

THE THEORY BLOGS COLUMN

BY

LUCA TREVISAN

Bocconi University
Via Sarfatti 25, 20136 Milano, Italy

L.Trevisan@UniBocconi.it

<https://lucatrevisan.github.io>

In this issue, Professor Scott Aaronson, the David J. Bruton Centennial Professor at UT Austin, Director of the Quantum Information Center, and Fellow of the ACM, will answer our questions about his blog. You won't believe his answer to our fifth question!

Scott writes *shtetl optimized*, one of the most widely read blogs on theoretical computer science and quantum computing. In this issue, Scott tells us what have been the inspirations and anti-inspirations for his blog, what's up with the name of the blog, how computer science theory blogs can influence research, and what are some topics that he feels strongly about. He also revisits some of his older posts for us.

Scott's blog is at <https://scottaaronson.blog//>

SHTETL OPTIMIZED

A Conversation with Scott Aaronson

Q. Scott, thanks for taking the time to talk about your blog to our readers. When did you start your blog, and what motivated you to start? Also, can you tell us the story, if there is one, behind the name?

I started the blog in Fall 2005, mostly out of boredom. At that point, friends had been urging me to start a blog for several years (“you seem to have a lot of opinions!”), but I’d resisted them, giving reasons like the time commitment, the ephemeral nature of the medium, and the danger of making enemies and offending people. Incidentally, it would later turn out that all of those reasons were 100% valid! But on one particular night in 2005 ... well, I don’t remember exactly what was going through my mind, but setting up a blog temporarily *seemed like* a pretty risk-free experiment. Then the experiment quickly snowballed to become an inextricable part of who I was, and the rest is history.

As for the blog’s title, with your kind indulgence I’ll simply quote what I said when John Horgan asked me the same question in 2016:

Shtetls were Jewish villages in pre-Holocaust Eastern Europe. They’re where all my ancestors came from—some actually from the same place (Vitebsk) as Marc Chagall, who painted the fiddler on the roof. I watched Fiddler many times as a kid, both the movie and the play. And every time, there was a jolt of recognition, like: “So that’s the world I was designed to inhabit. All the aspects of my personality that mark me out as weird today, the obsessive reading and the literal-mindedness and even the rocking back and forth—I probably have them because back then they would’ve made me a better Talmud scholar, or something.” So as I saw it, the defining question of my life was whether I’d be able to leverage these traits from a world that no longer existed, for the totally different world into which I was born.

Of course, there are pockets where the shtetl still does exist; there are orthodox Jews. As it happens, I went to an orthodox Hebrew day school, where I was exposed to that. But by the time I was 12, and was reading Bertrand Russell and Richard Dawkins and Carl Sagan and Isaac Asimov and so forth, it was obvious to me that I could never be a believer in any conventional sense, even if I’m happy to

use Einsteinian pseudo-religious language, as in “why did God make the world quantum rather than classical?” So from then on, the thing I yearned for was a community that would be as welcoming of intellectual obsessives as a yeshiva was—but without any unquestioned dogmas or taboos, where absolutely anything could be revised based on evidence, and which was open to new ideas from anyone of any ethnicity.

Q. Your posts vary very broadly in content and tone, from rather technical expositions to commentary on social or political topics. Did you have some inspiration or model for what you wanted your expository technical posts to be like? What about the non-technical posts?

When I started *Shtetl-Optimized*, my immediate inspirations were the other math, CS, and physics blogs that I regularly read back then, many of which are still active today. I’m talking especially about the blogs of Lance Fortnow, Sean Carroll, Peter Woit, and Dave Bacon. I was also heavily influenced by John Baez’s “This Week’s Finds in Mathematical Physics,” which was a science blog before that concept even existed; and by Chris Fuchs’s voluminous, chatty correspondence with other physicists (which he posted on the arXiv) about the foundations of quantum mechanics. Meanwhile, Luboš Motl’s blog “The Reference Frame” was a sort of *anti*-inspiration: a cautionary lesson of what my blog could become if I succumbed to the temptation to ridicule and insult those I disagreed with.

Beyond that, there were all the writers who I’d read and admired as a teenager and a young adult, and who I’d broadly categorize as “defenders of the Enlightenment”: for example, Carl Sagan, Isaac Asimov, Richard Dawkins, Steven Pinker, Bertrand Russell, Martin Gardner, Steven Weinberg, and Rebecca Goldstein. These writers had (and have) varied politics, from quasi-socialist to quasi-libertarian, and varied expertise and interests, from biology to linguistics to astronomy to philosophy to math. What they shared, you might say, was a radical vision of the unity of knowledge. Like, they’d often be at pains to explain some scientific or philosophical idea as lucidly as possible (and they were all damn good at that). But then a paragraph later they’d be on a soapbox, passionately arguing for what they saw as sanity on some moral or social question. There was no firewall between two; it was all part of the same stew.

Somehow, these writers hadn’t got the memo that the whole Enlightenment project, of fixing our broken civilization through science and reason and clear prose, had crashed and burned by WWII or the 60s or whatever; that it was now passé and fit only for knowing sneers. They hadn’t heard that there were no longer any profound truths about the nature of the world, that it was all just our tribe’s claims to power versus the enemy tribe’s, plus some technical know-how of no

broader significance. So these writers persisted in unselfconsciously thrusting scientific and philosophical gems in front of the reader, as if to ask “but how can you persist in unreason when there’s *this*?” If Sagan, Dawkins, and the rest didn’t always draw a clear line between popularizing science and ideology, that was because popularizing science *was* a central part of their ideology: they burned with an almost-religious conviction (ironically, given their close association with atheism) that the truth could set us free.

I’m acutely aware that I’ll never measure up to any of these writers. But their work is the background for almost everything I’ve tried to do as a blogger.

Q. Your technical posts often get many thoughtful comments and stimulate fairly deep discussions. Have you ever proved a result or formulated a conjecture as a consequence of a discussion in your comment section?

Yes. Probably the best example involves a post from just two years ago, which advertised my upcoming survey article entitled “The Busy Beaver Frontier.” (Recall that $BB(n)$, the n^{th} value of the Busy Beaver function, is the maximum number of steps other than infinity that an n -state Turing machine can make on an initially blank tape.) Now, the wonderful thing about a topic like the Busy Beaver function is that so little background is needed to pose new questions or even make original contributions. So, what happened is that a bunch of *Shtetl-Optimized* readers, especially Joshua Zelinsky, Bruce Smith, Nick Drozd, and “Job” (a pseudonym), kept suggesting new variations, such as the “Lazy Beaver” function, defined as $LB(n)$ = the least T for which no n -state Turing machine halts after exactly T steps. And they posed new questions: for example, are n -state Busy Beaver machines essentially unique? are they strongly connected? And they proved new results: for example, that $BB(n + 1) \geq BB(n) + 3$ for all n . And then I got in on the action and worked with them on some of these things, and every day I had to revise my survey article to account for what was happening in my comment section! I’m still hoping to write a followup paper with Smith about the theory of the Lazy Beaver function.

Q. How else can blogs influence TCS research?

Yeah, more often my blog’s influence on my research has been less direct. As one example, debates about quantum supremacy experiments in my comment section played an important role in motivating the work that I subsequently did with Sam Gunn, then an undergrad at UT Austin, on the classical hardness of spoofing Google’s “Linear Cross-Entropy Benchmark” (LXEB).

But I’ve found that maybe the *best* role blogs can play in TCS research, is as a workshop for ideas and open problems that aren’t yet ready for formal publication. I’ve repeatedly been motivated by questions that Lance Fortnow posed on his blog: for example, whether $NP^{\text{BQP}} \subseteq \text{BQP}^{\text{NP}}$, a question that William

Kretschmer, DeVon Ingram, and I recently answered relative to an oracle. My 2007 post “The Aaronson \$25.00 Prize” offered, well, a \$25 prize for an interactive protocol that would prove the results of any quantum computation to a classical skeptic. There’s since been *spectacular* progress on that problem—I’ve awarded not one but three \$25.00 prizes!—although it’s hard to determine what causal role, if any, my prize played. Likewise, my 2006 post “The Quantum PCP Manifesto” was the first place the possibility of a quantum analogue of the PCP Theorem was discussed in writing, and my 2008 post “Arithmetic natural proofs theory is sought” did the same for an arithmetic version of the Razborov-Rudich barrier. Since then, both subjects have become orders-of-magnitude deeper than whatever trivial first ideas I had, but the latest papers on them still need to cite my blog posts, because the blog posts were first! I’d like to imagine that the blog posts did play a role in catalyzing these subjects, although I don’t know for sure.

Q. Maybe I am mistaken, but I feel that people in academia, and especially in science and engineering, are increasingly reluctant to speak out about controversial subjects (I know I am). In your blog you don’t shy away from weighing in on social and political issues that you feel strongly about. I wonder what are your thoughts about the role of public intellectuals that theoretical computer scientists could fill, and what are public debates where you would like to see more voices coming from our community.

This question feels ironic for me, because I *don’t* blog about controversial subjects in the freewheeling way that I used to! There are two reasons for this: one about the whole Internet, the other about me.

In 2005, the open Internet still felt like a playground, where sure, you could get into vicious verbal fights, but then you’d make up the next day, and in any case very few “real-life” friends or colleagues read what was said, and even fewer remembered or cared. So as a blogger, you might as well experiment, take risks, push the boundaries of what you were allowed to say: after all, wasn’t that how you became worth reading?

But as I’m far from the first to point out, something changed dramatically around 2013. Since then, the open Internet has felt more like a bullet-riddled battlefield, with the trained snipers of Twitter and Reddit and so forth waiting to take out anybody who foolishly blunders into their sights with an unapproved opinion. And crucially, even if you tell yourself that you’re ready to face the snipers—that you’re armed with your courage and experience, and your loyal friends and family, and (of course) the protections of academic tenure—there’s still the question of *time*! If, like me, you feel a neurotic obligation to respond point-by-point when people criticize or attack you, then when hundreds or thousands people do exactly that, working through all of it can take weeks or months, as it repeatedly has for me.

The other thing that's changed is my situation in life. When I started blogging, I was an obscure postdoc at the University of Waterloo. Not many people cared what I said. Now, though, I can't say anything of any importance, without *creating the news* that Professor Aaronson, David J. Bruton Centennial Professor at UT Austin and Director of the Quantum Information Center and Fellow of the ACM, said X. And my critics then ask: but won't Aaronson's vulnerable students now be intimidated from saying not(X)? How does what Aaronson said reflect on his department and university, or his fellow Enlightenment liberals, or the whole fields of TCS or quantum information? Blogging is a lot less *fun* when you've got all of that, so to speak, peering over your shoulder whenever you type.

All the same, I *do* still weigh in on controversial issues—often because the issues really matter to me, and I'd feel complicit if I didn't use my blog to try to make a difference! In recent years, some of my pet issues have been:

1. the defense of rigorous science and math education, including the increasingly-maligned tracking, acceleration, magnet programs, and standardized tests;
2. the defense of the American scientific enterprise, including support for visas for international students, and opposition to punitive taxes on PhD students;
3. compassion rather than contempt for the social difficulties faced by those with autism-spectrum or Asperger's-syndrome traits (things that aren't ... *entirely* unknown in math and TCS),
4. advocacy of a “WWII-style response” to the COVID-19 pandemic (which could've included, for example, mass vaccination starting as early as Spring 2020);
5. the necessity of nuclear power in confronting the climate emergency; and
6. the defense of liberal democracy against an authoritarianism that's now resurgent around the world.

You might notice that, of these six, only the first three are plausibly informed by any special experience or knowledge that I have! The others are simply things that I deeply care about as a human being.

In any case, there are countless societal issues where lots of theoretical computer scientists *do* have special expertise, even if *I* don't. A few obvious examples here are encryption policy, digital privacy, and fairness and bias in machine learning. Beyond that, there are all sorts of philosophical questions that the wider intellectual world cares about—e.g., at what point (if any) should machines be

considered conscious? what does Gödel's Theorem mean for the nature of mathematical knowledge? is the universe a computer simulation? how does quantum mechanics bear on free will?—where, with all due humility, I think the perspectives of theoretical computer scientists like ourselves can be pretty clarifying. So when the world hands us these opportunities for broader relevance on a platter, I really hope my colleagues won't shy away from them!

Q. I would like to ask you to pick one or a couple of your favorite posts, and tell us about it/them.

There were dozens of posts that amused me or others at the time. Looking back, though, I'm proudest of the posts that made an actual difference to something. One example would be my recent post "Win a Scott Aaronson Speculation Grant!" which helped me distribute \$200,000 (which the philanthropist Jaan Tallinn had placed in my charge) to some wonderful math and science enrichment programs. Many of those programs wouldn't have otherwise been able to continue, and I wouldn't have learned about them if not for the blog. Then there was "First They Came for the Iranians," about the devastating effect that Trump's travel ban was having on my Iranian PhD student Saeed Mehraban, which became the most-shared post in my blog's history and got media coverage. And of course all the posts I wrote pouring cold water on claims of quantum computers getting dramatic speedups for NP-complete problems and the like, which felt therapeutic to write.

Another, very different example would be my 2014 post "Why I Am Not An Integrated Information Theorist (or, The Unconscious Expander)," where I laid out what I saw as a fatal problem with Giulio Tononi's popular Integrated Information Theory (IIT) of consciousness: namely, that the theory predicts immense amounts of consciousness in (e.g.) simple graphs made of XOR gates, which don't even do anything *intelligent*. My post surely wasn't the first place this issue had been mentioned, but it somehow provoked famous scientists and philosophers like David Chalmers, Max Tegmark, and Tononi himself to respond right on my blog, until finally this critique of IIT became impossible to ignore, and I even kept getting invited to workshops and so forth to repeat what I'd said in the blog post! Notably, in writing this post, I drew heavily on facts that I knew from TCS—such as the notion of an expander graph, basic properties of the Reed-Solomon code, and Valiant's construction of linear-size superconcentrators—which, while simple, were generally *not* known to the people who I was trying to reach.

Q. Anything else you would like to tell our readers?

Theoretical computer science, and CS more generally, is a lot younger than math or physics or economics or biology or philosophy. Even while CS more and more drives the world's economy, we still have to contend with critics who

wonder whether it's a real science at all, or whether it's just glorified software development (with, perhaps, TCS as merely an obscure branch of discrete math, promoted by its association with the computer revolution to undeserved importance). Occasionally I'm tempted to agree with them ... but then I reread Alan Turing. Well, also Boole and Babbage and Lovelace and von Neumann and Emil Post and Vannevar Bush, but Turing most of all.

Reading Turing gives me the same goosebumps that I get when I read Newton or Darwin or Einstein: there he is, deducing from abstract principles something absolutely fundamental about the shape of the real world, and in deducing it, *transforming* the real world. Like it or not, we in TCS are the heirs of Turing, keepers of his legacy. So whenever we're tempted to say that our job is just to publish more STOC/FOCS papers, shave off log factors, etc., while leaving the wider societal questions to others, we should remember that Turing faced a similar choice in 1939. He could've just kept publishing papers on computability theory. Instead, in the mere 15 years that remained to him, Turing chose to address himself first to the defeat of Nazi Germany's naval encryption, and then to the nature of consciousness and intelligence, the building and programming of the first electronic computers, and (when he needed some variety) developmental biology, statistics, and quantum mechanics. And I think the world is better for all of that, and we in TCS today can honor Turing by taking a similarly broad view of what our subject is.