

My Life as a Cyborg: 2026

Heather Grizzle · September 6, 2026

Meta description: A 2026 rewriting of My Life as a Cyborg examines cochlear implants, childhood autonomy, Deaf identity, assistive technology, and the governance lessons that now shape Heather Grizzle's work in sign language AI.

A 2026 Rewriting

Heather M. Grizzle

Originally written in 2010. Rewritten in 2026 to say what I now think the earlier essay was trying to say.

H.G. Wells sends a sighted man into a valley where no one can see, and the man assumes this makes him sovereign. Nunez has the proverb on his side and the anatomy to back it, and neither turns out to be worth anything, because the people of the valley have built a working world that does not require sight and have no framework in which his eyes constitute an advantage. What they have instead is a diagnosis. The community's physician concludes that the strange soft globes in Nunez's face are irritating his brain, and proposes their removal as a simple and easy operation, and an elder answers this proposal with thanks to heaven for science. The joke is precise and it is not really about blindness. It is about what happens when the authority to define a body sits entirely with people who do not have that body, and when the technical capacity to alter it arrives dressed as benevolence. I read that story as an undergraduate and recognized the operating theatre immediately, because I had been on the table.

I was born sensorineurally and prelingually deaf. Hearing aids were tried in the ordinary sequence and did nothing useful, and I understood the results of each audiological test as good news, since failure meant the world stayed quiet. In 1988 my ear, nose and throat

physician introduced my mother to the cochlear implant, and the next eight years consisted of my fighting her about it, sabotaging the device so that it would not function, telling the physician directly that I wanted it out of my head, and finally, at sixteen, running the external processor over with the first car I ever owned. I have never been able to describe that afternoon without sounding theatrical, so I will simply report that it did not feel like vandalism. It felt like the first time a decision about the equipment attached to my skull had been executed by me.

What I did not understand at the time, and what took me most of the following two decades to be able to name, is that almost nothing about how that device reached me was a matter of medicine. It was a matter of distribution.

The point of first contact was the defining feature. In the early nineteen-nineties, an American parent learned that their child was deaf in an examination room, from a physician, in the same conversation in which they were offered a device. Whatever else that arrangement is, it is a channel with a single supplier of information, and it opens at the moment of maximum parental distress. The physician was not required to describe sign language acquisition, or Deaf schools, or Deaf adults, or the developmental evidence on early language access, and generally did not, and it would be unfair to attribute this to individual malice when the structure produced it reliably across thousands of clinics. Manufacturers had encountered organized resistance from Deaf adults and had found a more receptive market in hearing parents who had, on average, never met a Deaf adult. A large share of pediatric implantation in that period was reimbursed through public insurance, which is to say Medicaid, a detail I got wrong in 2010 and which matters because it establishes that this was public expenditure operating without public deliberation.

Set beside that, the surgical question was almost secondary. The decision had already been made by the architecture of who spoke first.

In 2010 I tried to prosecute this through the Nuremberg Code, and I want to retire that argument rather than repeat it. The Code governs experimentation on human subjects, and pediatric cochlear implantation in 1990 was approved clinical treatment, not research. Reaching for Nuremberg let me express the size of what I felt while giving any competent opponent a way to end the conversation in one sentence. The ethical problem I was pointing at is real, but it lives in a different body of doctrine, and that doctrine is stronger.

Pediatric ethics does not pretend that a ten month old can consent. It distinguishes parental permission, which is what actually authorizes treatment for a young child, from the assent of a developing child, whose views acquire weight as capacity grows. It holds that a child's sustained refusal is morally significant and requires justification to override, and it treats the irreversibility of an intervention and the availability of alternatives as material to whether the override is warranted. That is the framework that fits what happened to me. I was not an infant who could not answer. I was ten years old and I answered. I was asked whether I wanted to hear, and I said no, and the answer was recorded as noncompliance rather than as information. The question a hostile reader should have to confront is not how a baby consents. It is what a clinical system owes a child who has said no repeatedly for eight years, and what it means that the system had no place to put that answer.

The medical claims I made in 2010 also need to be stated correctly, and correcting them costs me nothing, because the accurate versions are worse for the people I was arguing against.

On meningitis, I asked in 2010 where the lawsuits were, as if the answer were obviously nowhere. In fact the Food and Drug Administration issued a public health notification on 24 July 2002 concerning an association between cochlear implants and bacterial meningitis, and implants incorporating a positioner component were voluntarily recalled that same month, and a study published the following year in the New England Journal of Medicine, following a cohort of more than four thousand children implanted before the age of six,

found their risk of pneumococcal meningitis substantially elevated relative to the general pediatric population. Vaccination and monitoring protocols followed. So the honest account is not that harm went unnoticed. It is that a class of devices was distributed to children for years, and the risk was characterized and acted upon afterward, on a timeline set by adverse events rather than by prior evidence. That is a governance finding, and it is more damaging than the one I was reaching for.

On residual hearing, my 2010 claim was too absolute. Implantation does not necessarily destroy every trace of acoustic hearing in the implanted ear, and surgical technique and electrode design have since been developed specifically to preserve it. Loss of residual hearing remains a recognized and common outcome. The defensible statement is that the procedure carries a substantial and sometimes total risk to whatever natural hearing exists in that ear, that this loss is not reversible, and that the person who bears it permanently is a child who was not asked. Overstating the mechanism was unnecessary. The asymmetry between who decides and who lives with the outcome does all the work by itself.

There is one part of the 2010 essay I would not change, and it is the part I did not realize was the most important.

I wrote about the arrival of signed video on the internet, and I wrote about it with an enthusiasm that sits oddly next to everything else in the piece, and the oddness is the finding. Two technologies appear in that essay and they are not doing the same thing. The implant was placed in my head by other people according to their assessment of my deficit. Signed video was picked up by Deaf people and used to argue, in our own language, without translation into written English as the condition of being heard. One is a technology done to a population. The other is a technology used by one. The distinction has nothing to do with sophistication, cost, or clinical evidence. It has to do with who holds the decision.

I did not have that sentence in 2010. I had the feeling that produced it.

The other thing I was circling without naming was what happens when a child resists an assistive device in view of adults. In the videos that circulated then, and in the arguments that formed around them, a child's distress was read almost automatically as a symptom of something else. Poor discipline. A disordered home. Parental conflict. Manipulation by one parent against the other. Every available explanation ran through the adults, and the most parsimonious reading, that the child did not want the thing on her body and was saying so with the only vocabulary available to her, was the one nobody wanted to score. This is not a historical curiosity. Non-use is the most common outcome measure in assistive technology and one of the least examined. A device that is abandoned is recorded as a compliance failure far more often than it is treated as evidence about the device, and the person who abandoned it is treated as the variable requiring intervention.

Which brings me to why I am rewriting a personal essay from sixteen years ago at all, since I no longer work on cochlear implants and have no intention of returning to that argument.

I now evaluate artificial intelligence systems built to render signed language, and I do it in the setting where those systems are actually bought, which is institutional procurement. The technology is unrecognizable. The structure is not.

A system is developed by people who do not use it. Its capability is demonstrated to institutional buyers under favorable conditions. The buyers, who cannot themselves assess whether the output is comprehensible, accept the demonstration as evidence. Someone determines that the system's function constitutes access, and this determination is made without the participation of the people for whom access is the point. The system is then deployed into a hospital, a courtroom, a transit system, a classroom. And at the very end of that chain stands a Deaf person who was not consulted about whether the thing works, and whose report that it does not work arrives too late to affect any decision, and is frequently read as a preference rather than as data.

I have spent three years building instruments to interrupt that sequence, and the design decision I am most confident about is the one that took me longest to reach. In the evaluation framework I use, Deaf governance is assessed first, ahead of linguistic quality, transparency, accessibility and deployment context. This is not a weighting and it does not mean governance matters more than whether the signing is intelligible. It means that Deaf authority is the condition under which the other properties can be assessed at all. If the people who use a signed language had no role in defining what adequate output is, then a linguistic quality finding is a measurement of whether the system satisfies a standard that the wrong people wrote. You cannot fix that downstream. You cannot correct it with a better metric, and no amount of demonstrated capability substitutes for it.

The same principle explains why, in the consent instrument, voluntariness is not one dimension among several. It is a gate applied before the others are scored. A disclosure can be complete, specific and accessible, and none of that signifies if the person had no genuine alternative to accepting it. Where there is no real option to decline, everything downstream is qualified as constrained choice. I did not derive that from the literature. I derived it from being ten years old and answering a question that had already been decided.

This is what I would say to the version of myself who wrote the first draft of this essay, furious and twenty-something and reaching for the Nuremberg Code because it was the largest instrument in view.

The problem was never that the technology worked. In some respects it worked well, and there are people for whom it works and who wanted it, and their accounts are theirs to give and are not diminished by mine. The problem was that working was treated as sufficient. Capability was allowed to stand in for legitimacy, and once that substitution was permitted, every remaining question became a matter of engineering. Whether the device was wanted, whether alternatives were disclosed, whether anyone with the relevant body was

in the room when benefit was defined, what recourse existed if it failed, and what standing my refusal had, all of it was reclassified as sentiment.

I would put it this way now. No technology earns trust simply because it performs as designed. Trust depends on who defined the problem, whose evidence was permitted to count, who participated in the decision, what alternatives remained genuinely available, what happens when the system fails, and whether the people who bear the consequences retain the authority to decline.

In 2010 I knew what it felt like when my no was insufficient. I have spent the years since learning to state, in terms an institution is obliged to answer, exactly why it should not have been.