch Processing: From Data to Knowledge, 2024.

[43] Carson Kohlbrenner, Anna Soukhovei, Caleb Escobedo, Nataliya Nechyporenko, and Alessandro Roncone. Design, mapping, and contact anticipation with 3d-printed whole-body tactile and proximity sensors. arXiv preprint arXiv:2603.04714, 2026.

[44] Mikhail Koptev, Nadia Figueroa, and Aude Billard. Real-time self-collision avoidance in joint space for humanoid robots. IEEE Robotics and Automation Letters, 6(2):1240–1247, 2021.

[45] Mikhail Koptev, Nadia Figueroa, and Aude Billard. Neural joint space implicit signed distance functions for reactive robot manipulator control. IEEE Robotics and Automation Letters, 8(2):480–487, 2022.

[46] Adarsh K. Kosta, Alexis M. Burns, Caleb Escobedo, Siddharth Rupavatharam, Richard E. Howard, Lawrence Jackel, Daewon Lee, and Ibrahim Volkan Isler. Ultrasonic piezoelectric transceiver sensor for full surface contact localization, May 2026. U.S. Patent US 12,619,250 B2, assigned to Samsung Electronics Co., Ltd.

[47] Adarsh Kumar Kosta, Alexis Burns, Siddharth Rupavatharam, Caleb Escobedo, Daewon Lee, Richard Howard, Larry Jackel, and Volkan Isler. Acouskin: Full surface contact localization using acoustic waves. In 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 10821–10828. IEEE, 2023.

[48] Mike Lambeta, Po-Wei Chou, Stephen Tian, Brian Yang, Benjamin Maloon, Victoria Rose Most, Dave Stroud, Raymond Santos, Ahmad Byagowi, Gregg Kammerer, Dinesh Jayaraman, and Roberto Calandra. Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation. IEEE Robotics and Automation Letters, 5(3):3838–3845, 2020.

[49] Mike Lambeta, Tingfan Wu, Ali Sengul, Victoria Rose Most, Nolan Black, Kevin Sawyer, Romeo Mercado, Haozhi Qi, Alexander Sohn, Byron Taylor, Norb Tydingco, Gregg Kammerer, Dave Stroud, Jake Khatha, Kurt Jenkins, Kyle Most, Neal Stein, Ricardo Chavira, Thomas Craven-Bartle, Eric Sanchez, Yitian Ding, Jitendra Malik, and Roberto Calandra. Digitizing touch with an artificial multimodal fingertip, 2024.

[50] Przemyslaw A Lasota, Terrence Fong, and Julie A Shah. A survey of methods for safe human-robot interaction. Foundations and Trends® in Robotics, 5(4):261–349, 2017.

[51] Moonyoung Lee, Uksang Yoo, Jean Oh, Jeffrey Ichnowski, George Kantor, and Oliver Kroemer. Sonicboom: Contact localization using array of microphones, 2024.

[52] Martina Lippi, Giuseppe Gillini, Alessandro Marino, and Filippo Arrichiello. A data-driven approach for contact detection, classification and reaction in physical human-robot collaboration. In 2021 IEEE International Conference on Robotics and Automation (ICRA), pages 3597–3603. IEEE, 2021.

[53] V.J. Lumelsky and E. Cheung. Real-time collision avoidance in teleoperated whole-sensitive robot arm manipulators. IEEE Transactions on Systems, Man, and Cybernetics, 23(1):194–203, 1993.

[54] Quan Khanh Luu, Nhan Huu Nguyen, et al. Simulation, learning, and application of vision-based tactile sensing at large scale. IEEE Transactions on Robotics, 39(3):2003–2019, 2023.

[55] 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.

[56] Michael A. McEvoy and Nikolaus Correll. Materials that couple sensing, actuation, computation, and communication. Science, 347(6228):1261689, 2015.

[57] Kelly Merckaert, Bryan Convens, Chi-ju Wu, Alessandro Roncone, Marco M Nicotra, and Bram Vanderborght. Real-time motion control of robotic manipulators for safe human–robot coexistence. Robotics and Computer-Integrated Manufacturing, 73:102223, 2022.

[58] Philipp Mittendorfer and Gordon Cheng. Uniform cellular design of artificial robotic skin. In ROBOTIK 2012; 7th German Conference on Robotics, pages 1–5. VDE, 2012.

[59] Philipp Mittendorfer, Eiichi Yoshida, and Gordon Cheng. Realizing whole-body tactile interactions with a self-organizing, multi-modal artificial skin on a humanoid robot. Advanced Robotics, 29(1):51–67, 2015.

[60] Fangzhou Mu, Carter Sifferman, Sacha Jungerman, Yiquan Li, Mark Han, Michael Gleicher, Mohit Gupta, and Yin Li. Towards 3d vision with low-cost single-photon cameras. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 5302–5311, 2024.

[61] Yashraj S. Narang, Balakumar Sundaralingam, Karl Van Wyk, Arsalan Mousavian, and Dieter Fox. Interpreting and predicting tactile signals for the syntouch biotac, 2021.

[62] Stefan Escaida Navarro, Maximiliano Marufo, Yitao Ding, Stephan Puls, Dirk Göger, Björn Hein, and Heinz Wörn. Methods for safe human-robot-interaction using capacitive tactile proximity sensors. In 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 1149–1154. IEEE, 2013.

[63] Stefan Escaida Navarro, Stephan Muhlbacher-Karrer, Hosam Alagi, Hubert Zangl, Keisuke Koyama, Bjorn Hein, Christian Duriez, and Joshua R. Smith. Proximity perception in human-centered robotics: A survey on sensing systems and applications. IEEE Transactions on Robotics, 38(3):1599–1620, June 2022.

[64] Nataliya Nechyporenko, Caleb Escobedo, Shreyas Kadekodi, and Alessandro Roncone. Cat-rrt: Motion planning that admits contact one link at a time. In 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 4849–4856. IEEE, 2023.

[65] Nataliya Nechyporenko, Yutong Zhang, Sean Campbell, and Alessandro Roncone. Morphit: Flexible spherical approximation of robot morphology for representation-driven adaptation. arXiv preprint arXiv:2507.14061, 2025.

[66] Dong Hai Phuong Nguyen, Matej Hoffmann, Alessandro Roncone, Ugo Pattacini, and Giorgio Metta. Compact real-time avoidance on a humanoid robot for human-robot interaction. In Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction, pages 416–424, 2018.

[67] Gianpaolo Palma, Narges Pourjafarian, Jürgen Steimle, and Paolo Cignoni. Capacitive touch sensing on general 3d surfaces. ACM Trans. Graph., 43(4):103–1, 2024.

[68] Radhen Patel and Nikolaus Correll. Integrated force and distance sensing using elastomer-embedded commodity proximity sensors. In Robotics: Science and systems, 2016.

[69] Venkatesh Pattabiraman, Zizhou Huang, Daniele Panozzo, Denis Zorin, Lerrel Pinto, and Raunaq Bhirangi. eflesh: Highly customizable magnetic touch sensing using cut-cell microstructures. arXiv preprint arXiv:2506.09994, 2025.

[70] Daniel Rakita, Bilge Mutlu, and Michael Gleicher. Relaxedik: Real-time synthesis of accurate and feasible robot arm motion. In Robotics: Science and Systems, volume 14, pages 26–30. Pittsburgh, PA, 2018.

[71] Daniel Rakita, Haochen Shi, Bilge Mutlu, and Michael Gleicher. Collisionik: A per-instant pose optimization method for generating robot motions with environment collision avoidance. In 2021 IEEE International Conference on Robotics and Automation (ICRA), pages 9995–10001. IEEE, 2021.

[72] Alessandro Roncone, Matej Hoffmann, Ugo Pattacini, Luciano Fadiga, and Giorgio Metta. Peripersonal space and margin of safety around the body: learning visuo-tactile associations in a humanoid robot with artificial skin. PloS one, 11(10):e0163713, 2016.

[73] Alessandro Roncone, Matej Hoffmann, Ugo Pattacini, and Giorgio Metta. Learning peripersonal space representation through artificial skin for avoidance and reaching with whole body

surface. In 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 3366–3373. IEEE, 2015.

[74] Siddharth Rupavatharam, Caleb Escobedo, Daewon Lee, Colin Prepscius, Larry Jackel, Richard Howard, and Volkan Isler. Sonicfinger: Pre-touch and contact detection tactile sensor for reactive pregrasping. In 2023 IEEE International Conference on Robotics and Automation (ICRA), pages 12556–12562. IEEE, 2023.

[75] Siddharth Rupavatharam, Xiaoran Fan, Caleb Escobedo, Daewon Lee, Larry Jackel, Richard Howard, Colin Prepscius, Daniel Lee, and Volkan Isler. Ambisense: Acoustic field based blindspot-free proximity detection and bearing estimation. In 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 5974–5981. IEEE, 2023.

[76] Siddharth Rupavatharam, Richard Howard, Daewon Lee, Lawrence Jackel, Caleb Sebastian Escobedo, and Ibrahim Volkan Isler. Sonicfinger: Low-cost, compact, proximity and contact sensor for reactive positioning, December 2025. U.S. Patent US 12,496,723 B2, assigned to Samsung Electronics Co., Ltd.

[77] Alexander Schmitz, Perla Maiolino, Marco Maggiali, Lorenzo Natale, Giorgio Cannata, and Giorgio Metta. Methods and technologies for the implementation of large-scale robot tactile sensors. IEEE Transactions on Robotics, 27(3):389–400, 2011.

[78] Linda-Sophie Schneider, Junyan Peng, and Andreas Maier. Robot movement planning for obstacle avoidance using reinforcement learning. Scientific Reports, 15(1):32506, 2025.

[79] Carter Sifferman, Mohit Gupta, and Michael Gleicher. Efficient detection of objects near a robot manipulator via miniature time-of-flight sensors, 2025.

[80] Carter Sifferman, Yiquan Li, Yiming Li, Fangzhou Mu, Michael Gleicher, Mohit Gupta, and Yin Li. Recovering parametric scenes from very few time-of-flight pixels, 2025.

[81] Carter Sifferman, William Sun, Mohit Gupta, and Michael Gleicher. Using a distance sensor to detect deviations in a planar surface, 2024.

[82] Anna Soukhovei, Carson Kohlbrenner, Caleb Escobedo, Alexander Gholmieh, Alexander Dickhans, and Alessandro Roncone. Form-fitting, large-area sensor mounting for obstacle detection. arXiv preprint arXiv:2603.23725, 2026.

[83] Matthew Strong, Caleb Escobedo, and Alessandro Roncone. Volumetric data fusion of external depth and onboard proximity data for occluded space reduction. arXiv preprint arXiv:2110.11512, 2021.

[84] Marc Teyssier, Brice Parilusyan, Anne Roudaut, and Jürgen Steimle. Human-like artificial skin sensor for physical human-robot interaction. In 2021 IEEE International Conference on Robotics and Automation (ICRA), pages 3626–3633. IEEE, 2021.

[85] D. Um, B. Stankovic, K. Giles, T. Hammond, and V. Lumelsky. A modularized sensitive skin for motion planning in uncertain environments. In Proceedings. 1998 IEEE International Conference on Robotics and Automation (Cat. No.98CH36146), volume 1, pages 7–12 vol.1, 1998.

[86] Lac Van Duong et al. Large-scale vision-based tactile sensing for robot links: Design, modeling, and evaluation. IEEE Transactions on Robotics, 37(2):390–403, 2020.

[87] Charles Wade, Devon Beck, and Robert MacCurdy. Implicit modeling for 3d-printed multi-material computational object design via python. Proceedings of the 10th ACM Symposium on Computational Fabrication, 2025.

[88] Charles Wade, Devon Beck, and Robert MacCurdy. Implicit toolpath generation for functionally graded additive manufacturing via gradient-aware slicing. Additive Manufacturing, 2025.

[89] Charles Wade, Graham Williams, Sean Connelly, Braden Kopec, and Robert MacCurdy. Openvcad: An open source volumetric multi-material geometry compiler. Additive Manufacturing, 79:103912, 2024.

[90] Yeping Wang, Pragathi Praveena, Daniel Rakita, and Michael Gleicher. Rangedik: An optimization-based robot motion generation method for ranged-goal tasks. In 2023 IEEE International Conference on Robotics and Automation (ICRA), pages 9700–9706. IEEE, 2023.

[91] Kandai Watanabe, Matthew Strong, Mary West, Caleb Escobedo, Ander Aramburu, Krishna Chaitanya Kodur, and Alessandro Roncone. Self-contained kinematic calibration of a novel whole-body artificial skin for human-robot collaboration. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 1778–1785. IEEE, 2021.

[92] Daniel G. Wegerif and Douglas J. Rosinski. Sensor-based whole-arm obstacle avoidance for kinematically redundant robots. In Paul S. Schenker, editor, Sensor Fusion V, volume 1828, pages 417 – 426. International Society for Optics and Photonics, SPIE, 1992.

[93] Yuyang Wei, Bingqian Li, Marco Domingos, Zhihui Qian, Yiming Zhu, Lingyun Yan, Lei Ren, and Guowu Wei. Fully 3d printed flexible, conformal and multi-directional tactile sensor with integrated biomimetic and auxetic structure. Communications Engineering, 2(1):80, 2023.

[94] Nicholas Wettels, Veronica J Santos, Roland S Johansson, and Gerald E Loeb. Biomimetic tactile sensor array. Advanced robotics, 22(8):829–849, 2008.

[95] Anders Frem Wolstrup, Jon Spangenberg, Akio Yamamoto, Andrew Gleadall, and Gabriel Zsurzsan. Advances in 3d printed electromechanical sensors: Performance comparison, trends, and future directions. Additive Manufacturing, 106:104799, 2025.

[96] Xiaomeng Xu, Dominik Bauer, and Shuran Song. RoboPanoptes: The All-Seeing Robot with Whole-body Dexterity. In Proceedings of Robotics: Science and Systems, LosAngeles, CA, USA, June 2025.

[97] Shunnosuke Yoshimura, Kento Kawaharazuka, and Kei Okada. M3d-skin: Multi-material 3d-printed tactile sensor with hierarchical infill structures for pressure sensing, 2025.

[98] Wenzhen Yuan, Siyuan Dong, and Edward H Adelson. Gelsight: High-resolution robot tactile sensors for estimating geometry and force. Sensors, 17(12):2762, 2017.

[99] Gabriel Zöller, Vincent Wall, and Oliver Brock. Active acoustic contact sensing for soft pneumatic actuators. In 2020 IEEE International Conference on Robotics and Automation (ICRA), pages 7966–7972. IEEE, 2020.