The Alan Turing Institute

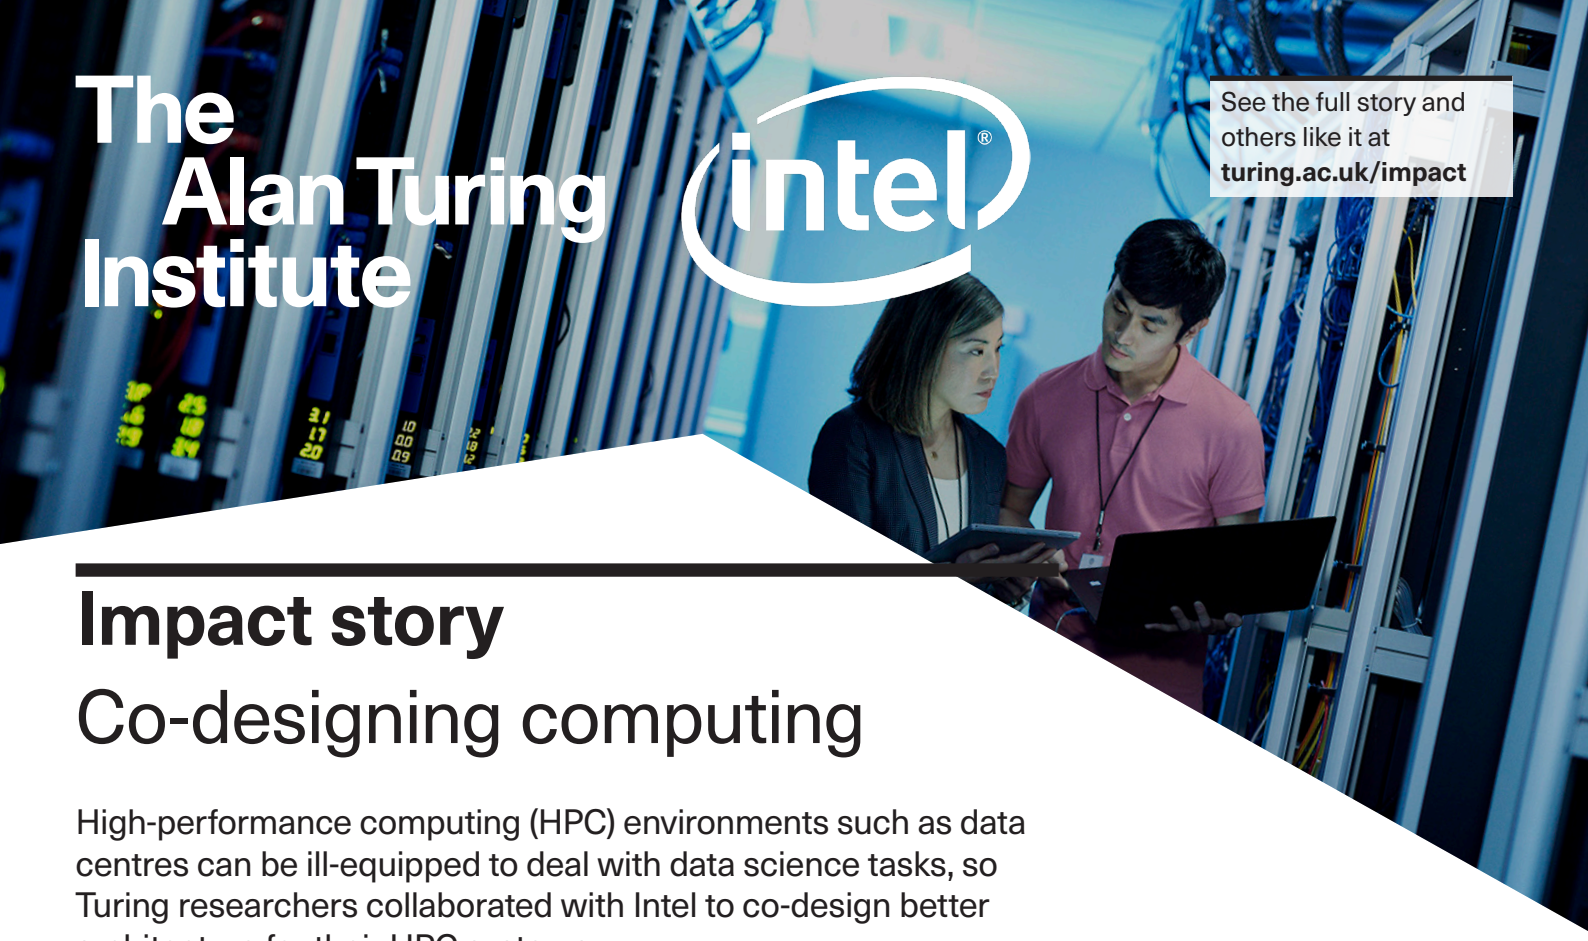

See the full story and others like it at **turing.ac.uk/impact**

Impact story

Co-designing computing

High-performance computing (HPC) environments such as data centres can be ill-equipped to deal with data science tasks, so Turing researchers collaborated with Intel to co-design better architecture for their HPC systems.

- In high-performance computing environments, well-designed algorithms and architectures allow huge data-analysis tasks to be performed.
- But for many data science tasks, these systems run at less than half their full capacity.
- Training multi-layered 'neural networks' on HPC architectures requires optimised communication between cores in the architecture.
- Researchers at the Turing worked with Intel's engineers to co-design improved architectures for their HPC systems, to make communication between machines more efficient.
- The collaborators examined Intel's HPC Omni-Path Architecture and identified where the code wasn't running efficiently. Resulting improvements boosted bandwidth and delivered a 10-fold increase in processing speed.
- To ensure accuracy when working with neural networks, the 32-bit data format is used because the simpler 16-bit format is insufficient for many machine learning problems.
- The team created a customised 16-bit format that could handle such problems, saving memory, making calculations much faster, and saving silicon space in the hardware.

Impact

- Collaboration between the Turing and Intel has helped Intel improve their products and services.
- Also enabled data scientists at the Turing, and the wider world, to manage and analyse massive datasets with greater efficiency.
- The code that this collaboration generated now shipped as standard in Intel products.

"We've seen a positive impact on our architecture as a result of our work with the Turing."

Anil Rao, Vice-President, Data Center Group and General Manager of Data Center Security and System Architecture at Intel

In distributed training, each machine in an HPC environment has to communicate effectively.