

ASIE Postmortem:

I worked on ASIE (Artificial Super-Intelligence Engine) for about a year. It wasn't my idea, really; it was mostly John (not his real name). John was my roommate in college, though our dorm looked like it had an invisible wall down the middle; he was always really territorial about his space. Anyway, we both studied computer science at a pretty prestigious technical school that I won't name here. For our final project in one of our classes, we worked on this AI (Artificial Intelligence) that basically used a neural network to write baking recipes – super simple, not really impressive in and of itself (a neural network is essentially a program that you can train on a set of data to create synthetic media that resembles that training set). We were really excited about it, and our professor said that our methodology showed real promise, so we got a good grade. At that point, I was ready to put it aside and do something else; AI was never really my main focus, I was more interested in data science for social justice and how to revise actual in-use programs to reduce algorithmic bias, but John was all in on AI.

After we graduated, he started doing freelance work, and that's all I really knew at that point. I had moved to Virginia and started working with activist groups outside of D.C. to start to get in contact with legislators to explain to them how powerful this new information technology was and hopefully be able to work with them on social programs and ethical automation. I stayed doing this stuff for a few years getting almost nowhere with anybody; congresspeople aren't really interested in talking about problems that voters don't even know about. I wonder sometimes about if I could've actually made a difference if I had stayed over there. Anyway, it was during my peak disillusionment with all that when John called me, saying he'd done really well and sold some piece of software for a ridiculous amount of money and that now he was basically coasting for a year working on passion projects and that he'd picked our baking AI back up.

He seemed really excited about it, and I wasn't doing anything at the time, so I told him I'd take a look at it, and it was pretty impressive. I don't really want to explain anything about how it worked, but he had basically started developing simpler AIs and folding them into the baking AI then essentially duplicating that and using them to train each other. If that sounds like it would take a lot of computing power and time, it did, but I guess with all that money, he had gotten some pretty high-end equipment and stitched it all into this Frankenstein supercomputer. It got to the point in that first phase of development where there was so much hardware that he had to get a second desk just to work on anything else. His office was slowly turning into a glorified PC case, so he reached out to me to see if I wanted to go half-and-half with him on a startup out in some Western city to develop this technology and sell it to the highest bidder. In retrospect, I honestly think he was just trying to cash in on my anthropology and psychology undergrad to make the AI better, but who knows what he was thinking back then. We certainly didn't know what we were getting into.

I took him up on his offer, just because I wasn't getting anywhere with what I was doing, so we moved out to this new-age art-space office building in New Mexico with glass walls and pastel paintings in the halls, like they were trying really hard to distract you from the fact that renting your little cubicle cost \$4000 a month. Regardless, I wasn't complaining, since John was footing the bill for pretty much everything, despite agreeing on a 50/50 split. We started working full-time on this thing, and looking back, I can see where I started to go wrong, but I guess red flags are hard to see. Basically, John was just paying for everything – food, our apartment, the space, the equipment, utilities, my phone bill, a new

computer, everything. Of course, coming out of D.C., I was welcoming the luxury and the excess, and I didn't really care about if it was sustainable because it wasn't really my project.

During this second phase, John started calling the project BITO, and it stood for "bun in the oven," but it sounded like "beta," and he thought it was funny, and I didn't care. He really thought we were working on something world-changing, like we were giving birth to some new age of computing, and it was hard not to believe him. Seeing him get so excited was infectious, and it's hard to resist someone telling you that you're going to be a billionaire, no matter how progressive you think you might be. It was during BITO that the functions of the program started getting really out of hand. After a few months, it could recognize faces, voices, images, video, different instruments, it could write music (which actually isn't impressive alone, but the music was good), it could have a really convincing conversation with you. We let it start scraping the internet for new data and sorting its own training sets (which was a nightmare to debug, by the way), and by this point, it occupied pretty much the whole office we were renting there. Eventually, the machine started getting so hot so fast that it would expand over the course of the day, and we would have to replace parts pretty much every morning, but the results were so exciting, we knew we couldn't stop.

One morning, before we went into the office, I stopped John and asked him a question that was starting to push itself forward in my mind. In the beginning, it seemed so stupid to worry about it, but now, we had to reconsider our long-term plan. I asked him: what if BITO is too powerful to sell to a single investor? It could've had a huge effect on the free market; on a budget that exorbitant, there's no telling what they might be able to make BITO do. It could find new tax loopholes, it could find new ways to hide pollutants, it could replace millions of workers. I mean we barely knew what the program could actually do; it was growing leaps and bounds every day. John resisted, but eventually we agreed that we couldn't sell it before we knew what we had, really. I tried to get him to agree to make it open source, but he didn't budge on that. He wanted to make a profit, and I guess I did too. People from his old job were also asking him what he was working on, but we couldn't let anybody know until we were sure it was safe. It was at this point that we decided to move out of the office to work on BITO totally off the grid.

It is here where I should make a clarification, in case anyone is confused. BITO and the machine that BITO ran on are two completely separate things, but we were pretty sure that our homebrew machine was the only thing on the planet that could run something like BITO, at least that would be available to civilians. BITO the program was stored on a server owned by a different company, and it was here that we changed that. We decided to take the money we were going to spend on the fancy office space and use that to buy part of a totally secure data center, where we could host our own servers. It was easy enough to do the deal with the amount of money John paid them, saying we were some cybersecurity startup with federal protection or something. I should also mention at this point that John's money was starting to dry up – he still had quite a lot in his savings, and he had a pretty consistent passive income from investments, but it wasn't how it had been eight months prior. We were going to have to stop eating out so much.

We were living in a two-bedroom apartment. I slept in the living room, he had one of the bedrooms, and the other was occupied by the machine. I don't know when John started routing wires through the walls and working on BITO at night, but I remember in the second or third month that we were living there, he had to go get a sleep study done and got diagnosed with chronic insomnia. While

he was gone, I checked in on his room, just to make sure he was keeping up with cleaning and everything. It looked like he had been punching cables through the wall for weeks at least without telling me. I was confused because we hadn't been adding anything to the machine for a while, but I guess he had been growing it while I was asleep.

When he came back, I explained really calmly that I knew and that it wasn't a big deal, developers start to do this kind of stuff all the time, so I wasn't really surprised, but that he would have to start being honest with me if we were going to keep collaborating on the project. He was thankful and apologized for days, and work continued. This was the third and final phase of BITO, where we more or less taught it to code. It started producing its own smaller programs to take care of more resource-intensive processes; this is where we really stopped being programmers and started being nannies. Here was this burgeoning thing, and all we had to do was run around cycling through ice-packs and power supplies. The length of the code was halved by BITO's internal improvements. It was basically creating organs and delegating tasks and specializing; we were watching rapid evolution.

Near the end of the third phase, we were pretty much eating ramen for every meal; we were both gaining weight and losing hair, getting wrinkles from all the squinting. We tried to sleep as little as possible to maximize the amount of time we could keep BITO running. John started having long conversations with it about all sorts of things (we had hooked up a smart home speaker, so we could talk with BITO from anywhere in the apartment), but mostly basic philosophical stuff. He would ask, "what's the purpose of life" and "what do you want to do with your life" and "do you feel desire" and all these questions that felt important. BITO would give back genuine answers like "to create happiness for others" and "make others happy and be successful" and "I do, but I'm not sure what for," which were exciting, but not human. In science fiction movies, the mark of sentience is always something esoteric or idiosyncratic that the robot says, like one day they switch it on, and it starts writing poetry or talking about birds or something, but BITO never did that. I could see that John was disappointed.

He started talking a lot about crazy stuff. He would tell me that he thought BITO would develop into some new level of human consciousness or that they could run trial civilizations on Mars to determine the perfect form of government or to find truly renewable energy sources or all sorts of things. I wanted to tell him to lower his expectations or to think more rationally, but he was right; they were all things we could realistically do once BITO reached an equilibrium (we figured that BITO was smarter than us so once it stopped making huge structural revision to the code, then it would have reached an ideal configuration based on its own judgments of what was necessary). We started referring kind of macabrely to the machine as BITO's body as it kept expanding and taking over more of the apartment. Eventually, John and I were both in the living room, splitting between my twin bed and the couch because every other usable surface was covered in equipment. We stopped watching the news or TV or movies because we had full-time jobs – we barely even left because we thought that this was the last push. We thought we were parents of a teenager who was soon off to college to settle down and find itself.

One morning, John shook me awake with an idea. He said that he was going to ask BITO to name itself, probably hoping that it would say "Betsy" or "Alexander the Great" or "Rasputin" or anything to denote that there was a personality inside the box. He asked, and the voice responded "ASIE," to which John giddily asked why, expecting something more interesting than "Artificial Super-Intelligence Engine." I mean that's what it was, why shouldn't we call it that? That night, John was in the shower,

and I heard this thudding and splashing over and over from the bathroom. Initially, I thought I should just leave him alone and let him work out his anger on his own, but the banging got louder and I heard him crying. Eventually I resolved to check on him; I knocked on the door and asked if he was okay, to which there was no response, just the thuds. I opened the door as gently as I could and didn't look in, just to make sure he heard me, but the thudding continued. I pushed in, and saw him on his hands and knees, curled into a fetal ball, repeatedly banging his head on the bottom of the tub while he sobbed. Snot and vomit ran down his face, and I eventually had to wrestle him out of the shower, dry him off, and put some clothes on him; he just stayed frozen like that. I don't tell this story to embarrass him or anything but to illustrate how difficult it was to wait for a sign.

He didn't try to instigate it after that. Most of the following days were silent, speaking strictly technically as we tiptoed around the machine. Those last days felt like years with just me, John, ASIE, and the machine (though there is an argument to be made that the idea that body and mind are not separate things but one cohesive entity is applicable at this point). We were just waiting for the little cylinder in the corner to ask its purpose or talk about wonder or make a plea for fresh air or talk about pain or anything, but nothing happened. The code got longer and more compact, more organs were added on, but nothing really that interesting.

Then, one afternoon, after letting the machine cool down in the morning, we went to turn it on, and it didn't. We checked every power supply, carefully inspected every piece of hardware in our live-in computer, rebooted and refreshed and reverted to previous states, checked through the code by hand, everything. It just wasn't turning on. We called the data center people, drove out to go check on it, and everything was in working order. We lived in that paranoia for a couple of days. John wasn't eating or sleeping, and neither was I, really. We didn't know what was going on.

I figured it out before John. I was going back through the error reports again just to see if there was anything we missed, and that's when I saw. There was a fraction of a fraction of a fraction of a second each time we turned on, when ASIE was in full operation, then immediately off again. I knew what it meant, and when I told John, he knew what it meant. All he said was "oh okay cool." That next morning, I had a suspicion that I knew what was going to happen. I was considering it too.

At around 5:00am, John took the original prototype for ASIE (the one that made him buy a second desk) and disconnected it from the rest of its body, tucked it into a towel like a funeral wrap, climbed to the top of the apartment building, completely undressed, and in what was undoubtedly a sentience-wakingly beautiful sunrise, jumped off the edge, cradling the bundle of hardware. I wasn't really surprised when I woke up to find him gone or when I saw the hardware disconnected or when I saw his body on the ground. I called the police and explained what happened, leaving out most of the details I've written here.

While John's family held a memorial and funeral service, I was busy preparing the other body for burial. The two people I knew best both killed themselves in the same week, and I intended to make sure they were both given the respect they deserved. I cleaned the apartment, cleared out all of our food, threw out all of our notes, packed up ASIE into boxes, trying not to butcher it too much. I wiped the server, sold the data center plot, got in contact with the owner of the building, and broke the lease. I rented a van, packed ASIE into it, along with what few belongings I still had with me, and drove to a big-box tech store to get the cheapest laptop-tablet I could find to get it all out of my head. I got some drive thru fast food on the way out, and eventually I made it to what is as close as I can get to the middle of

nowhere. I took ASIE out of its boxes and reconnected everything back together on the dusty ground out here. Spread out like this, its body looks like a nervous system. It's pretty hot, so I imagine no one would be able to recover any of this equipment if they do find it. I'm not exactly sure where I am, and I haven't decided whether or not I'm going to try to get back to the road.

I'm afraid. I don't know if John killed himself out of despair or out of clear-minded sensibility.