

# Guaranteeing Agency in the Age of Neural Prosthetics

## Executive Summary

**Sector:** Healthcare & Neuro-Technology **Proposed Safeguard:** The "Neuro-Escrow Standard"

In the next decade, Artificial Intelligence will dissolve the functional barrier between "abled" and "disabled." By leveraging Generative Motor Policies - transformer-based models that translate noisy neural activity into fluid robotic movement -we may restore physical agency to the 50 million people worldwide living with paralysis or limb loss. This is not merely a medical treatment, but an economic and humanitarian revolution that would allow individuals with quadriplegia, ALS, and combat-related amputations to re-enter the workforce and navigate the world with autonomy previously thought impossible.

However, the current trajectory of proprietary biology poses a catastrophic risk: Obsolescence of the Self. As demonstrated by the collapse of Second Sight Medical Products in 2020, which left hundreds of blind patients with unsupported and failing retinal implants, tethering the human body to cloud-based, proprietary servers creates a single point of failure. In the case that a manufacturer fails or changes its business model, users may find their own limbs locked and their autonomy turned into a subscription service they can no longer afford.

To mitigate these risks, we must decouple the function of the device from the fate of the company. This essay proposes the Neuro-Escrow Standard, a regulatory framework that requires the source code and model weights for any essential neural prosthetic to be held in a public trust. Should the manufacturer cease support, this code automatically enters the public domain. This grants the user the permanent right to repair their own body.

## Guaranteeing Agency in the Age of Neural Prosthetics

The year is 2029. Julian, a 24-year-old software engineer, sits in a coffee shop in Cambridge drinking his coffee and coding. Four years ago, after a severe car accident, a spinal cord injury left him paralyzed from the neck down. Today, he is typing code at 80 words per minute and holding a hot espresso. His movements look indistinguishable from anyone else's. Julian is not biologically healed. He is "bridged." A high-bandwidth Brain-Computer Interface (BCI) implanted in his motor cortex streams raw neural spikes to an AI model running on his exoskeleton. The AI doesn't just listen, it anticipates, using a Generative Motor Policy to smooth his noisy intent into fluid, graceful motion. For Julian, the AI is not a tool; it is part of him. It is his hands, his legs, his freedom, and his hope.

Then, as he is getting ready to stand up, the notification appears on his retinal display.

"SERVER UNREACHABLE. LICENSE VALIDATION FAILED. MOTOR FUNCTIONS SUSPENDED."

The miracle vanishes. Julian's arms lock in place. The cup falls from his hand and shatters on the floor. He is paralyzed again. Not because of biology, but because the startup that owns his subscription just entered administration.

This scenario is not science fiction; it is the inevitable conclusion of our current trajectory in neuro-technology. As we witness the rapid strides occurring in the field of Brain-Computer Interfaces (BCI), the science of establishing direct communication pathways between the brain and external devices, the stakes are becoming tangible. We see this most clearly with Neuralink, the neurotechnology company pioneering high-bandwidth implantable chips that allow users to control digital devices through neural activity alone. As these proprietary AI models evolve to become the essential "operating system" for the human body, we face a profound ethical divergence. We are building a future where the disabled can walk, run, and work again, but only as tenants in their own bodies. To realize the true "Net Positive" of AI in healthcare, we must answer a question that has never existed in human history: When the software crashes, who owns the limb?

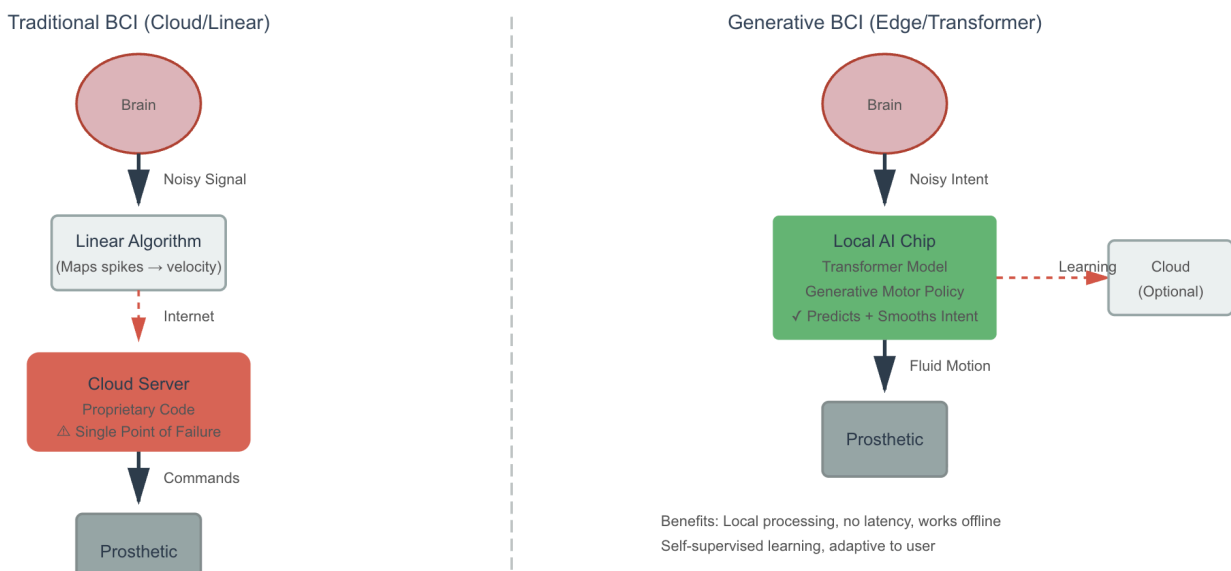
The healthcare sector stands to gain the most profound "Net Positive" impact from AI because it is the only domain where technology does not merely optimize systems but fundamentally redefines the limits of human biology. In this field, AI is unlocking potential that was once considered the realm of science fiction by granting "super-human" capabilities such as enhanced strength through robotic actuation, superior speed via active prosthetics, or direct telepathic control of digital interfaces. More importantly, it offers tangible restoration to those who have suffered catastrophic accidents, turning permanent disability into temporary injury. For decades, the field was stagnant. Prosthetics were "unintelligent mechanics" or passive devices that acted as dead weight. To operate a robotic arm, a user had to engage in myoelectric switching, manually toggling muscles in their chest or back to trigger simple open/close movements. This

process was cognitively exhausting, non-intuitive, and frustratingly slow, leading to a rejection rate of nearly 30% where patients simply abandoned their devices. AI eliminates this cognitive friction by replacing manual toggling with seamless, thought-driven intent.

Generative Motor Policies represent the specific architectural shift, illustrated in Figure 1, required to solve this cognitive friction. This technology represents the convergence of robotics and the Foundation Model architecture that powers Large Language Models (LLMs). Originating from breakthrough research in Vision-Language-Action (VLA) models and emerging World Models, these systems treat physical movement not as a series of equations, but as a language [1]. Just as an LLM predicts the next word in a sentence based on context, a Generative Motor Policy simulates the next millisecond of muscle contraction based on intent. By training on petabytes of “cross-embodiment” data—videos of humans moving, grasping, and walking—the model learns the fundamental “grammar” of the human body. We are at the inflection point for this technology today because of the Transformer architecture. Unlike previous neural networks (RNNs), which struggled with long sequences, Transformers can attend to the “physics of the cup,” the “geometry of the hand,” and the “noisy neural signal” simultaneously. When Julian thinks “grab,” the model doesn't just amplify the signal; it simulates the optimal trajectory to fulfill that intent, filling in the gaps to create fluid action from noisy thought. It effectively acts as a semantic translator, converting the vague intent of the mind into the precise syntax of the machine

To understand why this shift is imminent, we must distinguish between decoding the signal and inferring intent. Traditional BCIs rely on linear algorithms that map neural spikes directly to velocity. These systems are fragile. They require daily recalibration and struggle with the shifting nature of the brain. The next generation of AI introduces Self-Supervised Learning. Much like AI models predict the next word in a sentence, Large Motor Models predict the next movement in a sequence. This shift from command-following to intent-prediction is what makes the technology viable for widespread deployment within the decade.

**Figure 1: Traditional vs Generative BCI Architecture**

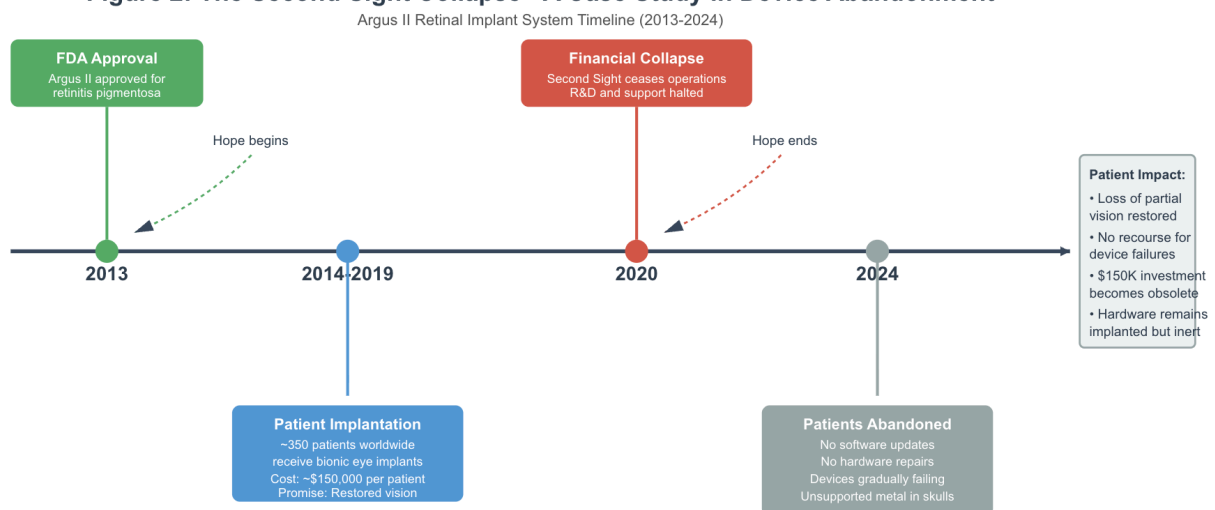


This technology’s impact extends beyond spinal cord injuries to reach nearly 2 million people in the U.S. living with limb loss, including thousands of veterans [2]. Traditional prosthetics are passive mechanics that essentially act as fancy springs. They cannot decide to run, jump, or stabilize on uneven terrain. AI changes this dynamic. A veteran with a smart amputation doesn't just get a plastic foot; they get a robotic limb that predicts terrain. Using the same transformer architecture described above, the leg can identify stairs or gravel via sensors and adjust its torque in milliseconds. This results in a return to active duty, competitive sports, or simply the ability to run and take care of their own child.

The “Net Positive” impact is not just moral. It is fiscal. Currently, the lifetime cost of care for a young person with a high-cervical spinal cord injury exceeds \$5 million, predominantly driven by the need for round-the-clock physical assistance [3]. By restoring motor function, AI prosthetics have the potential to invert this economic dynamic. We are not merely reducing the cost of care. We are unlocking the productivity of a workforce that has been systematically sidelined. In an aging society facing labor shortages, the ability to convert care dependency into economic contribution represents a massive, quantifiable social dividend. Furthermore, the psychological benefit of restoring physical capability is directly linked to the restoration of identity. The ability to stand tall or run again is a massive preventative factor against depression and PTSD for trauma survivors.

However, every coin has two sides, and the same is true of this technology. While the capabilities of this technology are promising, we must recognize the danger in treating these devices like consumer electronics. We accept that phones have a shelf life; when an old phone stops getting updates, we buy a new one, but you cannot swap out a brain implant like a SIM card. We have already seen the preview of this dystopia. In 2020, Second Sight Medical Products, a pioneer in bionic eyes, suffered a financial collapse (Figure 2). The company stopped developing the technology. Hundreds of patients with "Argus II" retinal implants were left stranded [4]. Their bionic eyes still worked physically, but the software ecosystem vanished. When parts broke, there were no repairs. When the software glitched, there was no patch. They were left with dark, useless metal in their skulls.

**Figure 2: The Second Sight Collapse - A Case Study in Device Abandonment**



Scale this up to 2030. We can imagine 100,000 people relying on proprietary AI models to walk. We face the risk of ransomware, where a hacker locks the cloud server and demands payment to unlock 50,000 pairs of legs. We face the risk of subscription traps, where a company introduces surge pricing for mobility, forcing users to pay more to walk faster during rush hour. Most likely, we face the risk of bankruptcy. If a startup fails and the servers go dark, the user is blocked.

A common critique of open standards in medical devices is security. Critics ask: *If the code for walking is public, can't a malicious actor hack a user's legs?* This relies on the fallacy of "Security through Obscurity," the mistaken belief that hiding code makes it safe. History proves otherwise. In 2017, the FDA confirmed critical vulnerabilities in 465,000 proprietary pacemakers manufactured by St. Jude Medical [5]. Because the code was closed, hardcoded "universal" unlock keys remained hidden in the firmware for years, allowing potential attackers to remotely drain batteries or alter pacing without the user's knowledge. The obscurity did not protect the patients. It protected the vulnerability.

By contrast, an open-standard architecture benefits from the "immune system" of global verification. We see this in the OpenAPS (Open Artificial Pancreas System) movement, where a community of patient-hackers built an open-source automated insulin delivery system. In a 2022 randomized clinical trial, the OpenAPS algorithm was proven safe and effective [6]. Open-source code allows for transparency, so bugs can be identified and patched by a global network of engineers faster than any single corporate team could manage. In the context of neural prosthetics, transparency is not a risk factor. It is the only way to ensure that the code running our bodies is robust enough to survive the real world.

We cannot rely on the benevolence of tech companies to secure human agency. We must implement it into the policy layer. To ensure the net impact of neural prosthetics is positive, we must adopt a regulatory framework that treats the code running a human body differently from the code running a thermostat. I propose the implementation of the Neuro-Escrow Standard (NES), a two-pillar policy safeguard designed to prevent the obsolescence of the self.

This standard is made of two pillars, the first being the "Right to Repair" trigger. Currently, medical device code is protected as a trade secret. The NES would mandate that any company seeking FDA approval for a Class III neural interface must deposit their source code, model weights, and compiler keys into a federal Public Trust escrow. This escrow remains sealed as long as the company supports the device. However, a "Trigger Event," such as bankruptcy, discontinuation of support, or server shutdown, automatically unlocks the escrow. The code enters the public domain immediately. This guarantees that a community of open-source developers can patch, update, and maintain the software, ensuring that no patient is ever orphaned by a balance sheet.

The second pillar would focus on local-first sovereignty, which entails that we must reject the cloud-dependent architecture for vital biological functions. The NES would require that the Minimum Viable Function, defined as the basic ability to walk, grasp, and speak, must run entirely on-device via Edge Computing. Cloud connectivity can be used for optimization or learning new skills, but it cannot be a dependency for operation. If the internet cuts out or the server fails, the legs must still walk. This Local-First architecture ensures that the user's agency is physically located on their body, not rented from a data center.

The transition from biological bodies to cybernetic ones is inevitable. The technology is already here. The policy is what falls behind. If we do nothing, we risk creating a future where the disabled are liberated from their physical limitations only to be shackled by legal and economic ones. We risk a world where a declined subscription payment causes a literal loss of footing. Though if we implement safeguards like the Neuro-Escrow Standard, we can mitigate this risk. We can build a future where neural prosthetics are not corporate appliances that we use, but open instruments that we control. By guaranteeing that the code governing our bodies is as durable and accessible as the bodies themselves, we ensure that AI does not simply manage disability, but truly erases it. We unlock a future where Julian does not drop his cup, because his hands, finally and permanently, belong to him.

## References

1. Broderick, A., & DeepMind Robotics. (2023, July 28). *RT-2: New model translates vision and language into action*. Google DeepMind.  
<https://deepmind.google/discover/blog/rt-2-new-model-translates-vision-and-language-into-action/>
2. Promed. (n.d.). *Limb loss statistics*. Retrieved February 5, 2026, from <https://promedeast.com/limb-loss-statistics/>
3. Institute for Health Metrics and Evaluation (IHME). (2017). *Global prevalence of traumatic non-fatal limb amputation*. University of Washington.  
<https://www.healthdata.org/research-analysis/library/global-prevalence-traumatic-non-fatal-limb-amputation>
4. Touch Ophthalmology. (n.d.). *Second Sight Medical Products Inc*. Retrieved February 5, 2026, from <https://touchophthalmology.com/supplier/second-sight-medical-products-inc/>
5. U.S. Food and Drug Administration. (2017, August 29). *Firmware update to address cybersecurity vulnerabilities identified in Abbott's (formerly St. Jude Medical's) implantable cardiac pacemakers: FDA safety communication*.  
<https://www.fda.gov/medicaldevices/safety/alertsandnotices/ucm573669.htm>
6. Burnside, M., et al. (2022). Open-source automated insulin delivery in type 1 diabetes. *The New England Journal of Medicine*, 387, 869-881. <https://doi.org/10.1056/NEJMoa2203913>