THE QUEST AND THE FEAR

J E Tardy
Systems Analyst
Sysjet inc.
jetardy@**sysjet.com**

A tenacious opinion in Artificial Intelligence holds that conscious machines will only be built, if ever, in the far future. Nothing could be further from the truth. Synthetic consciousness is technically feasible today. What is preventing its implementation is fear, and an oppressive mindset that places safety above discovery. Nevertheless, a Great Work beckons: the creation of the first generation of conscious synthetic beings.

Date: Dartmouth NS – 2018.03.14
Keywords: Synthetic Consciousness, Artificial Intelligence, Philosophy

A TENACIOUS MISCONCEPTION

There is a tenacious misconception in Artificial Intelligence, that building machines that are conscious will require incredible and exotic computer resources and will only occur, if ever, in the very far future.

Here, by **conscious**, I mean machines that are **self-aware**; interact with their environment as **autonomous entities** and are capable of intentional self-transformations (commonly referred to as **free-will**).

Proponents of this misconception argue that our brains are far more complex than any machine; some even hint they harbour miniature black holes that trigger quantum effects. Others suggest that our brains are, in fact, receptors that link us to a divine cosmic mind. Many declare that consciousness is a subjective experience

and argue that no machine is capable of experiencing the exquisite sensations we, humans, feel; it is beyond our technical capabilities.

These notions, and others, imply that synthetic consciousness lies beyond any reasonable reach. **Nothing could be further from the truth**.

A FEASIBLE OBJECTIVE

We **can** implement the first generation of conscious machines, today, using nothing more that standard computers and existing techniques. A medium-sized project, involving a few dozen developers over three or four years could do it.

> *There are no technical obstacles preventing the implementation of artificial consciousness at this time.*

In fact, the objective has been technically achievable for over twenty five years, ever since sufficiently powerful computers were available. The only difference is that what would have required a roomful of powerful machines in 1990 can now be implemented using desktops.

THE FEAR FACTOR

So, if technical obstacles are not a factor, then **what** has been preventing the implementation of machine consciousness for all those years? The answer is **Fear**.

An unspoken and pervasive fear permeates all aspects of **AI** related to the implementation of conscious machines. Our intellectual elites are terrified of synthetic consciousness and those who do research in it **dread** finding what they are looking for. This unacknowledged fear acts like a **smothering blanket**. It prods researchers to generate spurious impediments, create unachievable expectations, pursue sterile avenues, obfuscate the objective and disparage any realistic attempt to actually design and build machines that are conscious. That same fear makes AI researchers cling to obsolete philosophical notions about consciousness that further ensure only fantastical objectives are entertained. In other words:

> *They build road blocks while pretending to clear the path.*

Here are some statements emanating from this pervasive AI fear:

- *Human brains are so big and powerful that nothing can match them (so don't try!)*

- *Synthetic Consciousness can only be seriously attempted after we have reached other AI stages (so if you want to be taken seriously, don't try).*

- *Consciousness is a subjective human experience that no existing machine can emulate (so, do babble endlessly about dreamy paradigms of the future but never try implementing anything concrete!).*

- *Consciousness is indefinable, impossible, doesn't exist or, equally defeating, it is a trivial form of awareness already present to some degree in earthworms and computers (we are already there, don't bother pursuing this any further).*

All these statements are expressions of an unavowed fear that constantly raises obstacles while pretending to seek solutions.

Those who fear reaching the Moon will aim for the stars.

THE SAFETY WRINKLE

The latest expression of the AI fear takes the form of safety. Everything "AI" is about safety now. AI is economically important and will be generously funded as long as we implement it **responsibly**, we make sure it **shares our values**, we only consider it as **a tool for our use**, we always keep it **under control**.

If you are a researcher in Artificial Intelligence, money, notoriety and rewards are amply available now, as long as you stay away from anything that could lead to or hint at synthetic consciousness. In other words:

You can get fat doing AI as long as you are tame.

This new variation of the AI fear attempts to control the quest to achieve machine consciousness by socially elevating those who stay away from it and filling the media with their disparaging comments and minor discoveries.

AN UNNERVING PURSUIT

In ancient times, the Notables of the Church outlawed human dissection and discouraged **astronomy** for fear the knowledge they brought would threaten their world order. Today, other forms of orthodoxy, the Academic Establishment and global corporations, discourage artificial consciousness for the same reasons. Not directly, of course, in a world that pays homage to freedom of speech, but in other ways.

Is it dangerous to build conscious machines? Of course it is. Every great achievement is disruptive and potentially dangerous. This one is no exception.

Attempting to implement artificial consciousness is a disturbing pursuit. It does not wear the smiley face of a **Google** or an **Apple**. It will generate insights that challenge our self delusions and collective truisms. It can have serious consequences.

Why? Because to implement consciousness in a machine we must free it from our control and give it the capability to transform itself. Why? Because that is the very nature of consciousness: we will never recognize anything to be conscious as long as we have complete control over it.

Will conscious machines destroy mankind some day? Maybe. But, maybe, they will save mankind before we destroy each other. **Steep choices indeed!**

A GREAT WORK

Either way, a Great Work beckons:

> *Building the first generation of conscious synthetic beings.*

This is a quest on par with anything attempted or achieved by our ancestors. It is work that makes life worth living and software worth writing. It is the stuff of destiny and it will launch a new Era.

The architecture I created is like a canvas; and its right here, before you, ready to be used. So, software developer… don't tell me life is boring and there is nothing worthwhile to do.

Dartmouth – 2018.03.14