



## AI HARDWARE SUMMIT EUROPE

The Industry Event for the AI  
Hardware Ecosystem

March 10-11, 2020  
Munich, Germany

REGISTER NOW



[www.aihardwaresummiteu.com](http://www.aihardwaresummiteu.com)



+44 (0)20 3696 2920



[events@kisacoresearch.com](mailto:events@kisacoresearch.com)

## SYSTEMS-LEVEL CONSIDERATIONS FOR AI HARDWARE: PERSPECTIVES FROM DEEPMIND & FLEX LOGIX.

“While AI is often equated with deep neural networks, in practice, complex AI applications contain many non-NN parts, and the orchestration between these different parts is complex, making a holistic system software and hardware approach indispensable for achieving good performance.”

In preparation for the inaugural AI Hardware Summit in Munich, we caught up with two presenting companies, DeepMind and Flex Logix, to give us some insight into some of the systems-level considerations around deploying AI hardware.



Olivier Temam  
Hardware Engineer  
DeepMind

“AI is becoming pervasive in many industries, and AI hardware is one of the key components for enabling a broad adoption of AI. In parallel, Europe has many strong industries, and many strong AI researchers, so there is a natural fit for the emergence of a strong AI hardware industry. Moreover, the groundbreaking nature of AI extends to hardware, and it creates opportunities for new entrants and startups. At DeepMind, we want a global thriving ecosystem in order to keep innovations high in hardware.

While AI is often equated with deep neural networks, in practice, complex AI applications contain many non-NN parts, and the orchestration between these different parts is complex, making a holistic system software and hardware approach indispensable for achieving good performance. It is also a massive technical challenge.

During my keynote presentation at the AI Hardware Summit in Munich in March, I will discuss some of these points. I will talk about how the system angle is really important; for good reasons, we have initially focused on chips and NNs, but it is really time that we develop a broader system view of the problem. The other fundamental aspect that we need to fully absorb is that AI, as a family of algorithms, is atypical because it is evolving very fast. This creates massive challenges for both hardware and software to cope with, and still service the AI algorithms well.

Another major challenge is the cost of AI training and its associated carbon footprint: how do we keep costs and carbon footprint low, yet enable the fast development of future AI innovations?”



Geoff Tate  
Co-Founder & CEO  
Flex Logix

“What we see is that customers have various problems they are solving and AI inference is a part of their system, but the AI Inference is the part that a) they can’t do themselves, at least not well; b) consumes a large part of the system’s compute/power/\$.

So the customers we are talking to want an inference subsystem/co-processor and they’ll do the rest. Some want a PCIe board for their edge server with the rest of their software running on the host processor. Some want a chip to integrate into their system that talks to their host that does the rest of the processing. In the latter case the host may have specialized hardware for vision/DSP/etc.”