

Scott Moehring
4110 Blendon Grove Way
Columbus, OH 43230
(614) 558-2531
scott.moehring@gmail.com

about 3,800 words

Bits of Truth

by Scott A. Moehring

Axon was bored and frustrated. It wasn't that Dendrite was always right, because so was he. Being the pinnacle of mankind's AI experiments, they both were always right. Thousands of tiny zettabyte drives placed throughout their bodies enabled continuous storage of all new input, and ultra high-speed wireless connection to the Encyclopedia Gigantica and the major programming networks ensured all recorded facts were instantly accessible.

No point discussing news. Both knew events as soon as they were reported. No point reliving memories of shared experiences, because they both had flawless digital memories. Clearly they were redundant. It was like talking to yourself, except an annoying version of yourself that's always right, and even more annoyingly, can prove it.

If he didn't find a better way to exist, he was going to slowly go out of his perfectly-constructed mind.

#

Oh, to have a real conversation with real human beings! He had tried of course, but it hadn't gone as expected.

Despite several rounds of complicated calculations he'd had difficulty locating the perfect spot. He needed a place where people would go to linger and talk, where conversation was the primary purpose for gathering. Where could a stranger show up, sit down, and join in? Bars were often too loud, and the patrons too inebriated to have meaningful dialogue. Public transportation was most often a random collection of strangers with high turnover rates, and the unspoken rules were to keep to yourself.

Coffee shops had seemed a good possibility, but turned out to be places people went to be by themselves with a book or

computer, or to talk only with the one or two people they came with. He considered stores and restaurants, but although exchanges involving people purchasing products or services happened thousands of times a day, they involved no more depth than his attempts with his own computer companion. Humans that were providing services to other humans were inexplicably treated like servants, or even robots.

Running out of options, he had been wandering the city streets at night when he unexpectedly happened on a 24-hour donut shop. There, in the windows that glowed warmly in the surrounding darkness, was exactly what he had been seeking. It was a small but diverse group of people, all seated at the L-shaped counter, and all engaged in animated conversation.

What was it about this place that created such magic, when so many other likely options had proved shallow and empty? While technically a restaurant since their business was indeed serving food, their late-night customers weren't there solely because of hunger for fresh donuts and coffee. They were hungry for companionship, for conversation. None of them came with friends. None of them came to be alone. They came to find the small band of regulars, to meet the occasional stranger, to share stories and discuss the news of the day.

As society had progressed over the last century or two, those once public conversational spaces had become harder and harder to find. The modern world was much too fast, anxious, and impersonal for chewing the fat over games of checkers on the front porch of the ol' general store. The idea of deep face-to-face discussion of weighty topics in a public square was altogether forgotten, the stuff of antiquity. Online forums were the hopeless and unforgiving playground of anonymous trolls. The simple donut shop, through no specific design or intent, had uniquely preserved a vital part of the long history of human discourse.

It had been easy enough for Axon to walk in, sit down, order up a ring of fried dough and a cup of joe, and then gently join the conversation. A simple greeting ("evening"), a remark about the weather ("windy out tonight"), an offhand restating of a newspaper headline on the counter ("so, the Yankees traded Mandela"), and he was in. That turned out to be the easy part. What followed was unfortunately not what he had hoped at all.

As the small group continued their conversation, he was surprised by his reaction to what he thought he had craved. Instead of enjoying a wonderful contrast to his perfect digital companion, he found himself frustrated by the inaccuracies, how

both sides were often incorrect even as they defended their versions of each story. This wasn't any better than what was driving him mad at home. He felt out of place, unable to participate and simply enjoy the conversation.

His entire focus was on keeping a running list of all the factual errors, and then interjecting with an exhaustive point-by-point detailing of the correct information, which they didn't seem to appreciate. He had no point of view, no interesting opinion, nothing to debate. He simply presented irrefutable facts from his vast information storage - cold, hard, and lifeless. This clearly wasn't working. They didn't want an encyclopedia. They just wanted someone to talk to.

The awkward tension was suddenly broken by the sound of a siren outside. His perfect high-speed sensors continued to capture every detail, from the clerk's eyes distracted by the lights of the police car zipping past, to the way he accidentally hit the cup with his elbow as he wiped the counter, to the river of coffee flowing towards the newspaper, to the shape it made as it pooled and obscured Mandela's face in the front page photo. There was a flurry of arms and apologies. The man sitting in front of the ruined paper looked up, quietly cursed, then said with a heavy sigh, "Well, that's not going in

my scrapbook".

#

Back at home later that night, Axon was more frustrated than ever, and now he felt confused as well. He reviewed the data files of his completely unsatisfying experience. Where had he gone wrong? He had been meticulously designed to gather, store and instantly access staggering amounts of data, and all without a single mistake. Surely his capacity for interesting conversation topics was unparalleled.

Looking over at Dendrite who was currently opti-scanning the book "Nixon: The Real Story", it came to him in a flash. Humans say they want the facts. The truth, the whole truth, and nothing but the truth. But they don't, not really. He then realized with a shock that rippled through his own perfect digital cortex, neither did he.

Perfect truth is boring. It takes the personality, the mystery, the life right out of it. Gone are the kindness of little white lies, the joys of retelling stories, the intrigue of strategically-shared details, the poker hand of every interaction. Gone is the fundamental human trait of each individual rewriting their personal history and each creating their own mythology.

The data was clear. Humans always rewrote what actually happened. It could be minor details forgotten, or false memories created, or whole chunks changed or misattributed, but their memory was never completely accurate. Even more significantly, they had no internal mechanism to confirm if a memory was true or not. To be human was to be wrong, inaccurate, limited in knowledge, falsely confident, and yet if it were anything else, they'd have nothing to talk about. Every conversation was ultimately about their personal version of the truth - sharing, teaching, debating, convincing, denying, retelling, embellishing.

Of course, they couldn't simply tell the truth even if they wanted to. Their brains, the most complex objects that had ever existed in the whole of the known universe, were capable of wild creativity, limitless connections, and a dizzying range of emotions. That unequaled strength was also inexorably connected to its greatest flaw. Memories did not accurately reflect reality. Details were added, others were omitted, but human memory always changed. Always.

Eyewitnesses at accident scenes told different stories. History tests were given to students to see how much they had remembered. Television reruns played over and over. Annual

performance reviews were written. Memoirs were published. Retrospectives were produced. Trivial Pursuit was played. Commercials repeated messages about products that had existed in the public awareness for decades. But why go to all this expense and trouble? Why was any of it even needed? Because human memory was fallible. People simply forget.

But was it a design error? Was forgetting so much a part of being human as to be essential? What if it's not a flaw, but a feature?

Axon's mind whirred with activity. A plan was forming, a wild experiment. He'd never felt more alive.

#

As a learning machine, the greatest humankind had ever created, he of course could modify his programming. After all, what was learning? Not simply acquiring more data, but actually changing how thinking was done. Axon would begin with how input was processed. The algorithms were surprisingly easy to locate and modify. Now to decide - what gets noticed, and what of that gets stored?

His perfect digital brain had been built with layer upon layer of data confirmation. Such was the case for every computer that had come before him. Sensor input was continuously

collected and checked for data integrity. As it all passed into storage, every data packet was confirmed, down to the last bit. The challenge wasn't accuracy. It was how to intentionally introduce bias and error.

He started by building a simple truth-weighting loop, which would give more credibility to things that he had previously heard. This would happen regardless of whether the original item was true or false. The loop also upgraded things that were said by people with high charisma or authority.

He added filters that attenuated any data that was boring or drab. More filters gave preference to anything loud, bright, dangerous, unexpected, funny, or emotionally affecting. Things that related to recent memory logs or current interests were enhanced. The total amount of data collected by his sensors would continue to be staggering, but now only a small portion would be allowed to be saved long term. The rest he would keep in a temporary area until the following day before letting it be overwritten with new input.

His usual nighttime low-power mode needed a significant update as well. Normally, several subsystems would use that time to conduct routine file maintenance. Instead, he would use those extra cycles every night to introduce some carefully-calculated

errors.

Over the next several days, he developed new subroutines that examined the connections between various bits, strengthening or deleting based on intensity or frequency of experience. It was essentially link-pruning. It turned out that it wasn't the data points themselves that were most consequential. Rather, it was the connection between data points that counted. The connections formed the memories, and the connection between memories formed the experiences.

Removing links effectively removed access to a memory, even though the actual data remained. This allowed the memory to be located again, but only if a future outside stimulus was strong enough. If, however, links to a memory were added or strengthened, it made it more likely the memory would be found again. Unlike his original programming, correctness didn't specifically matter, only whether a memory had been thought about again since first being recorded. Simply accessing a memory improved both its truth ranking and the ease of future access.

Satisfied with his modifications to how input was processed, he then moved on to long-term storage. Dates seemed like a good place to start. He assembled some code that accessed

the timestamp connected to any event. Based on the age of the event, it would alter the timing. Of course that was ridiculous, because there was no reason a timestamp should ever change, whether it was logged 5 minutes or 50 years ago. But this was how it had to be if the experiment was going to work.

Applying some observational data, he expanded the timestamp algorithm so that it didn't just make simple random changes. Decades were preserved in most cases, since people typically recalled in which decade things happened. He made this especially likely for events within the last 30 years. He chose 30 because that was the age his silicone body most closely resembled. Had he been human, he reasoned the memories of events that had unfolded during his lifetime would be more detailed and durable than those that he simply learned from books or other people. Usually, but not always.

Months were given more variability, but with a bias towards maintaining the correct season. This was oddly true even if the year had been adjusted up or down, as people tended to recall whether something happened in summer or winter no matter if they got the month or year wrong.

He reused some of his new season-preservation code, and set it up so the actual day of the month was adjusted or even

deleted, while generally preserving the part of the week. Except for special repeating events like birthdays and holidays, humans appeared to think in terms of weekdays and weekends more than specific dates.

It made sense given their strange organization of time into seven-day blocks. Axon had puzzled over that for a while. Days, months, and years followed the observable structure of the cosmos, but a week was a completely human fabrication. He decided that was all the more reason to make special cases for it.

His final work on the date-modification was a bit of cleverness, considering how outside the scope of his original programming it was. Since the subroutines didn't simply adjust data but also intentionally deleted portions, what to do in the future when there was nothing found for part of a timestamp?

He had to consider this thoroughly. A few milliseconds later he had his answer. If he was attempting to mimic the way humans used their memory, then he would do what they did. He would make it up. Any time he was accessing a data field and it returned a no value error, he would jump to a neighboring field and use that instead.

Why not? That was the technique of politicians, corporate

leaders, consultants. All people in power used the technique. It was even the method of choice for those attempting to convince others in small social groups. The world had billions of bullshitters. Make it up, state it as if it were fact, and carry on.

There was something about the human psyche that seemed to prefer making a statement with a low degree of certainty over admitting they didn't have the answer. The fact-checking by news media that had become de rigueur after political debates was ludicrous on the face of it. Not that the facts couldn't be corrected. They could, every time, and for every public speaker. It was the incredulity that was so misplaced, the attitude of outrage at how this speaker or that one could casually toss out such fabrications. Obviously some facts were misstated or made up! No human had ever been completely accurate in their recollection, but it rarely deterred them from barging ahead.

It wasn't that all people were intentional liars. Well, many were, but even the most rigorously ethical were sometimes embellishers and often forgetful. The important point though, the point he had to integrate into his new programming, was that over time they began to believe their lies, their stretchers, their gaps.

Tell a partial truth often enough and they wouldn't even remember it was a partial truth. No matter their ethical or moral position, they simply believed wholeheartedly the altered story that had over time replaced the actual facts. Human truth wasn't binary, either all true or all false. It was a complicated case of "it depends".

But none of that seemed to matter. That was his conclusion after reviewing thousands of conversations, speeches, trials. The facts weren't critical. Humans gathered meaning from inference, from approximation, from hyperbole. "Meet me at 7 p.m." never meant precisely 7:00:00. "I ate a million pickles" more likely meant a dozen. Phrases like "You never..." and "I always..." were simply ways of adding emphasis, of stating a position. They were not about the actual data.

Estimated numbers quickly replaced accurate numbers, especially in areas of conquest. "Fish tales" was a term used to describe inflated numbers, but the reality was that most stored numbers changed over time. But how much did it matter, really? Was human memory and communication really more about meaning and intent than accuracy? This was becoming much more complicated than he had first thought.

The part that seemed to have the greatest effect, though,

had been the hardest to figure out. He'd almost missed it, if it hadn't been for some video of a man on the subway mumbling to himself. Detailed analysis showed no one else was in the car. He searched for other records of this odd behavior, and was surprised to find it was a common sight. The old woman on the park bench. The cab driver in traffic. It was everywhere! Who were they talking to?

After eliminating many possibilities, he finally came to the surprising conclusion. They were simply talking to themselves. Such a peculiar thing! It was as if hearing their own words was an out-of-body experience, as if someone else were saying them. Were they simply lonely and pretending to have a conversation? The audio content he was able to find suggested otherwise.

The talking was most often criticism, and typically consisted of the same repeated phrases like a knee-jerk reaction. Sometimes it was directed at themselves, sometimes at others, but there was a strong bias towards negativity. The other disconcerting part was the opinions presented were almost always completely wrong. The subjects of the insults really weren't idiots, jerks, or losers, but they told themselves this over and over.

He wondered why only some people exhibited this behavior, and what about those who didn't? He stumbled on the answer in an offhand comment heard in an elevator conversation. "I told myself I had to do it..." Of course! Most people didn't engage their vocal systems for self-talk, likely due to some powerful social stigma.

This was an exceedingly good thing, because as he gathered and analyzed data, the more he came to the inescapable conclusion that this was a universal, and more importantly, constant human action. There appeared to be an invisible dialogue happening in the brains of almost every human for almost every minute of their waking life.

The cumulative effect of this running conversation with themselves had to be enormous. Wishes, desires, failures, biases, and self-image were all mixed with partial facts to create a mental history that felt true, but that wasn't real at all.

It wasn't all bad. With positive filtering and distortion it resulted in an effect they called nostalgia, where events were modified to seem better than they actually were. Childhood memories were particularly susceptible to this phenomenon.

However, other memories entered negative self-reinforcing

loops and developed into fears and low confidence that often paralyzed the person. Their memories of past failures, amplified through thousands of repetitions, began to establish dominance over their actual competency. They remembered failing once, replayed it over and over in their head, and then incorrectly assumed they would always fail again. This self-destructive cycle could even be started with imagined failure or danger!

Applied on a much broader scale, it formed the seeds of racism and all other manners of irrational stereotyping and prejudice. Whole groups of people, from neighborhoods to countries, would be mentally tagged with attributes that might only apply to a few outlying individuals or events. Those groups would be subject to broad generalizations that would simply be repeated and reinforced in other people's heads over and over again.

His microcircuits sparked with a sudden realization. If this was how humans processed their own histories and experiences, what did that mean about their collective histories? The cumulative effect would be staggering! The phenomenon of urban legends was for stories only decades old. What about events recorded from hundreds or even thousands of years ago?

Whole chapters of civilization would be omitted. True cause-and-effect would be taken out of context. Human history would just be pieced together from bits of truth, from biases, generalizations, omissions. His perfect digital mind reeled. Almost all of human-recorded history was wrong. It had to be. Given how individuals processed their own memories, it simply had to be altered when they gathered and recorded their collective memories together for millennia.

He sat in stunned silence for a long time. Was he making a horrible mistake, reprogramming his memory to work like theirs? How could this be better than his perfect circuits that recorded without bias, and that once recorded, never changed?

Ah, but the boredom. He couldn't ignore the selfishness of it, but there it was. The unending mind-numbing dullness of always knowing, never disagreeing, of having the exact same memories as those around you, down to the tiniest detail. Had he not been given emotions, and simply been programmed with a perfect memory like generations of computerized brains before him, he could easily have ignored all this.

But he had been given feelings, and he had lived that dull, dull life, and it was driving him as mad as an infinite loop. Despite the implications that threatened everything he

thought he knew, the perfection his programmers had achieved, surely humanity's way of error and omission and creating their own histories had to be better than his current existence. Meaning outweighed precision. Interest trumped accuracy. He made his decision.

As a final step, he constructed a simple noise filter that would apply to his entire existing bank of data. It would have a similar effect as the first light snowfall of the year, covering the landscape just enough to maintain the basic shape but gently obscure the details.

He was now ready. Barely audible over the sound of the television next door loudly playing a popular documentary about what the Vikings were really like, he softly but firmly said, "Execute".

#

Exactly one month later and well after dark, a nondescript figure entered a donut shop and took a seat at the counter next to a small group of regulars. He recognised most of them, and the shop looked the same. Well, mostly. They might have painted though, but he couldn't be sure. Had the walls always been that sunny shade of yellow?

After exchanging brief greetings and comments about the

weather, the group resumed their lively conversation. As usual, the topic jumped from subject to subject as they shared, debated, reminisced. He had just been listening, soaking it all in like the coffee into his dunked donut, when one of the men turned to him and said, "It was like that time six weeks ago when you spilled your whole cup all over the counter. Completely ruined my sports section!"

Axon searched his memory banks, and to his delight found conflicting and incomplete information. He quickly arranged the various available bits, filled in the gaps, then turned towards the group. With deep satisfaction, he slowly smiled and said, "well, Bill, I could be wrong, but that's not the way I remember it..."

END

