From Prosthetic Constraints to Better Robot Hands

Arav Bhargava

4/22/26

1. Why dexterous manipulation is as much a data problem as a hand-design problem

I spent the flight home from San Francisco thinking through my conversations from my trip to the bay and how the work I have done in prosthetics maps onto robotics. This is my attempt to distill the angle I find most compelling: dexterous hands are not just a hardware problem, they are a data problem, and the most interesting opportunity may be building a hand-and-interface stack that captures the missing contact signal in human manipulation.

Over the last few years, I invented and patented a low-cost, size-adjustable prosthetic socket for low-resource settings. I worked directly with amputees, built EMG-controlled bionic hands actuated through tendons, worked with Morph Labs on neuroprostheses, and interned at Walter Reed, where I observed highly dexterous military prostheses and the broader challenge of trying to restore real hand function. Across all of it, I kept seeing the same thing: dexterity is not just a motor-count problem. It is an anatomy problem, a tendon-routing problem, a sensing problem, a feedback problem, and above all, a human-interface problem.

That is why prosthetics feels like a useful lens here. The real question is how to turn those lessons into better robot hands and, even more importantly, better data.

What feels increasingly clear to me is that neither of the two obvious paths is enough on its own. **Pure teleoperation is too inefficient for collecting the scale and diversity of data needed for truly general dexterous manipulation. At the same time, world models and robot foundation models are clearly improving, but they are still far from solving dexterity, especially in the contact-rich parts of manipulation, where small errors in touch, slip, and force correction matter.** Universal Manipulation Interface (UMI) matters precisely because it shows that portable, low-cost, information-rich human interfaces can produce useful robot-teaching data without requiring in-the-wild robots for collection. This can be extended in different ways, as I will discuss below.

My central belief is that the next important step in dexterous humanoid manipulation is not just building a more complex hand or just collecting more vision-based human demonstration data as many humanoid companies are currently attempting to do. It is building a hand and a human-interface system that can capture the missing contact signal in manipulation and capitalize on the ways humanlike five finger robotic hands can be trained based of our own hands as humans. The hand should be humanlike enough that rich human behaviors transfer cleanly, simple enough that it does not become overengineered (where the new data types just complicates things), and give enough rich data at the fingertips and wrist that it turns human behavior into better policy training data. In that framing, the goal is not to maximize "anatomical realism". The goal is to preserve the dexterous capabilities that are most significant, especially precision grasping and understanding how contact changes during grasping.

Ultimately, I believe that this would allow a method similar to UMI using gloves, tactile sensing, and perhaps even sonomyography (use of ultrasound to detect muscle movements) to build a UMI device that can gather much better data more efficiently.

2. Why use prostheses as a lens for robot arms

Prostheses matter because they force honesty about what a hand actually needs to be useful. A prosthetic hand cannot succeed because it looks good in a rendering or because it performs a few clean grasps in a demo. It has to be wearable, trustworthy, intuitive, and useful under repeated real interactions. That lesson transfers directly to robotics. A hand with many joints but poor contact understanding, brittle routing, high distal inertia, or weak data interfaces is still not a very good hand.

At the same time, I do not think this should become a prosthetics-centered argument. Prosthetics is useful as an influence, but, as I have found in my work, it is a slower industry with different incentives and constraints. Therefore, all of the high-tech tactile sensing research within prostheses is unfortunately years away from commercialization.

Robotics should borrow what is useful from it, not inherit its assumptions at face value. The prosthetics lesson that feels most important here is not "copy the human hand exactly." Instead, it is that humanlike design is valuable mainly because it makes human data easier to transfer. A robot hand can likely be simpler than a biological hand and still be extremely useful, as long as it preserves the right contact and grasp behaviors.

3. What the frontier is actually doing

The strongest signal from current robotics is that hands are becoming more perception-integrated, more tactile, and more explicitly tied to learning. Looking through each company's approach, we have:

Figure

In its Figure 03 announcement, the company said each hand includes an embedded palm camera with wide field of view and low-latency sensing, and described the hand system as redesigned around Helix. Figure also says it moved toward softer, more adaptive fingertips after concluding current tactile sensors on the market had real durability limitations for real-world use. In Helix 02, Figure described a full-body autonomy system that integrates vision, touch, and proprioception and highlighted palm cameras and fingertip tactile sensors for in-hand visual feedback and force-modulated grasping.

Unitree

Unitree has its Dex5-1 that is presented as a 20-DoF hand with 94 tactile sensors, replaceable fingers, and smoother backdrivability intended to make operation easier for reinforcement-learning training.

Sanctuary AI

Sanctuary publicly says its new tactile sensors are meant to improve dexterous manipulation and enable blind picking, slippage detection, and prevention of excessive force application when vision is occluded, and it explicitly ties tactile sensing to richer training data for embodied AI.

Tesla

Tesla has another unique approach, building and manufacturing their own hand. Optimus is one of the clearest examples of the belief that humanoid value will eventually be bottlenecked by manipulation and manufacturability.

Sunday

Sunday is especially interesting because it is not just building a robot hand, but a hand-centered data pipeline. Sunday uses a UMI glove. The company says its Skill Capture Glove and robot hand were co-designed to share the same geometry and sensor layout, with the goal of reducing the translation gap between human demonstrations and robot data. That makes Sunday one of the clearest examples of the idea that hand design and data collection should be treated as one problem, not two separate ones.

4. Why this is really a data problem

The strongest conclusion I came away with is that the most important thing to optimize first is not the physical hand itself, **it is the ability of the hand to collect rich human demonstration data.** UMI already proved that portable human interfaces can produce low-cost, information-rich data useful for transfer. But I also think the real gap is not generic human demonstration data. It is contact-rich human demonstration data.

Vision is clearly valuable, but computer vision alone is probably not enough for the parts of dexterity that depend most on real touch: initial contact, increasing pressure, incipient slip, unstable grasp, fine force corrections, and what the fingers are actually doing once the object is partly occluded. The point is not that vision is bad. The point is that vision is best at gross motion and scene context, while the most important missing signal in manipulation begins right when contact happens.

I was thinking about the recent developments in prostheses, trying to give feeling and stimulation to the amputee. Ultimately, if a robot is going to learn from people doing things with their hands, then we need to capture not only what the hand looked like, but what it actually felt and how contact changed over time. That does not necessarily mean a giant, fully instrumented glove covering every surface. In fact, I suspect the better research path is the opposite: identify the minimum extra sensing that gives a real jump over vision-only or standard teleop pipelines. My current guess is that the first meaningful upgrade is accurate hand pose, accurate wrist pose, and fingertip contact sensing. Fingertips are where the most information-dense contact transitions happen. Palm sensing and richer full-hand sensing may matter later, but fingertip contact plus hand and wrist pose feels like the highest-leverage first stack.

5. Tactile sensing: what is actually worth adding first

Tactile sensing gets discussed too vaguely most of the time. “Add touch” is not a useful design principle. The real question is what tactile information is worth collecting first, and at what cost in complexity.

At a high level, tactile sensing matters because it captures information vision cannot reliably recover through remote observation alone. Cameras can often estimate hand pose, object pose, and scene layout. They are much worse at telling you exactly when contact begins, how force is shifting across a fingertip, when an object starts to slip, or how grasp stability changes over time. Those are precisely the signals that matter in dexterous manipulation. Penn State’s recent work on an ultra-thin, flexible pressure sensor for robotic touch is relevant here because it shows the field moving toward packaging pressure sensing into realistic, skin-like form factors rather than treating touch as a bulky lab-only add-on.

There are several main tactile sensor directions. Capacitive sensing can be thin and sensitive but is more vulnerable to noise. Piezoresistive or resistive sensing is simpler and easier to integrate, though more vulnerable to drift and hysteresis. Optical tactile sensing is richer and better for contact geometry and slip-related signals, but bulkier and harder to package. Magnetic and hybrid approaches can be expressive too, but are often harder to calibrate robustly.

If I had to suggest a concrete first sensing direction, I would start with resistive or piezoresistive fingertip sensors, ideally in a small array, rather than jumping immediately to optical tactile fingertips or full-hand skin. For the first version of this project, the goal would not be to perfectly reconstruct contact geometry. It would likely be to capture the most important contact transitions in a way that is thin, durable, easy to integrate, and scalable to both a robot hand and a human data-collection interface. Resistive and piezoresistive sensing seem like the best first candidate because they can fit into compact fingertips and are well-suited for detecting the signals that matter most here: contact onset, changing pressure, unstable grasp, and possibly early slip when combined with good temporal processing. This likely won’t be perfect but may be an acceptable tradeoff if the first goal is simply to identify the minimum tactile signal that materially improves policy learning.

I would also not start with optical tactile sensing as the default first choice, not because it is unimportant, but because it may be too bulky and integration-heavy for the first version of a hand whose purpose is to scale human demonstration data. So the cleanest first proposal is to start with resistive or piezoresistive fingertip sensing for contact-state transitions, align it with accurate hand and wrist pose, and then use richer optical tactile sensing later as a deliberate comparison. That gives the project a clear progression: first prove that a minimal fingertip contact signal improves human demonstration data, then test whether a richer but more complex tactile stack adds enough value to justify the extra hardware burden.

6. The hand should be humanlike enough, not maximally humanlike

A second important conclusion is that more humanlike design is mainly valuable because it makes human data easier to transfer. I do not think the robot hand has to perfectly replicate the human hand to be valuable. It may be fine to simplify the kinematics or constrain the tendon layout, as long as the system still captures the dexterous primitives that matter most.

That pushes the hand design toward a middle ground. It should be simple enough to scale and robust enough to survive use, but still preserve useful thumb function, contact at the fingertips, controllable precision grasps, and enough wrist function to position and orient the hand naturally. The first target is not full human-hand capability. It is reliable precision grasping with meaningful contact awareness.

7. Human-hand interfaces already exist, but they are not yet the whole answer

What makes this especially interesting is that parts of the interface stack already exist. MANUS is explicitly positioning its gloves for teleoperation and humanoid robot training, emphasizing full-finger capture, low latency, and robot integration. That matters because it means the basic idea of using the human hand as the interface is already validated.

But the gap is that most of these systems are still primarily about kinematics, control, or teleoperation, not about collecting the richest possible contact signal from the human hand itself. That is why I think there is room for something like a hand-centered UMI. Instead of treating the glove only as a teleoperation device, it should be treated as a data instrument: something that captures hand pose, wrist pose, and fingertip contact in a form that can be retargeted onto a robot hand and used directly for learning.

There is also a natural extension beyond glove kinematics. MIT's recent ultrasound wristband work is especially interesting because it suggests a way to recover hand motion and possibly pre-contact intent from inside the wrist rather than only from external observation. MIT says the band tracks muscles, tendons, and ligaments in real time and wirelessly controls a robotic hand. That makes me think EMG and ultrasound are both worth mentioning, but as extension paths, not version-one requirements. My current view is that version one should be hand pose, wrist pose, and fingertip contact. Version two, if justified, could add EMG or sonomyography or ultrasound. The value of internal physiological sensing would be distinguishing similar-looking motions driven by different intent, capturing grasp preparation earlier, and studying how muscle activation changes across grasp strategies. But I would not make that the initial dependency.

I have seen this sonomyography approach to prosthesis control, and it ended up being highly accurate and much better than the EMG approach. That being said, it is important to make sure any "UMI" is easy to replicate and bring to many people to collect as much human hand data as possible from doing the specific tasks.

8. Project proposals

Project 1: A modular hand platform for testing tactile sensing, actuation, and wrist setups

This project is closest to the design direction I discussed in the Bay Area: a hand platform that is not just dexterous, but experimentally useful. The central idea is to build a hand that is humanlike enough to preserve the important structure of human demonstrations, but still simple enough to be scalable and mechanically realistic for a humanoid platform.

The design goals would be a tendon-driven architecture, controllable thumb opposition, fingertip-dominant contact, enough wrist freedom to position the hand naturally, and easy changeout of the subsystems we actually care to compare. The key is that the platform should make it easy to test different tactile sensor approaches, different actuation strategies, and different wrist architectures without rebuilding the whole hand every time.

I would use a staged integration path. Mechanically, I would build the hand structure, tendon routing, and wrist around a stable base geometry, keep fingertip modules swappable, and make the wrist and tendon-routing architecture configurable enough to test different setups without redesigning the palm every iteration. On the sensing side, I would start with off-the-shelf or easily manufacturable resistive or piezoresistive fingertip sensing, mount the sensors in modular fingertip cartridges behind compliant silicone, and route signals back through the finger to local PCBs or hand-level electronics. On actuation, I would first use a practical tendon-driven setup with readily available actuators and encoders, then leave room to compare compliance, backdrivability, and transmission choices before moving toward more custom actuation modules. In software, I would timestamp and synchronize fingertip sensor data, joint states, wrist state, cameras, and any glove-retargeted demonstrations, then convert raw tactile signals into usable features like contact onset, pressure change, contact duration, multi-finger contact pattern, and possible slip proxies.

The point of this project is not just to build a good hand. It is to build a hand that makes it possible to ask better questions: which embodiment choices actually matter for transfer, which sensor stacks are worth their complexity, and which actuation and wrist tradeoffs materially improve the usefulness of human-derived data.

Project 2: A hand-centered UMI for richer human demonstration data

The second project is a wearable interface inspired by UMI, but more explicitly centered on the actual human hand. The core idea is the following: today, a lot of hand data is effectively motion-first. I think there is room for a system that captures not just what the human hand looked like, but what it was actually doing at contact.

The best first version should probably capture hand pose, wrist pose, and fingertip contact information. That is focused enough to be realistic and buildable, but rich enough to test the real thesis of this memo.

I would also stage this project. In phase one, I would use existing pose-tracking infrastructure wherever possible, start from an existing glove or glove-like hand tracking setup, add custom fingertip contact patches, and synchronize the pose and contact streams in software. That avoids reinventing the whole glove before proving the core data hypothesis. In phase two, I would move to a more integrated textile or flexible-substrate design with embedded fingertip sensing, better wrist tracking, and improved comfort, calibration, and durability. In phase three, I would add EMG or sonomyography or ultrasound only if it becomes clear that intent separation matters enough to justify the added complexity.

The glove should function primarily as a data instrument, not just a teleoperation device. The point is not simply to let a user control a robot hand more naturally, though that may be a useful byproduct. The point is to create a pipeline for collecting abundant human demonstrations that are more informative for policy learning because they preserve the missing contact signal.

This project would test a very specific idea: the next jump in manipulation learning may come from contact-rich human data, not just more trajectories. If that turns out to be true, then the argument for more humanlike robot hands becomes much stronger, because their value is no longer mainly aesthetic or anatomical. Their value is that they make high-quality human data easier to capture, align, and transfer.

9. How these two projects fit together

These two projects are strongest when treated as one system. The hand by itself is not enough, because the whole point is to improve the usefulness of human data. The glove by itself is also not enough, because it needs a robot hand whose structure is similar enough for the data to transfer meaningfully. Together, they create a closed loop: a robot hand built to preserve the most important dexterous primitives, and a human interface built to capture the missing contact signal those primitives depend on.

That is why I think this is a more compelling direction than simply proposing a better hand or a better glove in isolation.

10. Why I think this matters

I am bullish on humanoids, but I think the most important part of that story is not just locomotion or even general reasoning. It is whether robots can become truly useful in the messy, contact-rich world. That means the hand matters enormously, as well as being able to collect abundant data efficiently.

But I think the lesson from the last few years is that the hand should not be understood only as a hardware challenge. It is a data challenge, a transfer challenge, and a question of how to preserve the most important structure of human manipulation without making the robot so anatomically faithful and mechanically complicated that it becomes hard to scale.

The opportunity is not only to build a better hand. It is to build a hand and a demonstration interface that can finally capture the missing contact signal in human manipulation and turn it into better robot learning. If that works, then more humanlike hand design becomes strategically useful not because the world morally demands five fingers, but because human data becomes more abundant, more transferable, and more informative.

That is the path I find most compelling.

References (Added anything I found helpful)

Chi, Cheng, et al. “Universal Manipulation Interface: In-The-Wild Robot Teaching Without In-The-Wild Robots.” *Robotics: Science and Systems*, 2024.

https://arxiv.org/abs/2402.10329

Figure. “Introducing Figure 03.” Figure AI, October 9, 2025.

https://www.figure.ai/news/introducing-figure-03

Figure. “Introducing Helix 02: Full-Body Autonomy.” Figure AI, 2026.

https://www.figure.ai/news/helix-02

Unitree Robotics. “Dex5-1.” Unitree Robotics.

https://www.unitree.com/mobile/Dex5-1

Unitree Robotics. “Dex3-1.” Unitree Robotics.

https://www.unitree.com/mobile/Dex3-1

Sanctuary AI. “Sanctuary AI New Tactile Sensors Enable Richer Sense of Touch.” Sanctuary AI.

https://www.sanctuary.ai/blog/sanctuary-ai-new-tactile-sensors-enable-richer-sense-of-touch

Sanctuary AI. “Sanctuary AI Equips General Purpose Robots with New Tactile Sensors.” Sanctuary AI.

https://www.sanctuary.ai/blog/sanctuary-ai-equips-general-purpose-robots

MANUS. “Data Gloves for Teleoperation & Humanoid Robot Training.” MANUS.

https://www.manus-meta.com/robotics

MANUS. “Mercedes-Benz Deploys Apptronik’s Apollo at Factory Floor Using Manus Gloves.” MANUS.

https://www.manus-meta.com/use-cases/mercedes-benz-deploys-apptroniks-apollo-at-factory-floor-using-manus-gloves

MIT News. “Wristband Enables Wearers to Control a Robotic Hand with Their Own Movements.” MIT News, March 25, 2026.

https://news.mit.edu/2026/wristband-enables-wearers-control-robotic-hand-with-own-movements-0325

Penn State University. “Q&A: Robots Can’t Feel; These Sensors Could Change That.” Penn State News, 2026.

https://www.psu.edu/news/research/story/qa-robots-cant-feel-these-sensors-could-change

Penn State Engineering. “Ultra-Thin Pressure Sensors Could Help Robots Feel.” Penn State College of Engineering News, 2026.

https://news.engr.psu.edu/2026/cheng-larry-pressure-sensors-robots-feeling.aspx

Inspire-Robots Store. “The Dexterous Hands.” Inspire-Robots.

https://inspire-robots.store/collections/the-dexterous-hands

Reuters. “China Humanoid Robot Half-Marathon Showcases Technical Leaps.” Reuters, April 18, 2026.

https://www.reuters.com/world/asia-pacific/china-humanoid-robot-half-marathon-showcase-technical-leaps-2026-04-18/

Reuters. “Tesla’s Cybercab, Optimus Output to Start ‘Agonizingly Slow’ Ramp-Up Later, Musk Says.” Reuters, January 21, 2026.

https://www.reuters.com/business/autos-transportation/teslas-cybercab-optimus-output-start-agonizingly-slow-ramp-up-later-musk-says-2026-01-21/