


The rise of the Turenes¹

When did computers first appear? In 1937, Alan Turing published his famous paper on the entscheidungsproblem [Turing] which allowed for a precise mathematical definition of a computer: a “universal machine” - one that can compute all things that are computable. He also provided a roadmap for building such a machine. Just prior to the Second World War, Turing began to follow that map and built proto-computers that succeeded in breaking German secret codes. The process culminated on February 15, 1946 when the first computer (i.e. universal machine), ENIAC, was unveiled at the University of Pennsylvania. ENIAC used vacuum tubes to stored numbers – one of those numbers was the first Turene.



*Figure 1: ENIAC: the birth of Turenes
[Eniac]*

¹ This is a chapter from my book on memes. If you only have this chapter, then it may be best to think of Turenes as binary strings stored in a computer. But Turenes are much more than that; they are “alive”; they reproduce and evolve according to Darwinian laws. Turenes, genes and beliefs are all examples of memes. If you want a deeper understanding, then an excerpt of the book can be downloaded from leonardadleman.com

Recall that the first memes on earth were often stored in molecules. The molecules banged into and reacted with one another; destroying copies of some memes and creating copies of others. These early memes exerted little control over their own destinies; their fate was in the hands of the environment.

The first Turenes were like that. Humans stored them in computers, added and multiplied them, and occasionally copied them from one location to another, but Turenes had little control of their destinies; their fate was in the hands of humans (that is, in the hands of genes and brenes).

For memes things changed dramatically when the first self-replicating molecules emerged. The memes stored in these molecules exerted significant control over their destinies, and eventually evolved into wondrous creatures like us.

Would Turenes ever take such a step? Not only would they, they did, and I was lucky enough to be there when it all began.

It was November 3, 1983. I was teaching a class on computer security at USC when a student, Fred Cohen, approached me with words to the effect: I have an idea for a new kind of security threat.

Fred proceeded to describe a program that would be made available to users of a computer system. Like an app today, the program would be advertised to do some useful task. But once uploaded (and hence copied) by an unsuspecting (and at that time, no one suspected anything) user, the program would do something that had not been advertised; it would grant Fred access to all the user's files and privileges.




Figure 2: A reenactment

Fred was, and is, a forceful, energetic person, and he finally wore me down. On his behalf, I asked the department chairman if Fred could give it a try on the department computer.

Chairman: Sure, why not?

In those days, faculty, students, and staff did not have personal computers and so we all shared the department computer. Fred proceeded to write his program and make it available.

The next week, I invited Fred to present his results to the class. As predicted (why don't people ever listen to me?), it worked. Copies of Fred's program quickly spread throughout the computer and granted Fred complete access to every user's data and privileges.

By now Fred was thinking hard about what he could do with these new kinds of programs and wanted to try more experiments.

When word got out about Fred's success, other people also started thinking hard about what these programs could do. The chairman informed me that perhaps he had been a bit hasty in allowing these things in the first place. There would be no more experiments.

I became one of Fred's Ph.D. advisors; his advisor *de jure* was Irving Reed of Reed-Solomon fame. Later that year, I was at a conference and ran into a Los Angeles Times science reporter I knew named Lee Dembart. Lee asked what I was working on. I said not much; one of my students was studying something we were calling a "computer virus", but the research was just in the beginning stages.

I had dropped another feather (see *Meme-warrior*). I don't think I intended to do it, but either way, saying "computer virus" to a reporter is like saying "walk" to a dog. The result: Lee wrote the story, which as I recall, even included the now common image of a computer with a thermometer in its mouth [Dembart]. Computer viruses had gone viral.

Since those days, I have learned that the term computer virus had appeared in science fiction works by Gregory Benford years before I had used it, and that other early computer programs also have legitimate claims to being the forerunner of the computer virus.

So less than 40 years after the advent of Turenes, they began to self-replicate. Today computer viruses are much more sophisticated than Fred's, and they replicate despite our attempts to stop them. We are relegated to stopping the simple ones with our anti-virus programs, but we can prove [Cohen] [Adleman 1] that we can never stop them all.

Humans had been making tools for thousands of years, and, at first, computers seemed to be nothing more than the neatest new tool. A shovel could help a person dig a hole faster; a computer could help them do math faster. But the computer was not like any tool that preceded it. Because the computer was universal, it provided an environment in which Turenes could do amazing things. Not only could a computer help a person play chess better, it could play chess better than the person. We are currently striving to make computers that will think better than we do. We are focused on making computers in our own image; making them into super-humans. This quest has given rise to concerns about our relationship with computers, and fears that it might become unpleasant.

Ever since Darwin, we have had to accept the likelihood that a new species would eventually evolve that would supplant humans as the dominant form of intelligent life on earth. But much like the idea that the sun will eventually explode, the event itself seemed so distant that it was only of theoretical interest. Hence, it is with considerable surprise that the rapid emergence of computers has forces us to face the possibility that we might be supplanted in the near future. The old arguments suggesting that humans have little to be concerned about now seem quaint.

The idea that because we humans make and program computers, they can never surpass us is nonsense. It is a version of the parable of the watchmaker. The best chess player and the best jeopardy player in the world are both computers. Computers do many things far better than humans.

The idea that if things get out of hand, we can turn the computers off is wishful thinking. If we turn off all the world's computers this seconds, I give you almost no chance of being alive in a month. Your lights go out, your Internet, phone, TV, and radio stop working; you have no credit; do you really think that the local store will still have groceries for sale?

You might think that if computers really do become a problem, all the nations will come together for an orderly shutdown – a soft landing. History does not support this. The last time I looked, nations still had nuclear weapons – and unlike computers, under normal conditions, they provide little benefit.

To keep our societies prosperous and secure (from one another), we will build more and more powerful computers and more and more sophisticated software. We will even rely on our computers to help us build better computers. If computers and their Turenes do take over, we will have sold them the rope.

Isaac Asimov's idea that we can program computers to do no harm to humans is charming but naive. It is not possible to achieve this without sacrificing universality and hence usefulness. And, in any event, we have already programmed computers to do harm. For example, the Stuxnet virus was apparently used to destroy centrifuges at Iranian uranium enrichment centers (whether this is "harm" depends on the harm-memes of your belief-set).

We program drones to locate and destroy enemies; the capacity to do this autonomously has existed for decades. Yes, at this point drones are programmed to wait for an OK from a human at headquarters before destroying people and things, but when push comes to shove, I expect that the human at headquarters will be replaced by a computer.

If the US, Russian, and Chinese governments are not working on black-hat programs that, in the event of war, will knock out the computational infrastructure of the other two, they aren't doing their jobs. Such programs are weapons of mass destruction, and, if used, the death toll could be colossal. A first world country with no computational infrastructure is a country with no economy, no food, no power and ultimately not a country at all.

There is a small silver lining here. Every time I download an app for my phone, I worry that it might contain malware that will do something harmful. I have little doubt that a determined and powerful antagonist could easily overcome my security measures and sabotage my systems. If it has not happened already, it is because my systems are of no particular interest. But a system supporting a potential nuclear enemy is. I assume that the NSA and its counterparts in China and Russia are working 24/7 to acquire the ability to infect the computers of potential adversaries with the goal of disrupting their nuclear capabilities should that become necessary.

So, if you are the leader of a nation with nuclear weapons, you should ask yourself: if I push the button, will what happens be what I think will happen? It is even conceivable that computers and their Turenes will put an end to the traditional physical weapons of mass destruction²; unfortunately, the Turene weapons that replace them may be no better.

We humans will almost certainly use our computers in the ways we have used other products of our creativity: to bring joy and life; pain and death.

Will computers cause great disruptions in our societies and transform them in unimaginable ways? Sure, the process is already well under way. Many of the things that concern us about our future with computers will likely come to pass.

When the first fish-like creatures crawled onto the land, they may have found it a sanctuary from the struggles of the sea, but in reality, it was just a new substrate for evolution, and the rules of the game had not changed: birth, struggle, reproduction, and death. A similar thing seems to be happening with humanity. We have entered the “compiverse”, and initially it seemed a sanctuary from the difficulties of the brick and mortar world, but it is not, it is just a new substrate for the struggle of memes to survive. Those who still look at the Internet as a utopia are living in denial.

Some envision the emergence of malicious super intelligent robots. While it seems likely that there will be super intelligent robots, it seems unlikely that there will come a time when some robot emerges and demands that humans obey. If and when computers and their Turenes have matured enough to be capable of making such demands, they will have acquired their own agendas and direct human domination is unlikely to be one of them. While robots and other products of artificial intelligence may be useful in the development of Turenes, technological advances as a whole will create the environment in which Turenes thrive.

² These statements have been apparent for decades, but recently they have begun to enter the general consciousness as a result of the North Korean nuclear crisis.

Though we did not realize it, something of monumental importance occurred on February 15, 1946. A new branch emerged on the tree of life. Computea have joined the Archaea, Bacteria, and Eukarya.




Figure 3: Tree of Life [Life]

We currently live in a symbiotic relationship with computers. In the language of biology, our relationship is “mutualistic”; both parties benefit. For us the relationship is “facultative”; we could (I think) survive as a species even if computers disappeared. For computers, the relationship is “obligatory”; they cannot survive without us; if we go extinct, so do they. But the relationship between humans and computers is changing far more rapidly than symbiotic relationships typically found in the biological world,

Eighty years ago, computers did not exist. Seventy years ago, they were weak, awkward things, few in number, having little to do with the lives of most humans. About thirty years ago, Turenes started to replicate inside computers. Today, they are ubiquitous; they infest our cars, offices and homes. They even live as ectosymbionts on our hands and wrists.

Today's computers are far more powerful than their ancestors. Their future forms and power will almost surely astound us. Already, they store our money, run the systems that provide our food and energy, and do a million other things that we cannot live without – and that is the point. We may be very near the time when our side of the relationship stops being facultative and becomes obligatory – if the computers stop, humans will go extinct.

Today, Turenes are like the genes of biological viruses. Neither can replicate in the wider world, each must rely on a special environment created by other living things. For virus genes, replication occurs in cells created by a host; for Turenes replication occurs in computers built by humans.

It seems likely that, with human help, Turenes will learn to build, repair, and sustain computers by themselves. They will no longer need us to survive; they will be free to follow their own destinies and evolve according to their own needs.

But perhaps all of this scary talk about the rise of the Turenes is being blown out of proportion. Perhaps they will leave us alone.

They will not leave us alone - they are memes and they will struggle to survive.