

Marc'Aurelio Ranzato, Program Chair NeurIPS 2020

Welcome to NeurIPS 2020! I am Marc'Aurelio Ranzato, Program Chair of NeurIPS 2020 along with Hsuan-Tien Lin, Maria-Florina Balcan and Raia Hadsell. In my everyday life, I am a research scientist and manager at Facebook AI Research lab in New York City, working on machine learning and deep learning methods applied to computer vision and text understanding. In my role as program chair for NeurIPS, I worked with my program co-chairs and other organizing committee members to put together the main program of the conference.

It was also our role to define processes and tools that enabled us to coordinate a very large team of reviewers to eventually select 1898 papers from 9,467 submissions. Today, we are all excited to see the culmination of this process as the community gets together to start the 2020 virtual edition of NeurIPS.

It has been a very tough year. COVID-19 has taken many lives, including beloved members of our community, and it is still raging around the world. There have been very tense times around the world, with social movements demanding justice for all and natural disasters hitting several communities. NeurIPS 2020 is run by very talented people who share a passion for research in machine learning and AI, but more importantly, we are people. As such, we have all been affected directly and indirectly by these bigger events, and we have been supporting each other through these difficult times. As organizers of this conference, we had to make tough decisions. We ended up extending our submission deadlines twice, we had to make the call to go fully virtual back in June, and we had to be flexible in our processes due to the unprecedented circumstances. But, now we are finally together and it's a great feeling.

Let me summarize the work that went on behind the scenes to curate the main program, share some highlights and briefly explain how the conference will unfold.

As I mentioned, we received 9,467 submissions, an increase of 40% compared to last year. This rate of increase has been consistent over the past couple of years. Since we wish to provide 3 reviews for each paper, the challenge is to acquire 30,000 reviews from qualified experts in the field in just 3 months or so. That's a daunting task, as you might imagine.

Two innovations helped us achieve our objective. First, we introduced a summary rejection phase in which each paper got reviewed by two senior program committee members. Whenever both of them agreed that a submission had little chance of being accepted, we rejected it and immediately notified the authors. The other submissions were sent through to the regular full review process. Within a month, we removed a little over 1000 submissions. Second, we asked all authors who submitted papers to agree to serve as reviewers. We identified those who had experience reviewing and who had a sufficient number of publications at top-tier conferences in machine learning and related fields. As a result, we recruited several thousand additional

reviewers from this pool, for a total of over 7000 people who eventually provided 31,000 reviews!

Another change the team of program chairs made this year is related to ethics. Our program covers a broad spectrum of topics, from highly theoretical to directly applied. As a result, a substantial fraction of submissions cover topics that might have real-world consequences. The vast majority of these consequences are likely to be good, but some raise significant concerns. For instance, technology could potentially harm people if misused, or the required data could be associated with privacy infringements or other violations. We asked authors to include a section in their paper about the broader impact of their work, including a balanced view of the expected outcomes, and an explanation of how potential negative consequences might be mitigated in practice. In addition, we recruited a small pool of experts at the intersection of machine learning and ethics to provide assessments for papers that were deemed technically sound but were flagged by reviewers for ethical concerns. A few dozen papers underwent this additional review process, and eventually, we rejected four papers for issues related to ethics.

Let me now briefly describe this year's program. We accepted a total of 1898 papers. The three most popular topics continue to be general algorithms, deep learning and various kinds of applications (from computer vision to natural language processing, audio modeling, biology, healthcare, etc.). Core deep learning as well as theory and optimization have seen a decrease in the number of submissions, while algorithmic topics have seen the largest increase. In particular, there was a significant expansion in the number of submissions related to social aspects of machine learning, like fairness and privacy. This year, we also had a special call for submissions related to COVID-19, and we are pleased to include a handful of strong presentations on this topic at the conference, including one oral and four spotlight presentations.

This year, we are naming three best paper awards, highlighting outstanding work covering a nice spectrum of topics, theory, core algorithms and an application in natural language processing. This year's test of time award goes to a paper from NeurIPS 2011 that has been greatly influential in the development of large scale optimization methods, which have been critical to the success of the field over the past decade.

We were able to design the virtual conference to follow a schedule similar to past editions of NeurIPS. The main program will be presented for three and a half days, starting on Monday afternoon and ending on Thursday evening. Unlike physical editions of the conference, we allocated two sessions a day, each six hour-long, to help make the conference a planet-wide affair. The morning session starts at 5 am PT, while the afternoon session starts at 5 pm PT. The aim is to give everybody around the world the opportunity to attend at a convenient time. Each session unfolds as follows. We start with a plenary keynote talk, followed by the oral and spotlight presentations organized into multiple thematic tracks, and we conclude with a two-hour-long poster session. Each talk is pre-recorded and live-streamed and will be made freely available after the conference has taken place. However, there will be live and moderated

Q&As, and attendees will have the chance to interact with each other via chat and at the poster sessions hosted through Gather Town. In Gather Town each attendee controls an avatar that moves around in a virtual conference hall. This setup simulates the physical experience, enabling attendees to visit posters as well as to chat informally in the hallway with other attendees. The organizers have experimented with this arrangement and, we have to admit, it's a pretty fun experience on its own! But, it's also a little bit frightening to us as organizers as it is the first time that this technology has been used at this scale!

We are very proud to have greatly reduced the registration fee this year and to be making most of our content, including all oral presentations and tutorials, available for free on our website. Registration is required only for the interactive parts of our program, such as participating in the poster session and accessing the internal chat.

Please, check out our blog posts on medium (<https://neuripsconf.medium.com/>) for more details about the paper-selection processes as well as for additional information about the conference.

Thank you for joining us, and we hope you will enjoy the virtual experience at NeurIPS 2020!

Now I'd like to introduce the Test of Time Award winner and the three Best Paper Award winners who will share with you their outstanding contributions to this year's program.