The Abacus Medal 2022: Mark Braverman

By Rachel Thomas, produced as part of the ICM coverage on plus.maths.org

Mark Braverman of Princeton University has been awarded the 2022 Abacus Medal at the International Congress of Mathematicians. The Abacus Medal is awarded for "outstanding contributions in Mathematical Aspects of Information Science". The Medal used to be called the Rolf Nevanlinna prize, and is awarded every four years at the International Congress of Mathematicians.

Braverman won the prize for his "path-breaking research developing the theory of information complexity." He said his overall goal is to understand computation, both because it is a fundamental intellectual pursuit, just like studying black holes or prime numbers, but also because computation is now embedded in our daily lives. "Computers and communication are now so cheap they're part of many devices - even a toaster is a computer!"

Mark Braverman (Photo Lance Murphey)

We were lucky to speak to Braverman in the run-up to this year's Congress, which is held as a fully virtual event with only the prize ceremonies and lectures taking place in-person in Helsinki, Finland. He told us about the role of communication in computation, and why a mathematical view can help you understand how to solve problems while sharing as little information as possible.

From a mathematical theory of communication

"Shannon's classical information theory is a beautiful area that explains communication," says Braverman. "There are still open problems, but communication is mathematically extremely well understood."

Braverman is referring to the mathematician Claude Shannon, who developed many of the key ideas used in digital communication today in his groundbreaking 1948 paper, *A mathematical theory of communication*. Shannon realised that binary digits, better known as *bits*, lay at the heart of information technology. Any type of information, be it pictures, music or words, can be encoded in strings of these 0s and 1s. Shannon worked out the minimum

number of bits you need to encode the symbols from any alphabet, be it the 26 letters we use to write with, or the numbers that encode the colours in a picture.

Given an alphabet of n symbols and a probability distribution telling you the probability p_i with which a symbol i occurs in a text made out of those symbols, the number

$$H = p_1 \log(1/p_1) + p_2 \log(1/p_2) + p_3 \log(1/p_3) + \ldots + p_n \log(1/p_n)$$

is called the *entropy* of the distribution (see this article to find out more). Shannon proved that the average number of bits needed per symbol cannot be smaller than the entropy, no matter how cleverly you encode them.

If you're in the business of sending messages long-distance, then Shannon's entropy is a useful number to know. If you know you can transmit some number, say *C*, bits per second, and that the symbols in your message require around *H* bits per symbol on average, then you'd guess that you can transmit around *C/H* symbols per second on average. Shannon showed that this is indeed correct: by picking a clever way of encoding your symbols you can guarantee an average transmission rate that's as close to *C/H* per second as you like. Conversely, no matter how you encode your symbols, it's impossible to transmit at a faster rate than *C/H.* This neat fact is known as Shannon's *source coding theorem*.

To a mathematical theory of computation

The clear mathematical framework developed in information theory has led to communication being very well understood, with clear mathematical bounds on things like how much information you can transmit.

"You can use Shannon's entropy to define channel capacity," says Braverman. The *channel capacity* is the maximum rate at which information can be communicated over some medium, such as via email, over the radio or downloading to your phone. "Mathematically it makes it easier and you get very precise answers to questions like 'How long would it take me to transmit a billion bits over this medium?' You divide a billion by your channel capacity and you get your answer."

In contrast, a similar mathematical framework is not nearly so well developed for studying computation. Braverman compared the precise answers available for communication, to the open fundamental questions in computation, such as *computational complexity*: how long a computational task, such as factoring a number, will take. Not only do we not have any algorithms that can reliably factor numbers in a reasonable amount of time, we don't even know if such an algorithm could exist. (You can read more about computational complexity here.)

Braverman's groundbreaking ideas was to bring communication into the picture. He is being awarded the Abacus Medal for his "development of the theory of *information complexity*, the interactive analog of Shannon's information theory." The goal of this new field of *information complexity* is to apply the mathematical framework from information theory to computational settings. "In hindsight it's almost obvious that information theory should be a core tool, and hopefully now it's a little more central to our understanding," says Braverman.

Communication complexity

One example of this approach is studying the *communication complexity* of a task, where instead of asking how long the task will take to compute (the computational complexity), you instead ask how much communication you need to compute the task.

“Communication is an important part of computation,” says Braverman. There are many scenarios where a task is distributed, say between a number of servers holding a distributed data set, and you need some sort of coordination between the different servers working on the task. “In practice it’s often not the local processing that is the bottleneck [in a task], but moving the data back and forth between the worker and manager computers.”

A simple example is one where two parties, say you and Braverman, are trying to do some task together. Perhaps you have a file *X*, and Braverman has a file *Y*, and you both need to know if these files are identical. You could always just send your whole file to Braverman to compare with his, but communication complexity asks if there is a way of achieving this task that requires fewer bits of information to be communicated between you.

One approach could be for Braverman to send a *hash* of his file to you. A *hash* of a file is the result of applying an agreed mathematical function to the digitally encoded file. (You can read more here.) Hash functions are designed so that if two hashes agree, you can be pretty sure that the two files that were hashed are identical. This technique is used in the *check digits* of bar codes that allow scanners to check if the bar codes have been read correctly. “I could hash my file, send the hash to you, and you could compare the hash with your hashed file, and if they match, except in a vanishingly small probability, we can assume that the [files are the same].”

The last digit of a bar code is the check digit. The check digit is the answer to a mathematical function of the rest of the digits, allowing scanners to validate if the barcode has been read correctly.

To define things more formally, the exchange of information needed to complete a task is called a *protocol*, a kind of formalised conversation where the speakers interact with each other, their input each time depending on the conversation that has already taken place. The *communication cost* of a protocol is the number of bits communicated during the completion of the task. The hash of a file will be much smaller in size than the file itself, so this protocol would have a far smaller communication cost than the protocol that just

involved you sending your whole file to Braverman to compare with his file. And the *communication complexity* of a task is the smallest possible communication cost of a protocol that completes that particular task.

The new approach of looking at computation in terms of communication, to which Braverman has made important contributions, has brought new and deep insights into theoretical computer science.

Information complexity

An important aspect of computation these days is the possibility for two or more parties to share aspects of their data without revealing too much information to each other. "It's a major objective to be able to do computation without revealing information," says Braverman. For example, regulators might require access to data from private organisations to prevent or detect crime, but should not be given access to all information a company holds. This led Braverman to think about revealing information from a theoretical standpoint. The information cost of a protocol is the amount of information the two speakers learn about each other's inputs in order to successfully complete a task. And the information complexity of a task is the smallest possible information cost of a protocol for completing that task.

"In information complexity the goal is for the parties to teach each other as little as possible about their files," Braverman says, referring to our file comparison example above. "And it turns out, at least for problems involving two parties, you get a nice picture in the sense that the information complexity behaves in a very similar way to Shannon's information entropy."

Shannon's information theory applied to one-way communication, whereas the framework built by Braverman describes interactive communication in a way that provides a more fundamental understanding of computation. Braverman has proved important results in information complexity, understanding it's linked to communication complexity, and showing that it can be thought of as the interactive analog to Shannon's information theory. This work is relevant to other settings which depend on interactive communication including applications that are important in real life such as memory requirements in streaming algorithms, distributed error correction and information security.

Fundamental understanding

A prestigious prize such as the Abacus Medal, Braverman feels, also brings with it the responsibility to help set the direction of his field. "Theoretical computer science is very exciting because it's between the deepest maths and applications that are developing at breakneck speed," says Braverman. "But it's kind of tricky - how do you maintain long term focus while staying relevant, while also not getting pulled into the latest trend?"

This dual nature of the field is something that Braverman feels might be unique: "Some fields have this feeling of a Buddhist monastery – where it is looking to the very long term. And some fields exist in the present, driven by immediate needs such as the high commercial demand for applied machine learning research, or important needs like developing COVID vaccines or cutting edge cancer therapies. Here it's kind of both. Theoretical computer science might be unique in that it is simultaneously very slow and very fast."

For those of us outside of computer science we might think the biggest motivation would always be the applications, but it is the field's deep questions that motivate Braverman. "My main goal is a mathematical tool for understanding computation," he says. And this understanding could shed some fundamental insights.

"We've seen the stars for hundreds of thousands of years, we've had the natural numbers for 5000 years, and we've only been thinking of computation for about 100 years. But somehow, if we met aliens, they'd have looked at black holes, they'd have thought about prime numbers, and they probably have thought about computation. It's hard to imagine being advanced and not realising that [computation] is this basic process that you can abstract. Computation is basic, and it's important to understand its properties in the same way it is important to care about if the Universe is expanding or whether the Riemann hypothesis is true."

Marianne Freiberger and Rachel Thomas, Editors of plus.maths.org, interviewed Mark Braverman in June 2022.

This content was produced as part of the collaboration between plus.maths.org and the London Mathematical Society. You can find all our content on the 2022 International Congress of Mathematicians here.