

The People Who Understand Bottlenecks

Jensen Huang, Blender topology, and why the future still has to pass through the physical world.

2026.05.14 11:02 DENNIS HEDEGREEN ANALYSE V1.0

[HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/THE-PEOPLE-WHO-UNDERSTAND-BOTTLENECKS/](https://hedegreenresearch.com/articles/the-people-who-understand-bottlenecks/)

Jensen Huang, Blender topology, and why the future still has to pass through the physical world

I was listening to Jensen Huang talk about bottlenecks and realised I had stopped listening to the business point.

Not because he was boring.

The opposite.

My mind had gone somewhere else.

To a bottle.

A real bottle.

Or more precisely, a bottle in Blender.

A wide body. A narrow neck. A shoulder between them. A simple object, until you actually have to model it properly. Then the problem appears. The neck is not the hard part. The transition is.

You cannot move from one mesh logic to another for free.

If a surface begins with three clean quads and suddenly needs to become five, the mesh has to pay somewhere. An edge has to redirect. A pole has to appear. A flow has to change. The topology has to absorb the difference.

If you do it badly, the surface tells on you.

The highlight breaks.

The curve pinches.

The model looks wrong before the viewer can explain why.

Systems are like that too.

A bottleneck is not just the narrow point where things slow down. A bottleneck is where the whole structure reveals how it handles pressure.

That may be why I have more respect for Jensen Huang than I have for most tech CEOs.

Not because he is rich.

Not because Nvidia became the symbol of the AI boom.

Not because the leather jacket became a costume for accelerated capitalism.

But because he talks like someone who understands that the future has a shape.

AI is often sold as if it floats above the world. Models. Intelligence. Agents. Reasoning. Magic little boxes that answer questions and write code and draw pictures and promise to reorganize work.

But intelligence does not float.

It runs somewhere.

It consumes power.

It produces heat.

It waits on memory bandwidth.

It needs networking.

It needs chips.

It needs data centers.

It needs factories that can build the chips.

It needs machines that can build the factories that build the chips.

The more people talk about AI as software, the more important it becomes to listen to the people who remember that software still has to pass through matter.

That is where Nvidia's story becomes interesting.

Nvidia did not simply wake up one morning and become useful to AI. For years, the company built around a different belief: that general-purpose computing would not be enough for the future. Some problems are too parallel, too heavy, too physical, too hungry for ordinary CPU logic.

So Nvidia built not only chips, but a path.

GPU.

CUDA.

Libraries.

Developers.

Systems.

Networking.

Data centers.

A full stack around accelerated computing.

That is the part people miss when they say Nvidia sells the shovels in the AI gold rush.

They do sell shovels.

But they also sell the handles, the map, the mining method, the training camp, the replacement parts, and the shared language all the miners learned to use.

A shovel is a tool.

A stack is a dependency.

And that is why bottlenecks matter.

When AI suddenly became the thing everyone wanted, the question was not only: who has the best model?

The question became: where does the pressure go?

It went to GPUs.

Then memory.

Then networking.

Then data centers.

Then energy.

Then cooling.

Then TSMC.

Then ASML.

Then export controls.

Then cloud contracts.

Then talent.

Then inference cost.

Every time the world found one bottleneck, another appeared underneath it.

This is not a failure of AI.

This is what happens when an idea becomes infrastructure.

Ideas are light.

Infrastructure is heavy.

A chatbot can appear instantly in a browser. But the system behind it is not instant. It is layers of physical, financial, and technical dependency. The visible interface is smooth because the hidden topology is complex.

That is the bottle again.

A clean bottle looks simple because the transition work has been hidden inside the form.

A bad bottle shows the stress.

The shoulder collapses.

The neck pinches.

The geometry cannot carry the change.

In Blender, you learn quickly that the narrow point is not the whole problem. The real work is in the transition zone. How does the wide body become the narrow neck without lying to the surface?

In AI, the same thing is happening.

Human ambition is wide.

The passage is narrow.

Everyone wants intelligence everywhere: in phones, browsers, cars, robots, hospitals, factories, schools, offices, homes, farms, and governments.

But the world has to route that ambition through actual constraints.

Compute.

Energy.

Chips.

Cooling.

Capital.

Supply chains.

Law.

Trust.

Time.

The future does not arrive as an idea.

It has to pass through a topology.

That is the sentence I keep returning to.

Because it explains more than AI.

A poor person's life has topology. Money becomes the narrow passage, and everything else deforms around it.

A municipality has topology. Caseworkers become the narrow passage, and citizens, forms, deadlines, rules, budgets, and risk all press into that point.

A market has topology. If one supplier becomes the narrow passage, the entire industry reshapes around that supplier.

A mind has topology too. If fear, shame, trauma, or uncertainty becomes the narrow point, then thoughts and decisions begin to flow around that pressure.

So when Huang talks about bottlenecks, I hear more than an executive explaining supply and demand.

I hear someone describing the shape of pressure.

And maybe that is why he seems different from the normal tech-CEO theatre. He does not sound like a man trying to convince you that the future will be frictionless. He sounds like someone who has spent decades studying where friction moves next.

That does not make Nvidia harmless.

A company sitting at the bottleneck becomes powerful. Sometimes too powerful. The shovel seller in a gold rush does not need to own the gold to shape the rush. If everyone needs the same tools, the toolmaker becomes part of the terrain.

The danger is not only price.

It is lock-in.

CUDA is not just software. It is habit, training, documentation, libraries, optimization work, developer muscle memory, research code, production pipelines, and procurement logic. Once enough of the world has learned to build through one stack, the bottleneck becomes cultural as much as technical.

That is a different kind of power.

A chip can be replaced in theory.

A dependency stack is harder to move.

And underneath even that sits another layer most people never see: the machines that make the chips possible at all.

ASML in Veldhoven does not appear in the chatbot interface.

But without advanced lithography, there are no leading-edge chips.

Without leading-edge chips, the bottleneck moves before the model even begins.

This is where respect has to stay awake.

Understanding bottlenecks does not make a company morally good. It makes it strategically dangerous. The person who understands where the world bends can help build the future. They can also decide who gets to pass through the narrow point first.

That is why respect cannot become worship.

Still, there is something valuable in the way Huang frames the world. He keeps pulling AI back down to earth. Back to factories. Back to power. Back to systems. Back to the material stack beneath the demo.

That matters in a culture addicted to interfaces.

We keep staring at the smooth surface.

The app.

The answer.

The agent.

The keynote.

The model name.

The benchmark.

But the real story is often lower down.

Where the mesh changes.

Where the flow redirects.

Where the hidden structure has to absorb pressure.

Where the bottleneck is not just blocking the system, but shaping it.

Maybe that is the lesson from a bottle in Blender.

Maybe it is also the lesson from Nvidia.

The future is not only invented by the people with the biggest ideas.

It is built by the people who understand where those ideas get stuck.

And sometimes the most important person in the room is not the one promising a frictionless future.

It is the one pointing at the bottleneck and saying:

There.

That is where the world has to bend.

RELATION MEMORY

MATURES FROM [A Moratorium for AI Data Centers](#) · Extends the physical-infrastructure AI question into bottlenecks, topology, and accelerated-computing constraints.

NEARBY [Google Had a Brake](#) · Both articles read AI power through infrastructure constraints rather than model spectacle alone.

CANONICAL URL: [HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/THE-PEOPLE-WHO-UNDERSTAND-BOTTLENECKS/](https://hedegreenresearch.com/articles/the-people-who-understand-bottlenecks/)
PDF VIEW GENERATED: 2026-07-11T15:35:46.040Z