

Pushing Computer Vision Boundaries Beyond 4K



Marshall Choy is the Vice President of Product at SambaNova Systems, responsible for product management and go-to-market.

In the context of machine learning image processing and analysis—resolution is everything. An image's resolution can enable a more detailed, meaningful analysis that results in greater understanding. To this end, a high-resolution image will contain more information and detail than a low-resolution image of the very same subject.

Today, higher-resolution image processing requires significant computational capabilities. So much so, in fact, that training models to use these high-resolution images has rendered current state-of-the-art technologies unusable.

When it comes to high-resolution processing, legacy architecture gridlock is holding back research and technology advances across numerous use cases, including in areas such as autonomous driving, oil and gas exploration, medical imaging, anti-viral research, astronomy, and more.

Surpass the Limits of the GPU

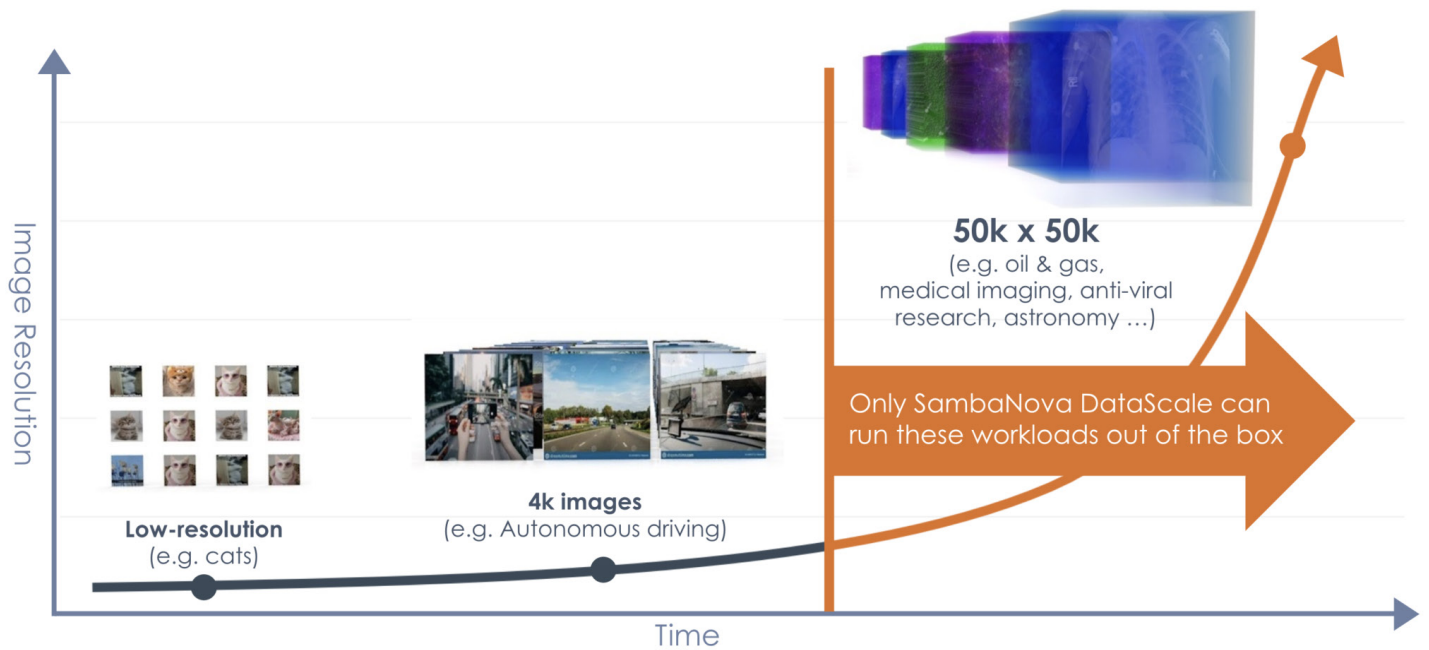
SambaNova Systems has been working with industry partners to develop an optimized solution for training computer vision models with increasingly growing levels of resolution—without compromising high accuracy levels. We take a “clean sheet” complete systems approach to enable native support for high-resolution images. Co-designing across our complete stack of software and hardware provides the freedom and flexibility from legacy GPU architecture constraints and legacy spatial partitioning methods.

Adding More GPUs Isn't the Answer

If you consider images in the context of AI/ML training data, the richer and more expansive your training information (i.e., images), the more accurate your results can be.

Using a single GPU to train high-resolution computer vision models predictably results in “Out of Memory” errors. On the other hand, clustering multiple GPUs brings all the challenges of disaggregation of the computational workflows onto each individual GPU in the cluster to aggregate GPU memory.

In this case, this is not merely clustering a few GPUs in a single system, but aggregating hundreds, if not thousands of GPU devices. In addition, conventional data parallel techniques that slice the input image into independent tiles deliver less accurate results than training on the original image.




Train Large Computer Vision Models With High-Resolution Images


Massive Data: A single SambaNova DataScale™ system—with petaflops of performance and terabytes of memory—is designed as a Dataflow architecture. This co-design of software and hardware properties is built to enable high-performance processing of a range of complex structures such as high-resolution images, pushing computer vision boundaries far beyond 4k.

SambaFlow™ Software: The SambaFlow software stack transforms deep learning operations to work seamlessly. SambaFlow native software support for tiled input images, intermediate tensors, and convolution overlap handling are all automated. The results are equivalent to the non-tiled version and require no changes to the application or programming model.


Robust Architecture: SambaNova's Reconfigurable Dataflow Architecture is critical for efficient processing of input image tiles and is fully materialized in device memory, unlike with non-Dataflow devices, such as GPUs.



Massive Data
SambaNova Systems
handles massive
amounts of data



SambaFlow™ Software
Automatically tiles the
input image for deep
learning operations



Robust Architecture
Efficient processing
materialized in
device memory

Re-Think What's Possible

The resulting solution is unlimited by capacity and is capable of processing images of any size on a single DataScale system. End users then have the option of scaling up additional DataScale compute resources to further reduce training times while maintaining high levels of utilization and accuracy.

The high-resolution image processing breakthroughs achieved on SambaNova DataScale allow organizations to cut years of development time, significantly simplify architecture, and ease programmability. All this while yielding state-of-the-art results and capabilities.

SambaNova provides state-of-the-art technologies to support NLP, high-resolution computer vision, and recommender models. To learn more, request a meeting.

Request a Meeting

sambanova.ai

