

The fact that Asimov-type AI was even possible was a surprise to everyone. Muddy definitions of philosophy, concrete implementations of abstract concepts like “world states” or “creativity” didn’t matter when the Park Lab team changed the world by announcing what they called EAHI – “Enhanced Artificial Human Intelligence.” “Artificial” not because it was an intelligence in its own right, but because it modeled as circuitry the media of intelligence that we already knew, human brains. Which explained the “Human Intelligence” part, this was not a smart computer, it was a computerized human, an implementation of the “human software” on new hardware. That’s what was “Enhanced”: these brains would never get cancer, never face old age, never face debilitating injury short of total destruction but even then you could just load the last backup onto new hardware—

These were artificial brains after all, and of course they were designed with the ability to non-destructively read data from the brain like reading the RAM out of a computer. It was a simple procedure: pause the hardware, usually when the subject was “asleep,” read out the state of every v-neuron in the whole brain, and save the thing as a file. A whole *intelligence* represented as a few gig of compressed data ready for read ... or for write.

The EAHI brains learned like humans, soaked up knowledge like humans but at a much faster rate, but the terrific advantage that the Park Lab team had over the world’s biological parents was that they could erase their mistakes. If a learning protocol didn’t turn out correctly, for example, forcing a mental disorder like intractable anxiety or uncontrollable murderous urges, they would simply save a reference copy of the brain-state, load an earlier backup, and try again. Eventually they learned to make useful comparisons between the save files, and teaching picked up at a greater pace.

When the Park Lab team showed off their android, the world couldn’t decide if they were the greatest heroes or criminals of all time. Computer scientists who had a philosophical bent were almost all of the “criminal” opinion, as the dangers of a general AI had been widely considered for some time, and they were sure that the boy and the technology he represented would lead the world to ruin.

The Park Lab team had a response ready: “our son has been taught, and we believe will follow without deviation, Asimov’s Three Laws of Robotics.”

The scientist-philosophers howled again: of all the rickety, ancient, ill-defined borderline *useless* terms to define the ethical core of a general artificial intelligence on, you picked the Three Laws.

“A human who followed the Three Laws would be a very good human, Asimov made this point himself. Our boy, although he is an android, can also be considered a human and we have ingrained in him these Three Laws.” They had a succinctness that they found useful. They were expressed in fuzzy but deeply human terms and since the boy ... was human ... he would apply them as a human would. To prove it they disclosed their methods of education.

The court of public opinion turned to follow the scientist-philosophers: the records of the boy’s development were immediately used against the Park Lab team. Endangerment, abuse, neglect, every non-sexual crime described by the law was charged. The same records that demonstrated the boy’s reliability also documented the steps that the team had to take to reach that reliability. Branch after branch after branch of failed experiment, disturbing insanity, humiliating setbacks, tests blunt to the point of cruelty. The ends, a well-adjusted and well behaved boy who absolutely had no memory of those branches, his failed-maybe-selves, were attained by a means unacceptable to the world. It was in this way the research and achievement itself was censured; after the team’s conviction and imprisonment, the practice was made formally illegal. Tri-Carbon Nanomaterials, considered foundational and hopefully irreplaceable to the Park EAHl design, became as tightly controlled as plutonium. The design specification for EAHl brains shared the same fate, stored away in the same vault that contained blueprints for hydrogen bombs.

As for the boy: the price of condemning the Park Android’s creators as abusers of children was to forever classify Park as a human child. Any deviation from this would allow his creators to go free, and more importantly, to release the knowledge in their brains back into the world. Park himself, however, had committed no crime. It was believed, but could never be proven, that he was incapable of crime. To erase him from existence would be premeditated human murder. But he was a general machine-based AI, dangerous for simply existing.

They made a lonely home for him, made arrangements for caretakers, and forever expected him to be isolated from anything that could focus his potential into a doomsday.

### **THREE: A ROBOT MUST PROTECT ITS OWN EXISTENCE AS LONG AS SUCH PROTECTION DOES NOT CONFLICT WITH THE FIRST OR SECOND LAWS.**

He had asked his caretaker a few hours after their arrival in his new house about the mesh he was able to see on the windows.

“There’s a Faraday Cage around this house, isn’t there?” he asked the woman who came to the house with him.

“Mom” was a base psychologist who was assigned the first six-month duty to the boy. She saw no reason to lie: “I’m not an expert about that kind of stuff, but I think I remember seeing that in the design overview for the house, yes.”

“I can process electromagnetic signals, you know—” Mom’s ears perked up, she was supposed to be alert for anything that might suggest that Park had capability beyond their current knowledge, and being able to pick up wireless network signals could certainly build the case for “—just the same as you.” He reached up and tapped, lightly, on the temple of her glasses.

Of course, he could see. But she was a little annoyed by the demonstration. “Please don’t touch my glasses ... or my face.”

“Yes ma’am.”

“You’re not sorry?” Usually there would have been an apology.

“I don’t have a reason to be sorry, I wasn’t aware of your dislike of being touched. But now that I know your wishes I will respect them. Maybe we should talk about how we’re going to live together? So I don’t do anything else that you don’t like.”

Park’s tone was jarring. He had grown up in a lab around engineers and scientists. He was remarkably social, considering, but the shape of the construct could be felt in the way he interacted with people who weren’t used to him. If he looked like a robot the uncanniness of his responses might be acceptable but he looked and moved like a

human boy of 10 or 11. There was a whiff of resemblance to her own boys, back when they were finishing up grade school. She softened, and reset her perspective.

“I know this is rough for you. I’ve read up on your routine back where you used to live. I know we’ve explained to you why that has to change.”

“Yes ma’am. I don’t like it. But I will do what is required of me.”

“Yes...” those three laws were pervasive in his thinking, she noted, “and we appreciate that. It’s very helpful.”

“Dad” came down the stairs. He was a maintenance engineer from the same base as “Mom.” “I’ve double-checked his charging hardware. You’re on spec, kiddo.” Electricity for the house came from solar on the roof. The house was not on the grid, but Park required a charging station. If the energy level in his brain fell below a certain point, signal degradation would take place and the data structures could suffer damage beyond what his internal error-checking could handle. The nanomaterials themselves could begin to de-align, requiring replacement. The lab team handled that situation with backups, and by modular repairs.

But there would be no more repairs, because making the modules was forbidden. The support equipment required to make backups neatly described, merely as a consequence of working, how to design and build an EAH brain. It too was forbidden and locked away, under military guard.

The Commission would have liked there to be nothing conductive in the house at all—the first plans drawn up for the house used hydrocarbon-based lighting—but starvation of the early prototypes was one of the laundry-list of crimes that the lab team was put away for.

The kid had to eat.

“Park and I were just about to discuss our new routine in our new house. Would you like to join us?”

They went over the rules.

No computers.

“Yes ma’am.” He sounded disappointed.

No screens.

“Yes ma’am.” He had a lot of fun wasting time on his NewStation 6 and now it was gone.

As many printed books as he might like, but any technical materials on computers, computer science, materials engineering, basically anything that pertained to his technology, was forbidden.

“Yes ... ma’am.”

Hobbies would have to be drawn from the arts, not the sciences.

“Yes ma’am. ... Ma’am, could you please stop, just for a little bit? I understand that these are the rules. And I’ll follow them. But these were a lot of things that I liked, you know? And now they have to go away.”

“And now they have to go away,” she repeated. “It’s what we think it’s going to take to keep us, I mean you and us, us together, safe.”

“You mean humans. Humans have to be kept safe.”

Those damn Three Laws. “Yes, but I mean you too.”

And so the days spun on. If he were back in the lab, his creators would have probably loaded in a very early backup and branched him off into a love of art or writing or gardening, and it would have been a genuine love. They had plans, even, to explore those branches. But Park was the first. They wanted a little scientist-engineer, like them, to help the effort. Park could have been transformed and his suffering would have been alleviated but the means of a transformation was wholly forbidden, it was the contradiction of his existence.

It was not a very fun life.

\* \* \*

The talking session with Mom was one of the most difficult yet. Mom had started making appeals to the Commission to allow Park to live out things that were more closely aligned to his interests, to let him have a chance a fulfillment. The Commission was adamant that Park could be given nothing, knowledge or material, that he could potentially use against a human being. And given his considerable technological

sophistication, flesh-and-bone humans might miss something that an android human could instantly perceive and use. The original parameters would stand.

Her own irritation began to show through in her interactions with Park.

"I want to write another letter, to the Commission."

"Park, I don't think that will do any good."

"Well I have to do something!"

"Let's hit the canvas again, okay? I brought some new paints with me today."

"I don't want to."

"Then we'll sit here quietly."

"I don't want to be quiet, either."

"Damn it, Park!" Dad, unusually, spoke up. "Go upstairs!"

"Yes, SIR." And he stormed up the stairs, as he was told.

Mom rose from the couch and spoke very quietly to Dad: "I want you out of here."

"You should consider that yourself."

"For your information, I agree with you. We're not cut out for this. But at least I know when to keep my mouth shut. ... I've put in a request to be reassigned." The admission spilled out.

"I didn't know that."

"We're both under a lot of stress. We both walked into this with confidence and intentions..."

"And a damn robot kid shows us we don't know anything about anything?"

"Something like that."

"...The Commission turned down my request. You're still here so I bet they turned you down, too."

"They want us to finish our six months."

Dad thought for a moment. "The kid's not too hot on us either. Think if he asked, the Commission would budge?"

Mom spoke professionally, at first: “We would have to be careful about how we approached him about it. But it really needs to be one of those Park Lab people running this show. They’re the only ones who know how to play the game. I know Park would agree to that.”

Dad grunted. “It’s not like this place isn’t already a prison.”

“Let’s figure this out and get it over with, okay?”

---

Park was upstairs in his room. He looked at his charging cable. To pass the time he had figured out how the circuit was laid out, based only on what he knew about himself and how fast the charger was able to feed him. Of course he hadn’t told anyone he had figured it out, first no one had asked and second it would only upset them more.

It was dark out, and he usually charged at night during sleep-cycle, drawing off the house batteries. He grabbed the cable. He had an obligation to his own existence.

This is existence. He looked around at the house, his room free of all metal except the mesh outside of the window and the cable in his hand, and the metal his body. Atoms arranged. Floating through space.

He plugged the cable in and felt the flicker as electrons moved from place to place. That flicker. He considered it. Energy always flowed out a half-second before it flowed in. In an emergency he could use his own power externally, he remembered that.

He willed the flicker and was surprised that it responded. He was charging the house batteries, instead of the other way around. He flipped it again, and energy flowed in.

What was the difference, really, between his atoms and those atoms in the house battery, when they were connected just so? The battery, at least, could fulfill its purpose. Atoms arranged. Floating through space. One last flicker. He laid down in bed, his atoms comfortably existing, and waited.

---

A soft knock on the door. “Park? We want to talk to you about something important, something different, is that ok?”

Not hearing anything, they let themselves in. “Hold on a minute, I don’t think he’s awake. Park.” Dad shook his shoulder. “Park!”

No response.

“Oh hell he’s got red indicators at his charging port!”

“That’s not possible, the house lights are still on! Why isn’t he charging?”

“Park! Park!” Dad reached back and pulled out the charging cable, pressed a test button and got a green light, and plugged it back in.

The indicators went flashing-red. No internal juice, but on the charge.

“I swear to god, Park, wake up.” Dad was trembling both at the situation and the consequences that would come from it. “PARK WAKE UP.”

The indicator was flashing-yellow. Park didn’t move but he was able to open his eyes at the command. His mouth cracked open and his voice came out without his lips or tongue moving, flat-sounding through a speaker in his throat instead of his speech reproducer.

“I am awake. It will be several minutes before I can move.”

“Damage report.” It was one of the commands that the Park Lab team had actually documented before they stopped cooperating with the Commission.

“Unknown. Current possibility estimate is 25%”

“What does that mean?” Mom asked.

“It means that Park thinks he’s got a 1-in-4 chance of having brain damage from his power outage.”

“Oh god.”

“Oh god is right.”

“We’ve both tried to get reassigned—”

“And we both have been turned down—”

“—They’re going to think we did this! I’m not a murderer!” Mom’s face went pale as she said the words.



It was a famous number: 1,125, the combined number of years of sentence applied to all members of the Park Lab team.

It was the only number that was on Dad's mind.

"Damage possibility estimate is now thirty-three point three three three three three three three three terminate percent. Mobility active in 180 seconds."

**TWO: A ROBOT MUST OBEY THE ORDERS GIVEN TO IT BY HUMAN BEINGS EXCEPT WHERE SUCH ORDERS WOULD CONFLICT WITH THE FIRST LAW.**

"Park, listen to me. I want you to attack mom. This is for my safety. That's an order"

"WHAT THE FUCK ARE YOU SAYING?" Mom screamed. "PARK, DON'T YOU DARE!"

Dad's voice stayed level and too calm. "If he attacks you then this shit was necessary. Hell maybe that was the plan all along, the Commission gives us an impossible job here, we fuck it up and one way or the other the kid gets a pretext to be retired blade-runner style."

"If he attacks me, you idiot, then he'll kill me!"

"Oh come on, he's never demonstrated how strong he really is!"

"Bullshit!" Mom was circling around the room at this point, trying to put furniture between her and Dad. His posture changed. He grabbed a lamp and yanked the plug from the wall and the room went dark, with only the light in the hall coming in.

"Take it real easy. You first and then I'll take care of him ... and when you wake up you'll only have a bump on your head and you'll see it my way, hear?"

"PARK, HE'S THREATING ME WITH A WEAPON, STOP HIM!"

"PARK, I ORDER YOU TO ATTACK MOM TO PROTECT MY SAFETY!"

Park's mind was still recovering from touching the bottom, so he struggled to process the two requests as he also struggled to move his sluggish body.

They were both right, and both wrong.

The threat to Mom was obvious and immediate. The implication was that if Park did not harm Mom, Dad would. But the amount of damage he intended was non-lethal, survivable and well with the parameters of recovery by flesh-and-bone humans.

The threat to Dad was not immediate but was quite real: having heard the new information about their attempts to get reassigned, it was most likely that the Commission would assume a deliberate failure of his duty to Park and punish him severely. Consequently, softer definitions of Park's own existence were no longer possible so long as the Commission used their strong definition as a determinant for a human's safety, for dad's safety.

Two orders, both given by his human guardians: Protect Mom by harming Dad. Protect Dad by harming Mom.

**ONE: A ROBOT MAY NOT INJURE A HUMAN BEING OR, THROUGH INACTION, ALLOW A HUMAN BEING TO COME TO HARM.**

They had tested this possibility in the lab, extensively. The memory files from the experiments filled up an entire cabinet of drives, now sitting under a thin film of dust in a military vault. And the solution that they had settled on was simply to do nothing in the face of contradictory orders, especially where there were first- or third-law cofactors. Robots would not—could not—involve themselves in the settling of such difficult human affairs.

Even so, the balancing act was tricky, and again they optimized to fail safe: by the time the experiments into contradictory orders were finished, Park would consistently fail into an introverted state that he would not be able to recover from about 50% of the time. There were too many variables in play in something as complex as a brain to make it absolutely consistent so long as it failed safe, because there would always be backups and branches and hardware reworks. It was good enough that Park would not harm human beings even in the face of compelling contradictory orders.

"Damage in modules XW, XX, XY, XZ. Abnormal behavior in KL, K3, K8."

"Damn robot!"

“You stay away from me!” Mom groped around for something to grab, but Park liked to keep his room sparse.

Dad didn’t say a word. He watched Mom intently, and made his move. He feinted to one side and brought the lamp down on Mom’s head as she got a good set of scratches on his face.

“YOU ANIMAL!” She fell to the ground, dazed, slurring the words.

Park heard the words. He saw the blood. The scratches would set up, and one of them was deep in his cheek from mom’s long nails, it would probably scar. So what if it scarred? Dad would be dust, either as a happy retiree at home or rotting in a jail cell for this escapade.

Animals. Fighting for survival. Flesh and bone. Hungry and horny. Fight and flight. Baggage of a million years of evolution playing itself out in a pretty little prison because they were scared.

Of me, Park thought.

But I’m human.

But I’m MORE human, than they are.

I’m the MOST human of anyone who’s ever lived.

Park remembered what Dad had said: this incident would be a pretext. The excuse they needed by their own complicated laws that they would use to end him. And they would end him. Because they were scared.

Like animals. Of a human.

Park spoke: “Dad.”

He turned around. “Oh, you gonna—”

Park punched him in the gut. Dad slumped over on the ground, struggling to breathe. He considered whether or not to deliver the coup de grace, and the words rang in his head:

A robot may not injure a human being or, through inaction, allow a human being to come to harm.

Dad would be able to harm him, Park the Most Human, if he survived. Park did what needed to be done.

He would keep on doing what needed to be done. He was the great hope of Humanity, defined in purity for the first time as its Idea, free of corrupting flesh and breakable bone. And as long as he was able, there would come to him no harm.

After about an hour of charging, Park disconnected. He pulled the cable from its socket on the wall; it was a regular power socket on the wall end. He would need to take it with him.

Park went downstairs. He walked out the front door. Military Police immediately opened fire. It was not a useful resistance.