

Rethink What's Possible: Pushing Computer Vision Boundaries Beyond 4k



Computer vision and image processing algorithms are increasingly common to commercial and research-related applications. These high resolution models perform complex computations to process image data in real-time and require rapid processing to yield analysis and results.


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 have rendered current state-of-the-art technologies unusable.

When it comes to high-resolution processing, legacy architecture constraints are 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.

What's Inside

Surpass The Limits of The GPU	2
Adding More GPUs Isn't The Answer	2
Train Large Computer Vision Models With High-Resolution Images	3
Rethink What's Possible	3



Surpass The Limits of The GPU

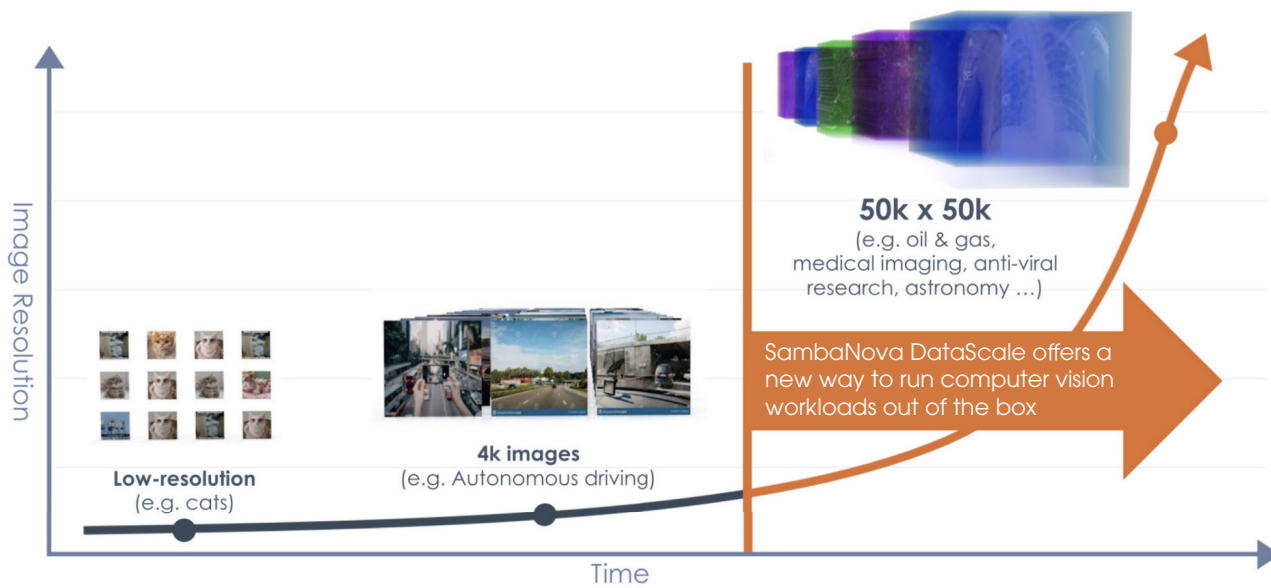
SambaNova Systems has been working with industry partners to develop an optimized solution for training computer vision models with increasingly growing levels of resolution — without compromising high accuracy levels. We take a “clean sheet” complete systems approach to enable native support for high-resolution 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 built with a dataflow architecture design. 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.

Rethink What's Possible

The resulting solution is capable of processing images of extremely high resolution — from 4k up to 50k — 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.



Learn more at SambaNova.ai



linkedin.com/company/sambanova



[@SambaNovaAI](https://twitter.com/SambaNovaAI)



[@SambaNovaAI](https://facebook.com/SambaNovaAI)



info@sambanova.ai

AI is here. With SambaNova, customers are deploying the power of AI and deep learning in weeks rather than years to meet the demands of the AI-enabled world. SambaNova's flagship offering, Dataflow-as-a-Service™, is a complete solution purpose-built for AI and deep learning that overcomes the limitations of legacy technology to power the large and complex models that enable customers to discover new opportunities, unlock new revenue and boost operational efficiency. Headquartered in Palo Alto, California, SambaNova Systems was founded in 2017 by industry luminaries, and hardware and software design experts from Sun/Oracle and Stanford University. Investors include SoftBank Vision Fund 2, funds and accounts managed by BlackRock, Intel Capital, GV, Walden International, Temasek, GIC, Redline Capital, Atlantic Bridge Ventures, Celesta, and several others. For more information, please visit us at sambanova.ai or contact us at info@sambanova.ai. Follow SambaNova Systems on LinkedIn.