

Neural Architects workshop, ICCV 2019

Deep Neural Networks (DNNs) now represent a fundamental building block of machine perception methods and are a core technology for autonomous driving. At the 2019 International Conference on Computer Vision (ICCV) in Seoul, South Korea, Momenta hosted a workshop that aimed to bring together some of the world's leading researchers to discuss the open challenges relating to state-of-the-art DNN design.



The opening keynote "Neural Architecture Search and Beyond" was delivered by Barret Zoph, a senior research scientist at Google Brain. The talk was extremely popular!

The workshop, which was organised through a collaboration between researchers at Momenta, the University of Oxford, Tencent and Google, sought to engage the community on many of the different research directions that are relevant to DNNs. The day consisted of a series of invited talks, together with three oral presentations and 22 poster presentations from researchers focusing on architecture design.

Barret Zoph (Google Brain) launched the proceedings by laying out the tremendous promise of fully-automatic machine learning through a technique known as Neural Architecture Search (NAS). NAS, which

was pioneered by Barret and others, has shown remarkable results on a wide range of competitive vision benchmarks. Barret also highlighted some of the challenges NAS faces in achieving more widespread adoption: "One issue is the enormous compute consumption" he said, describing some of the breakthrough early research in this area. Nevertheless, he was particularly optimistic about recent work that aims to achieve "efficient NAS", bringing down the compute cost to a level that would allow more researchers and practitioners to explore these ideas.

Next, Prof. Iasonas Kokkinos (UCL/Ariel AI) set out his vision for developing architectures that can tackle many tasks concurrently. These multi-tasking architectures offer significant computational advantages in modern computer vision systems, but getting them to work in practice is challenging. If multi-tasking is implemented naively, Iasonas noted that "rather than having this universal, omnipotent network, we end up having this dilettante that tries to do both and does nothing well!" He showed how attention mechanisms, inspired by the human vision system, can be used to effectively address this issue.

Prof. Alan Yuille (Johns Hopkins) then offered a critical review of the limitations of existing DNN designs. In his talk "Deep Nets: What have they ever done for Vision?" he showed that DNNs have benefited the community tremendously, but they are not a panacea. He pointed out several failure cases for current state-of-the-art models and suggested that new approaches were needed: "Images in the real world, once you go beyond faces and beyond local parts, are combinatorially complex and I don't see how our current methods can deal with that. As a community that's something we need to think about and study."

Following this, Sara Sabour (Google Brain) outlined how capsule-based networks can address several of the issues for existing architectures. One of their key advantages is their ability to generalise: "You don't need to learn new weight parameters to handle new viewpoints" she said, describing the mechanism for capsules. She showed a series of promising experimental results which show that capsules could offer significant performance gains but stressed that there was still a lot to be done for this approach to reach its full potential. Nevertheless, capsules have been receiving growing interest in the community and it seems likely that this trend is set to continue.

The final keynote of the day was delivered by Ross Girshick (Facebook AI Research). Ross noted that research was shifting from the design of the architecture to the design of architecture generators. "I think that there's a lot of power in looking at populations of networks rather than individual networks" he commented. He also presented recent work on using randomly-wired network designs, showing that they can be surprisingly competitive on many benchmarks.

In addition to the keynote talks and a panel discussion, the workshop featured three oral presentations and a vibrant poster session where 22 further peer-reviewed works were presented.

Video recordings and presenter slides for each presentation can be found on the workshop website: <https://neuralarchitects.org/>

The workshop was undoubtedly a major success and we hope to support further editions in future!

About Momenta

Momenta is building the “brain” for autonomous vehicles. Our deep learning-based software in perception, HD Map, and data-driven path planning enables the realization of full autonomy.

Momenta pursues a two-leg strategy of autonomous driving for mass-production passenger vehicles (Mpilot) and full autonomy (MSD) for mobility service vehicles. Data, data-driven algorithms, and closed-loop automation between Mpilot and MSD are the key to achieving autonomous driving at scale.