

TCS ON THE WEB

BY

STEFAN NEUMANN

TU Wien

Erzherzog-Johann-Platz 1

1040 Vienna, Austria

stefan.neumann@tuwien.ac.at

<https://neumannstefan.com>

Clément Canonne is a senior lecturer at the University of Sydney. He works on distribution and property testing, the theory of machine learning, and differential privacy. He recently received the 2024 Caspar Bowden Award for Outstanding Research in Privacy Enhancing Technologies (joint with Gautam Kamath and Thomas Steinke) and was a co-organizer of the Sublinear Algorithms program at the Simons Institute in Berkeley, California.

Clément is also one of the most important theoretical computer science influencers on Twitter/X with more than 30,000 followers.¹ In his guest column, he answers our questions on how he started tweeting (or “microblogging” as people used to call it) about theoretical computer science (TCS), and how we can use social media to promote TCS to a broader audience.

¹https://twitter.com/ccanonne_

PROMOTING THEORETICAL COMPUTER SCIENCE ON SOCIAL MEDIA

A Conversation with Clément Canonne

Clément, thank you for taking time for this interview. You do not have a blog, but you are one of the most-influential theoretical computer scientists (TCS) on Twitter/X with more than 30,000 followers.¹ So, I thought featuring you in this column would be a good idea. You joined Twitter in 2011, do you still remember what motivated you to go there? What has changed since then?

I think I joined Twitter on a whim, not being very clear what that was good for – then several years passed, during which I forgot about my account. I became more active around 2019 or so, as it seemed some TCS researchers were posting there; and then, the pandemic happened, and I started spending a significant amount of time on Twitter. It was a way to stay connected with other researchers, and what was happening, and to have at least some meaningful TCS-related interactions and discussions beyond my living room at the time.

Now, a lot has changed, for clear reasons: first, the change of ownership of the platform, which has affected the form and content in many ways. Another more personal aspect is my move to Australia, which has meant a slightly different experience: now including people from a different continent, but also a change of rhythm and content I see and gets brought to me by Twitter, due to the time zone. It's an interesting shift.

Is there a reason why you share your thoughts on Twitter/X and not in a more classic blog format?

Good question! I really enjoy some of the TCS blogs (e.g., Boaz Barak's "Windows On Theory", to name one among others), but I think I lack the discipline to write long form in a regular fashion. I also think Twitter has (or used to have?) the appeal of conciseness, which is a valuable exercise in itself: how to convey a point or idea in very few characters. I'm not claiming all the ideas or points I'm conveying are good, or worth writing, though... I don't think I'd have

¹https://twitter.com/ccanonne_

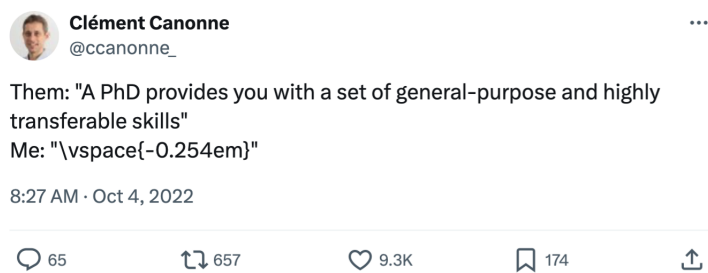


Figure 1: One of Clément’s viral tweets.

much credibility there, given that my Twitter account is roughly 50% puns, 50% technical content, and 50% dubious math.

For what it’s worth, I tend to like long form writing in a different context, that of writing lecture notes or surveys; and I am an occasional writer on blogs such as <https://differentialprivacy.org>.

Are there some posts or interactions that you would like to highlight?

Not necessarily the most informative, but I remember once waking up and making a silly CS-related pun. I think it was about LaTeX, something about negative `\vspace` (see Figure 1, the editor). It went sort of viral, I guess LaTeX woes are universal in our field: three weeks later, I received a message from a friend from undergrad in France about it, telling me my name popped up in a conversation with a colleague. Not sure what to make of that, except that I didn’t have “being talked about in my home country because of a LaTeX tweet” on my bingo card.

Maybe more informative though is the fact that having some reach (moderate, I won’t kid myself, but still) can be useful – for instance, I want to believe that repeatedly highlighting the Australian student visa issues on Twitter (back in 2021–2022) did make a difference. It definitely ended up gathering some amount of media attention, and (hopefully?) helped a little.

A lot of theory researchers mostly use social networks to promote their own research and new results in their own area. However, your content has a much broader scope, ranging from classic textbook algorithms, over probability puzzles to new results in your area. Do you have a specific process of what you post about?

I wish I had! Honestly, I feel like most of my Twitter posts come from various things I read, or hear about (talks, lecture notes, textbooks, interesting nuggets in papers) and want to share. There is no real process beyond the common thought

that “if I had known that 10 years ago, I’d have liked it.” I have a tendency to like neat, self-contained “tricks” or results, and my working assumption is that others will want to learn about them too. (There’s also a whole other part of my posts that comes from jokes or puns which pop up, also often reading papers or textbooks. No real process there either...)

Has your active presence on Twitter/X helped you to find new research ideas or even new co-authors?

It has definitely made me think more deeply about some questions: for instance, after posting something and seeing the comments and discussions, realizing that one assumption in the phrasing was not obvious, or at least less natural than I thought. I guess it’s part of a broader picture related to science communication in general, where we don’t fully see our own thinking patterns and biases until we have to explain an idea to a broad audience.

In terms of research ideas, the main advantage of Twitter has been the exposure to other people’s research. It’s hard to keep track of all the interesting stuff that happens, but people actually explaining their new paper in a few tweets (which is hard!) is a very good way to at least get some of it. Recently I ended up using a result from a paper I’m confident I’d have completely missed had it not been discussed on Twitter: the title would not have rung a bell, the abstract wouldn’t have seemed relevant, and I probably would have never read it.

Sometimes my impression is that in TCS we have a lot of great results, that would deserve more visibility in the general CS community. Do you have some suggestions to other TCS researchers on how we can improve our outreach on social media?

That’s quite tricky, though related to the previous answer. I don’t really have a good idea beyond stuff that worked for me somehow, and things I enjoy from others. Generally speaking, I enjoy to learn things, even if they might seem simple: one “cool trick” from a paper goes a longer way than a very technical and overwhelming statement. And it’s nice to spend time mentioning and discussing such great results even when they’re not yours (especially when they’re not): not to say there cannot be any self-promotion on social media, of course! But at the end of the day, the goal is to make as many people as possible aware of all the incredible stuff there is, regardless of whom it came from.

Before we conclude the interview, could you please let us know what your current research is about?

One thing that has been on my mind for a few years now is what happens to a lot of classical statistical or computational learning questions (including “well understood” ones) when you see everything as a “resource”. The first resource we learn about in CS undergrad is time, and that’s the first foray into algorithm analysis. But memory is a constraint as well, and communication in distributed settings, and randomness, too. And once you’ve started, you realize you can view privacy as another, and robustness to adversarial corruptions, and maybe even the type of “queries” that are allowed, and how much adaptivity your algorithms can have. Now you have a range of constraints and resources, and analyzing their interplay and the various tradeoffs is fascinating, mathematically and algorithmically. There’s so much of it we don’t understand, even for the most basic tasks!

Finally, is there anything else you want our readers to know?

Wombats are really cute animals. Seriously, they’re like little fluffy marsupial juggernauts.

Oh, also, and on a more serious note: if you are not already, please check out the TCS Blog Aggregator at <https://theory.report/> to stay informed about TCS-related papers, blog posts, announcements, and job postings. I’m not affiliated with it, it’s just a great resource.

Thank you for this nice interview!