



Whitepaper

Edge AI in Robotics

How Foundation Models for Physical Interaction Are Enabling Tactile Intelligence

October 2026

Physical AI



Executive Summary

The Frontier of Physical AI

Robotics is entering a shift. For decades, industrial automation relied on pre-programmed, rigid trajectories. These systems performed exceptionally well in sterile, structured environments where target parts occupied exact, unvarying coordinates. But today's commercial landscapes demand more-adaptable machinery. The next major paradigm shift in industrial and commercial robotics centers on Physical AI and dexterous manipulation.

A major technical challenge facing humanoid hands and advanced robotic grippers is giving robots physical interactivity. Robotic hands must adapt instantly to objects of unknown shapes, varying surface frictions, and delicate structures. If a robot gripper applies too much force, it destroys the item. If it applies too little force, the item slips and falls. Achieving the precise balance requires a continuous stream of high-frequency feedback that allows the system to adjust its physical behavior dynamically. This whitepaper details a joint solution that bridges the physical intelligence gap: the integration of the open Synaptics Astra Edge compute processor platform with TacnIQ.ai's pre-trained tactile AI models. By running lightweight, physical neural networks locally on specialized Edge silicon, we establish a production-ready Edge platform for dexterity functions. This platform enables robotic systems to react to tactile inputs with millisecond latency, bringing tactile reflexes to the next generation of dexterous robots. It represents a breakthrough that turns complex, unreliable tactile sensing into a predictable, deployable, and commercial reality.

Why Tactile Sensing is Challenging but Essential

Historically, the robotics industry attempted to solve physical manipulation almost entirely through computer vision. Vision algorithms have advanced rapidly, yet they face severe structural limitations during physical contact. Vision models struggle with the physical reality of occlusion: the moment a robotic hand approaches an object to grasp it, the hand itself blocks the camera's line of sight. The camera can't see the exact points of contact, making it impossible to detect micro-slips or localized surface deformations.

Pure vision systems also suffer from high false-positive rates due to changing lighting conditions, shadows, and reflections. Most importantly, cameras cannot measure physical forces. A camera can identify the shape of a container, but it can't evaluate how heavy it is, how slippery the surface feels, or how much internal pressure the object can withstand. Robots fail because they cannot feel what they touch, even when they possess perfect vision.

The Severe Labeled Data Deficit

Why has tactile sensing remained stuck while computer vision has achieved massive production scale? The answer lies in a severe labeled data deficit: the zero percent wild data bottleneck. Modern computer vision relies on transfer learning, built on massive open datasets like ImageNet, which contains over fourteen million hand-labeled images. Every commercial vision model today inherits that structural database, allowing developers to fine-tune pre-trained vision models for specific tasks with minimal training data.

Tactile AI, on the other hand, possesses no such equivalent. There are zero large-scale, commercially usable, in-the-wild tactile datasets. Tactile sensing data remains confined to isolated research laboratories, captured using highly custom, non-standardized sensor setups under strictly controlled conditions. Because of this, every robotics lab must attempt to build its tactile algorithms nearly from scratch, a process that increases the likelihood of failing when the system encounters unpredictable, messy real-world environments.

Signal Complexity and Prolonged Development Timelines

The physics of touch are incredibly complex. Tactile sensors generate high-frequency, exceptionally noisy, and uncharacterized raw data streams. A single grasp involves infinite physical variables:


- normal force
- lateral force
- sliding friction
- thermal conductivity
- microscopic vibrations

Translating these raw electrical waveforms into clear, actionable physical states (like slip, or contact compliance) requires immense mathematical and software sophistication. Most development teams lack the specialized expertise in haptic signal processing and machine learning required to parse these signals.

Consequently, building a functioning tactile processing pipeline from zero increases traditional engineering schedules. Development teams typically spend 18 or more months on expensive cycles of raw data collection, custom sensor characterization, and manual AI model training. These elongated timelines introduce development risks, strain R&D budgets, and ultimately delay commercial product launches. Manufacturers and robotics integrators simply cannot afford to spend nearly two years proving out a basic sensory input loop before they can even begin working on their core robotic applications.


TACTILE INTELLIGENCE FOR ROBOTICS

A pre-trained model that adapts across tactile sensors with a short fine-tune.



FROM TOUCH TO ACTION

One compact on-device pipeline



FINE-TUNE, DON'T BUILD

~~18 months~~ > 4-6 weeks

Get to deployment faster

Source



Why Tactile Reactions Belong at the Edge

To design an efficient sensory architecture for a robot, it's beneficial to look to biology. Human physiology divides cognitive tasks into centralized and localized functions: the subconscious versus core brain division. When you interact with the physical world, your conscious brain (the core brain) handles high-level cognitive planning: identifying the object, calculating its coordinates, and deciding to pick it up. The fine-grained physical adjustments that prevent you from dropping the object run as a local, reflexive subconscious reaction.

For example, when a glass of water starts to slip from your fingertips, you don't wait for the sensory signal to travel to your conscious brain, undergo cognitive evaluation, and return a command to tighten your grip. If you did, the glass would smash on the floor. Instead, your peripheral nervous system triggers an instant, localized reflex that tightens your fingers in milliseconds.

Robotic architectures must mirror this biological reality. High-level path planning, semantic object classification, and general task sequencing belong in the robot's central processor, running the main robot operating system. But the critical, real-time physical adjustments during a grasp must run as a subconscious reaction at the far Edge: directly within the gripper or fingertip assembly itself. If a captured object begins to shift or rotate, the control system must adjust its gripper motor forces instantly. The grasp loop simply cannot wait for a round-trip communications link to the central computer, nor can it risk getting blocked by other processes running on the main operating system.



The Necessity of Millisecond-Level Edge Inference

Achieving a reliable haptic reflex requires millisecond-level Edge inference. Processing tactile signals locally at the far Edge protects the critical feedback loop from latency spikes, network jitters, and operating system overheads. If a slip reaction is delayed by even a few milliseconds, the physical state of the object changes irreversibly. The object slips out of reach, rotates into an ungraspable orientation, or falls. Processing haptic data locally at the fingertips ensures deterministic execution, allowing the gripper to respond with millisecond-level latency. This localized Edge architecture makes the physical contact loop fully autonomous, robust, and completely independent of any software lag or hardware bottlenecks occurring elsewhere in the robotic system.

TacnIQ.ai's Pre-Trained Physical Foundational Model

To solve the tactile data deficit, TacnIQ.ai applies a novel cognitive thesis: decoupled pre-training. Just as computer vision models learn the general substrate of vision (identifying Edges, colors, shapes, and textures) on ImageNet before they specialize in identifying highly specific objects, TacnIQ.ai teaches machines the fundamental substrate of feeling before they attempt to handle specific target items. You don't need a custom dataset of micro-connectors to train a tactile model. By exposing a model to a massive variety of physical touch waveforms, the neural network learns to generalize the basic laws of physical contact: friction, compliance, and movement. Once the model understands this baseline haptic substrate, specializing it for a specific industrial task becomes a simple, fast fine-tuning process.

Model Optimization for Low-Power Microcontrollers

To deploy this intelligence at the far Edge, TacnIQ.ai developed TIQ-1: a distilled, lightweight physical foundation model. While traditional foundation models are massive and require power-hungry cloud GPUs, TIQ-1 is specifically optimized to run with millisecond-level latency on low-power, resource-constrained embedded microcontrollers and Edge SoCs. The model delivers high accuracy while operating within tight memory and power footprints, allowing developers to embed tactile intelligence directly into robotic end effectors.

The Structured Three-Stage Learning Methodology

The TIQ-1 model generalizes touch across different physical environments and sensor modalities using a highly structured three-stage pipeline:


- **Intuitive Exploration:** This stage represents the “baby’s first touch” phase. Just as a baby learns about the world by touching everything in sight, TacnIQ.ai pre-trained its foundation model on a massive, wild data moat. This proprietary dataset contains over ten million unique contacts and six thousand hours of real-world touch waveforms. TacnIQ.ai captured this data across diverse commercial environments using wearable sensing nodes called Backy™. The dataset spans logistics, construction, e-commerce, hospitality, and healthcare.
- **Decoding Physics:** The second stage translates raw, unstructured tactile waveforms into clearly labeled, ground-truth physical primitives and micro-states. The model decodes two core primitives: slip and compliance (softness versus hardness). It also applies disturbance rejection, filtering out incidental false positives that could otherwise trigger unnecessary force changes. Each primitive represents a fundamental mechanical interaction, providing the system with a multi-dimensional understanding of physical touch.
- **Application Layer:** In the final stage, developers use the pre-trained physics embeddings as a standardized haptic SDK. They fine-tune the model for specific downstream tasks, such as pick-and-place assembly, parcel sorting, or general dexterous humanoid manipulation. This structured pipeline turns open-ended, complex AI training into a predictable, plug-and-play development experience.

The Synaptics Astra Platform: An Accessible Compute Baseline

A highly optimized tactile model like TIQ-1 requires an equally efficient, developer-friendly Edge compute platform to reach its potential. The Synaptics Astra system on chip (SoC) provides this ideal compute baseline. Astra represents a welcome departure from traditional embedded processors. It features an open, highly scalable, and standards-based neural processing unit (NPU) architecture. The Synaptics Torq™ AI engine integrates Google Research's Coral NPU into a tiered, heterogeneous subsystem designed for low-power, always-on Edge AI workloads. This developer-friendly design avoids the proprietary black boxes, restricted assembly languages, and restrictive non-disclosure agreements that have long plagued the semiconductor industry and slowed down evaluation.


To simplify integration, Synaptics delivers a suite of open-source AI tooling designed to accelerate model evaluation, streamline on-device compiler optimization, and prevent industry fragmentation. Developers import standard neural network models, compile them for the Astra NPU, and profile their execution speeds in real time. This open ecosystem gives software engineers and data scientists complete visibility and control over how their models run on the hardware.

This hardware-software synergy is critical for localized far-Edge AI. By interfacing Synaptics' reference capacitive sensing modules (such as the SN6012T capacitive force sensor with SPI output) directly with the Astra SoC, developers build an ultra-low-latency processing pipeline. The haptic signals travel directly from the sensor over SPI, undergo noise filtration on the processor, and feed immediately into the distilled TIQ-1 model running on the Torq NPU. This tight, on-device integration operates with extremely low-power consumption, allowing robotic hands and grippers to run continuous haptic reflexes without draining battery reserves or generating excessive thermal loads.



OPEN DEVELOPER TOOLCHAIN

Evaluate, compile, and profile standard neural network models for Astra.



Business and Operational Impact: Quantifying Efficiency

For manufacturers and robotics integrators, deploying a pre-trained foundation model on specialized Edge hardware delivers massive, quantifiable business advantages. The most significant metric is compressed time-to-market. After all, historically, building a custom tactile processing pipeline required up months of open-ended research. With TacnIQ.ai's pre-trained model running on the Astra platform, developers deploy a functioning pilot in just four to six weeks, and scale to full commercial deployment within four months. This collapse in development schedules reduces engineering overhead, protects capital, and allows companies to launch physical AI products far ahead of competitors.

Unprecedented Resource and Training Efficiencies


This speed is driven by the massive resource efficiency of the TIQ-1 model. Because the model already possesses a rich, pre-trained understanding of touch physics, it requires 95% less labeled training data to reach full accuracy on downstream tasks. The model also converges twenty times faster than models trained from scratch. In practice, this means developers don't need to spend months collecting and labeling tens of thousands of custom grasp waveforms. A small fraction of task-specific training data is all it takes to achieve production-grade accuracy, dramatically reducing data acquisition costs and technical risks.

Sensor-Agnostic Abstraction and Investment Protection

Another critical operational challenge in robotics is the sheer fragmentation of tactile sensor hardware. Every sensor modality possesses a unique physical bottleneck:

- Capacitive sensors struggle with parasitic noise
- MEMS sensors are fragile to mechanical shock
- piezoresistive sensors suffer from hysteresis and non-linear drift
- magnetic sensors suffer from cross-talk and electromagnetic interference

Traditionally, if an engineering team decided to change their sensor hardware mid-development, they had to discard their entire AI model, rewrite their firmware, and start data collection from scratch. TacnIQ.ai solves this fragmentation with its unified, sensor-agnostic software layer. The model's physics embeddings transfer training near-zero-shot across capacitive, MEMS, magnetic, and piezoresistive sensors. TacnIQ.ai acts as a unified AI substrate, abstracting the underlying physical hardware. This sensor-agnostic design provides investment protection. Developers swap hardware sensors at any time to meet new cost or packaging demands without losing their accumulated software development and training data. It gives companies the freedom to choose the best sensor for the job while maintaining a consistent intelligence layer.



Reference Implementation: Closed-Loop Gripper Implementation

To prove this technology is ready for production, Synaptics and TacnIQ.ai developed a closed-loop reference implementation. The demonstration addresses high-mix, low-volume goods packing, which represents a classic, highly challenging warehouse and e-commerce distribution scenario. In this environment, a robotic gripper must process an endless stream of highly variable household and personal care items back-to-back. The system must handle a delicate, highly deformable toothpaste tube and a heavy, rigid one-kilogram shampoo bottle, each requiring a different grasping strategy. For further details, read the whitepaper, *Gaining the Upper Hand: A Blueprint for Edge Tactile Sensing and Predictive Robotic Dexterity*.

In traditional automation, a robot would struggle with this mix. A rigid, pre-programmed grip force would likely crush the soft toothpaste tube, spilling its contents, or fail to secure the heavy shampoo bottle, causing it to drop as soon as the arm accelerated. Solving this through computer vision alone is impossible because the camera cannot evaluate material stiffness or detect microscopic sliding movements between the gripper surface and the object.

Closed-Loop Edge Inference in Action

The reference implementation features a two-finger industrial gripper equipped with real-time capacitive tactile sensors that feed touch signals directly to the Astra processor. The distilled TIQ-1 model runs locally on the Astra NPU, creating a continuous, self-contained feedback loop. The entire grasping operation executes with seamless, real-time intelligence:

- 1. Initial Contact Detection:** As the gripper fingers close on an irregular, asymmetric or uneven surface or shape, the capacitive sensors transmit initial contact signals to the Astra processor. Within milliseconds, the model infers the compliance of the object, detecting the exact moment of contact without applying excessive force.
- 2. Minimal Force Grasp:** Recognizing the delicate, deformable nature of an object, the controller commands the motor to apply the absolute minimal grip strength required to lift the item safely. The item is transported and packed without being crushed.
- 3. Real-Time Slip Prediction:** When transitioning to a heavier object, the gripper initiates a standard contact grasp. As the robotic arm begins to lift, the capacitive sensors pick up high-frequency micro-slip waveforms. The TIQ-1 model running on the Astra NPU instantly infers that the object is in an unstable state.
- 4. Proportional Force Adjustment:** Within milliseconds, the Astra processor signals the gripper motor to proportionally and dynamically increase the grip force. The force increases just enough to stabilize the heavier item, preventing it from dropping. This entire process occurs locally at the far Edge without any lag or round trips to the central robot computer.

Conclusion and Future Roadmap

Tactile intelligence has historically required robotics developers to build sensor characterization, data collection, model training, and application logic largely from scratch. Pre-trained tactile models change that development model. Combined with local Edge compute, they provide a path toward placing low-latency physical intelligence directly within robotic grippers and end effectors.

The immediate opportunity is closed-loop control for industrial grippers, where contact compliance and slip information can improve manipulation of variable objects. As tactile models, sensing technologies, and embedded compute continue to mature, the same architecture can extend toward increasingly dexterous multi-finger robotic systems, adding physical states such as roll and spin that become relevant once a robot must reorient objects in-hand rather than simply grip and hold them.

Get started

To discover more about Astra, visit <https://www.synaptics.com/products/embedded-processors>.

To learn more about TacnIQ.ai's TIQ-1 model, contact TacnIQ.ai at tacniq.ai/contact.



Copyright

Copyright© 2026 Synaptics Incorporated. All rights reserved.

Trademarks

Synaptics, the Synaptics logo, AI Native Synaptics logo, Astra, and the Astra logo are trademarks or registered trademarks of Synaptics Incorporated in the United States and/or other countries. All other trademarks are the property of their respective owners.

Contact

Visit our website at www.synaptics.com to locate the Synaptics office nearest you.

Notice

Use of the materials may require a license of intellectual property from a third party or from Synaptics. This document conveys no express or implied licenses to any intellectual property rights belonging to Synaptics or any other party. Synaptics may, from time to time and at its sole option, update the information contained in this document without notice.

INFORMATION CONTAINED IN THIS DOCUMENT IS PROVIDED "AS-IS," AND SYNAPTICS HEREBY DISCLAIMS ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, AND ANY WARRANTIES OF NON-INFRINGEMENT OF ANY INTELLECTUAL PROPERTY RIGHTS. IN NO EVENT SHALL SYNAPTICS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES ARISING OUT OF OR IN CONNECTION WITH THE USE OF THE INFORMATION CONTAINED IN THIS DOCUMENT, HOWEVER CAUSED AND BASED ON ANY THEORY OF LIABILITY, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, AND EVEN IF SYNAPTICS WAS ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. IF A TRIBUNAL OF COMPETENT JURISDICTION DOES NOT PERMIT THE DISCLAIMER OF DIRECT DAMAGES OR ANY OTHER DAMAGES, SYNAPTICS' TOTAL CUMULATIVE LIABILITY TO ANY PARTY SHALL NOT EXCEED ONE HUNDRED U.S. DOLLARS.