DELL Technologies

AI deep learning extends data science horizons

Data scientists at Stanford University are exploring new terrain for deep learning models using massively parallel processing.

Customer profile

Leading Research University

Information Technology | United States

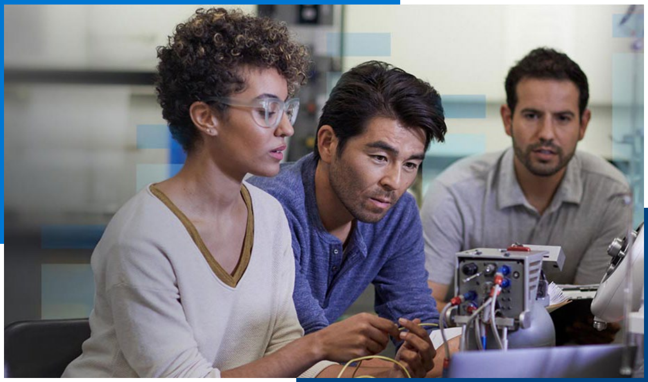

> "We've opened new ways of thinking about our AI approaches and how we can use the Dell Precision and NVIDIA technologies to try them out."
>
> **Jure Leskovec**
> Associate Professor of Computer Science, Stanford University

Organization needs

To find patterns and insights in multi-terabyte-sized sets of complex, unstructured data requires huge amounts of parallel processing power. Stanford University's data scientists were using consumer-grade gaming graphics processing units (GPUs) but needed more GPU-accelerated processing power, with central processing units (CPUs) helping to keep the GPU utilization rate maximized while running deep-learning AI models.

Organization results

- Broadens the scope and size of problems researchers can solve.
- Enables previously impossible inferences from data.
- Reduces AI algorithm development and machine learning times by up to 50%.
- Provides supercomputer power in a go-anywhere mobile laptop.
- Handles data sets of up to 1.5TB.

Solutions at a glance

- Dell Precision 7750 Mobile Data Science Workstation
- Dell Precision 7920 Tower Data Science Workstation
- NVIDIA Quadro RTX Powered Workstations

Data science is the discipline that drives one of the world's fastest-growing — and transformational — technological developments: artificial intelligence (AI). While AI's historical roots go back to the 1950s, today's data scientists have pushed its capabilities much further than even a few years ago by applying ever-increasing amounts of massively parallel processing power to their data methods and learning models.

Such is the case at Stanford University, where Jure Leskovec, associate professor of Computer Science, and Rok Sosič, a senior researcher, lead a team of 40 other data scientists. Individually and in teams, Leskovec and Sosič are exploring applications of AI in the form of network analytics and deep learning for a wide range of problems. Network analytics doesn't refer to IT networks, but rather to finding complex patterns among relationships between numerous data entities.

"Our mission is to do groundbreaking research to tackle a variety of extremely challenging research problems. We do this by developing new methods to learn from data, and then use those methods to derive fresh insights from the data that can be applied in practical problems," says Leskovec, noting that the Stanford lab is the world's leading developer of AI models for life sciences applications.

Deep learning reveals complex network effects

To illustrate AI-based network analytics, Sosič cites his team's research into the side effects of polypharmacy, saying, "When a person takes multiple prescription drugs, the interactions of those drugs and the dynamics of those interactions can vary, depending on the person's disease(s), genetics, gender, age and so forth."

Sosič further explains, "Traditionally, extremely costly clinical trials examine large cohorts of patients to draw inferences and conclusions about drug efficacies. However, such trials are not feasible in the case of patients who use multiple drugs because there are too many possible combinations. What we do is develop a deep-learning methodology to analyze the complex networks of these interactions along with the varied characteristics of proteins that these drugs act on. This approach provides the pharmaceutical domain with greater insights into how their products work and, more importantly, how they can potentially work better."

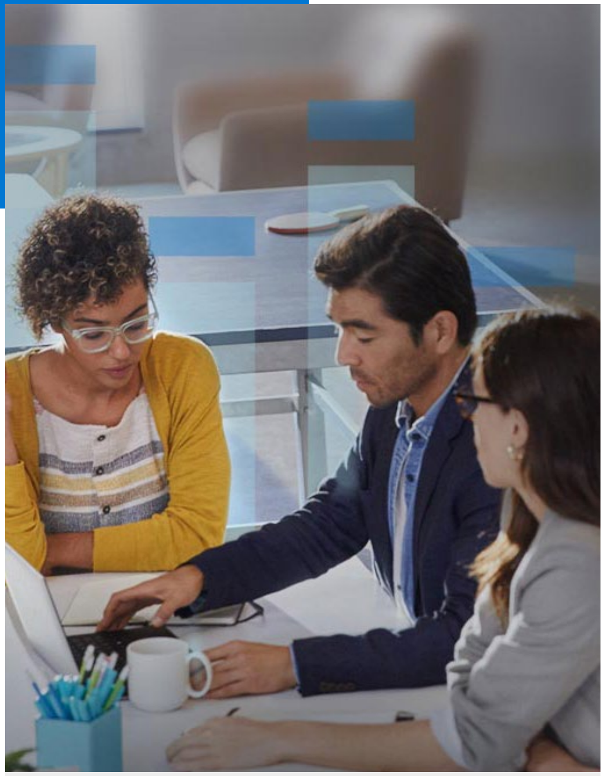

"Using the Dell Precision 7920 Tower Workstation with the NVIDIA GPUs, we can now explore much bigger and more complex problems in data sets up to 1.5TB in size."

Rok Sosič
Senior Researcher,
Stanford University

Comes ready-to-use, **simplifying setup** and saving time.

To conduct their research, which typically involves multi-terabyte-sized data sets, the lab's team members tap into a large bank of servers, which were previously equipped with consumer-grade GPUs. These GPUs were originally designed for gaming's video rendering demands to take that processing burden off a gaming PC's CPU.

"We loaded as many of these GPU cards as we could into our servers, but we lacked enough RAM per GPU, so the limits to their processing power became the constraining variable in extending the reach of our research," says Sosič.

Broadening research horizons via massively parallel processing

To deliver a quantum leap in the parallel processing power available to the team, the Stanford Data Lab acquired a Dell Precision 7920 Tower Data Science Workstation (DSW). A collaboration with NVIDIA and other leading technology providers such as Canonical (the company behind Ubuntu, the worldwide de-facto Linux OS used in AI on workstations), the Dell Precision 7920 Tower DSW comes as a fully integrated, ready-to-use AI hardware and software package. This package simplifies user setup, saving days over a bare-metal, build-it-yourself approach, not to mention the configuration efforts that would otherwise be required.

Its thermal, bus design and engineering make the Dell Precision 7920 Tower DSW capable of handling up to three NVIDIA Quadro RTX 8000 GPU cards, each with 48GB of memory and 4,608 parallel-processing cores. Collaborative efforts between Dell Technologies and NVIDIA engineering teams have aligned the performance of the DSW's CPU — available in dual processors and up to 2TB of RAM — with GPUs to keep utilization maximized throughout an AI model's training run.

The lab team uses the Dell Precision 7920 Tower DSW with NVIDIA RTX 8000 GPUs as a shared resource for developing their AI algorithms and machine learning. "It's been a giant leap forward for us," Leskovec says. "We've opened new ways of thinking about our AI approaches and how we can use the Dell Precision and NVIDIA technologies to try them out."

Sosič adds, "Using the Dell Precision 7920 Tower Workstation with the NVIDIA GPUs, we can now explore much bigger and more complex problems in data sets up to 1.5TB in size."

> "We ran some benchmarks against our servers with lower-end GPUs and found the Dell Precision 7750 with NVIDIA GPUs cut the time by as much as 50 percent."
>
> **Rok Sosič**
> Senior Researcher,
> Stanford University

> “In fact, having your own personal Dell Precision DSW running NVIDIA RTX GPUs is like driving a Ferrari compared to the trucks, which are shared, that our other lab servers would represent.”
>
> **Rok Sosič**
> Senior Researcher,
> Stanford University

Go-anywhere supercomputing fits in a backpack

The Stanford University Stanford Data Lab also acquired a Dell Precision 7750 mobile DSW, a 15-inch laptop with an 8-core CPU, 64GB of RAM and an NVIDIA Quadro RTX 5000 GPU with 16GB of RAM. Like the Dell Precision Tower DSW, this laptop is a fully integrated, ready-to-use AI device with special thermal engineering to dissipate the heat generated when both the CPU and GPU are operating flat-out on an AI model.

Sosič and Leskovec note that the Dell Precision 7750 mobile DSW is especially fast compared to their servers. “We ran some benchmarks against our servers with lower-end GPUs and found the Dell Precision 7750 with NVIDIA GPUs cut the time by as much as 50 percent,” Sosič says.

According to Sosič, the Dell Precision 7750 mobile DSW is signed out to one of the team’s data scientists to enable him to develop his AI models at home or wherever he might choose. “Mobility like this can benefit our researchers who are working from home so they can develop their models using the Dell Precision 7750 and then upload them to our data center for further training,” he says.

In Sosič’s view, the fully integrated, AI-ready Dell Precision DSW technology combined with powerful NVIDIA RTX GPUs are helping the Stanford Data Lab apply the parallel processing needed to continue its breakthrough research. “We’re able to draw data inferences we couldn’t before,” Sosič says. “In fact, having your own personal Dell Precision DSW running NVIDIA RTX GPUs is like driving a Ferrari compared to the trucks, which are shared, that our other lab servers would represent.”