



OPINION AND ANALYSIS

The Borrowed Intelligence Economy

When packaging other people's minds became an industry—and inventing real intelligence became an oddity among its pretenders

In 1984, a small Dutch company began work in a rain-leaking shed near the Philips facilities in Eindhoven. Its project was to develop the photolithography machines used in semiconductor manufacturing. [1](#)

Its name was **ASML**.

Four decades later, ASML had become the only company shipping commercial extreme ultraviolet lithography (EUV) systems, among the most complex machines in modern industrial history.

Inside one of these machines, microscopic droplets of molten tin are fired into the path of successive laser pulses, fifty thousand times per second. The material becomes plasma that emits light at a wavelength of 13.5 nanometres. That light then travels through an optical system of extraordinary precision to help etch microscopic circuits onto silicon wafers. [2](#)

The machine is not a triumph of mechanical engineering alone. It is the result of optics, plasma physics, materials science, control engineering, software, and precision metrology converging in a single system.

More striking still, ASML did not invent all these sciences, or every component in its machine. It built a specialised ecosystem of suppliers and partners, including Germany's ZEISS for optics. The American company Cymer collaborated with ASML on developing the EUV light source before ASML acquired it in 2013 to accelerate that development. [3](#) [4](#)

What it did possess was the engineering knowledge required to combine those components into a system the world had never before been able to produce in this form.

The Netherlands therefore became a power in this industry not because it bought the latest chipmaking machines, but because a company founded on its soil came to produce a machine required by the makers of the most advanced chips.

Here lies the difference between possessing technology as a commodity and possessing the knowledge that makes its production possible.

While ASML was building its place in the global economy through the accumulation of knowledge, the age of artificial intelligence began to produce a class of company that wanted the same standing by a much shorter route: rent intelligence from another company, place an elegant interface over it, then announce on its website that it had developed proprietary artificial intelligence technology.

Perhaps this is the first time when, in some circles, changing a product's name on a screen has been enough to turn a user of technology into its inventor.

Wrapper Factories in the Age of Rented Minds

Generative artificial intelligence has created a historic opportunity to build applications. Alongside it, however, it has created an equally prosperous market for manufacturing technical titles.

A startup today can use a large language model (LLM) developed by OpenAI or Anthropic, connect it through an application programming interface (API), then add retrieval-augmented generation (RAG), a set of prompts, and tools for task execution through tool calling.

With a little workflow orchestration, the product appears as an artificial intelligence agent capable of reading documents, summarising meetings, responding to customers, generating reports, and interacting with certain enterprise systems.

Up to this point, we are looking at an intelligible software business. It may be useful and profitable.

Then comes the interesting turn.

The company begins speaking of a “proprietary AI engine,” an “advanced inference architecture,” and “owned intelligent models.” It may even elevate its description of itself to a deep-technology company, while the capability it is describing resides, at its core, inside GPT, Claude, or another commercial model.

They did not invent the mind. They invented a new name for it, then asked the market to recognise their scientific parentage.

Some of these companies possess little more of an “AI engine” than an API key, a software development kit (SDK), and instructions sent to a model hosted by someone else.

Yet their founder speaks enthusiastically about the company’s cognitive architecture, reasoning abilities, and technical independence.

The scene deserves a little ridicule. When the model provider changes its token pricing, the economics of the supposed invention change with it. When the service connection fails, most of its intelligence disappears. When the supplier releases a new feature, months of “revolutionary development” become a free function available to everyone.

This is an industry in which a company can declare its cognitive independence in the morning, then receive an error message in the evening: **429 Too Many Requests**.

How many independent minds cease to think when their monthly token allowance runs out?

The phenomenon is known in technology circles as model API wrappers, or more specifically, GPT wrappers.

They can take the form of software-as-a-service (SaaS) products, agent platforms, or vertical artificial intelligence applications for medicine, law, education, and financial services.

Changing the sector does not change the origin of the capability.

A medical product that uses Claude to analyse a document does not become an original medical model because its interface wears white. A legal agent that calls GPT does not become a discovery in legal reasoning because it has learned to draft memoranda in formal prose.

Nor does adding a thousand documents to a retrieval database, or connecting the agent to ten enterprise systems, mean that the company invented a language model or a new learning algorithm.

It may have built an excellent product. But an excellent product is one thing; an invention in artificial intelligence is another.

The problem is not the use of other people’s technology. The entire history of industry rests on accumulation and the mutual use of inventions.

The problem is turning your dependence on someone else’s knowledge into a claim that you created that knowledge.

In the traditional economy, we call this importation, integration, or product development. In the artificial intelligence economy, some have begun calling it “pioneering research and development.”

At that point, the issue is no longer mere marketing excess. It is a deliberate blurring of what innovation means and of the kind of industry we want to build.

The Market for Lemons That Sells Genius

In 1970, the American economist George Akerlof published his celebrated paper *The Market for Lemons*, which became one of the foundational texts for understanding information asymmetry in markets.

Akerlof examined how a market can fail to distinguish between good and poor goods when the seller knows what the buyer does not, and how this can lead to adverse selection, weakening the market’s ability to reward genuine quality.

I believe we are looking at a contemporary version of the same problem. This time, the commodity is not a used car, but a claim to possess intelligence.

Two teams stand before an investor.

The first spent a year studying model architectures, improving training algorithms, testing memory-management strategies, and developing new approaches to language processing.

The second spent a few weeks building a beautiful interface over a commercial model.

Both teams demonstrate their products. The answers are polished, and the two presentations look similar.

The second team’s presentation may even be more attractive, faster, and better able to impress the committee.

Because many investment presentations test visible outputs more than they understand the technical provenance of the capability, both teams receive the same description: two leading artificial intelligence companies.

This is where a small economic crime begins, one that grows each time it is repeated.

When the market cannot distinguish between those who produce knowledge and those who rent it, it begins to reward the performance of innovation instead of the production of innovation.

Naturally, companies able to deliver dazzling presentations at low cost proliferate, while serious research laboratories appear slower, riskier, and less impressive in the short term.

A researcher trying to reduce memory use by 20 percent may spend months testing hypotheses, while another founder can present an agent that speaks five languages within days by relying on a model they had no role in developing.

The first may be closer to creating a technology the world needs. The second may be closer to creating a product a customer needs.

Each has a different economy.

Confusing them, however, sends a dangerous message to the market: **do not trouble yourself with developing the mind; learn to sell it under your own name.**

That message does more than distort investment decisions. It reshapes the ambitions of a generation of researchers and engineers.

Markets do not merely reward talent. They also teach talent what it should aspire to become.

If the route to being called an “artificial intelligence pioneer” passes more often through wrapping GPT than through developing an algorithm, architecture, or training technique, it should not surprise us when interface companies multiply and serious laboratories become harder to find.

We are then suffering not from a shortage of companies.

We are suffering from a surplus of claims to originality and a deficit in producing it.

Why the Model Maker Will Devour Its Borrowed Children

The literature on platform economics describes a strategy known as platform envelopment, studied by Thomas Eisenmann, Geoffrey Parker, and Marshall Van Alstyne. [6](#)

It occurs when a major platform expands into adjacent functions and markets by drawing on its technical infrastructure, its user base, and its ability to integrate services.

What is happening in artificial intelligence today reminds me of this dynamic.

In March 2025, OpenAI announced the [Responses API and Agents SDK](#), with built-in tools for web search, file search, and computer use, plus an SDK for orchestrating agents.

These were functions that startups had offered as standalone products or distinctive technical features. They then began turning into capabilities supplied by the model platform itself.

The equation is brutal in its simplicity.

If your competitive advantage is a function that the supplier from whom you rent intelligence can add in a software update, you are not building a technological fortress. You are renting temporary space on someone else's land.

Your product may be successful today, but its success does not mean that you control the conditions of its survival.

As foundation models improve, and as they become better at tool use, multi-step reasoning, and more autonomous agentic execution, the space that wrapper companies can sell as an exceptional addition will shrink.

Application companies may remain commercially successful because they understand their customers, master systems integration, or build strong distribution networks.

Those are advantages in software and business. They are not certificates proving that the company invented the intelligence its product uses.

A company whose greatest technical achievement is concealing the name of the original model will one day find itself forced to conceal a more embarrassing fact: the model no longer needs it at all.

I believe this sorting process will be one of the market's most important transformations in the years ahead.

The economic value of many generic artificial intelligence wrappers will weaken, while the value of companies that own original algorithms, more efficient technologies, specialised models, or cumulative industrial data and knowledge—assets that cannot be rented through an API call—will rise.

At the same time, application companies will come under greater pressure to prove that they offer value beyond the functions available in the underlying platforms.

As for those who present every artificial intelligence agent as a scientific discovery, the coming years may be less tolerant of their marketing eloquence.

In technology, a single update can demolish an investment story built over years.

Real Inventions Happen Behind the Screen

To understand the difference between making intelligence and making interfaces for it, we do not need to contemplate the world's largest language models.

Sometimes an invention is more profound and less theatrical.

At the 2022 Conference on Neural Information Processing Systems (NeurIPS), researcher Tri Dao and his colleagues published the [FlashAttention paper](#).

Their goal was not to design a new chatbot. It was to address a problem in the implementation of the attention mechanism used in the Transformer architecture.

The researchers found that a substantial part of the slowdown came not from the calculations themselves, but from the movement of data between high-bandwidth memory (HBM) and the fast static random-access memory (SRAM) inside the processor.

Through an input/output-aware algorithm based on tiling operations, they reduced data movement, improved the performance of exact attention, and achieved a speed-up of as much as three times in one of the paper's experiments.

They did not add a new button.

They changed the way a fundamental mathematical operation was executed. Broad systems across the model industry then benefited from that knowledge.

In 2023, researchers at the University of California, Berkeley presented [vLLM and the PagedAttention algorithm](#).

The inspiration came from an old idea in operating systems: virtual memory paging.

Instead of wasting large areas of memory to store the attention key-value cache (KV cache), the researchers developed a more efficient way to manage that data during inference.

Their published experiments reported throughput improvements of between two and four times over specified baseline systems at comparable latency levels.

Consider the difference between these works and a company announcing that it developed an artificial intelligence platform because it connected a commercial model to a database.

Both involve software.

In the first, however, there is an addition to the knowledge by which models operate. In the second, there is an application of existing knowledge.

That distinction is the essence of deep technology.

Invention does not require training an enormous model from scratch. It may lie in a novel architecture, a learning algorithm, a model-compression mechanism, knowledge distillation, a method for specialising models or continually pretraining them, or a discovery that reduces resource consumption while preserving performance.

Science does not advance by reinventing the world every morning. It advances by discovering a new relationship, overcoming an old constraint, or combining distant bodies of knowledge into a structure that opens possibilities which did not exist before.

This is why I like to distinguish between two companies that may look alike from the outside.

One asks: How can we make the current model perform a task for our customer?

The other asks: How can we develop the technology that makes this task possible in a way that was not available before?

The first is looking for a product.

The second may be building an industry.

Where technical value is created

Dimension	An application built on a model	A contribution to the technology itself
Question	How can the model perform a task for a customer?	How can we improve the technology that makes the task possible?
Focus of work	Interfaces, systems integration and workflows	Architecture, algorithms, training and evaluation
Potential value	A useful product and a sector or business advantage	Knowledge and technical assets that other products can build on

Misraj and Kawn Lab: From Consuming Intelligence to Producing Knowledge

This brings me to what I consider most important for the future of artificial intelligence in Saudi Arabia.

We need companies that do not treat artificial intelligence merely as a service to call, but as a scientific and engineering field to which they can contribute.

Among the local examples in which I see signs of this direction are **Misraj AI and its research laboratory, Kawn Lab**.

What deserves attention in this ecosystem is that it has not confined itself to building applications and agents on commercial models. It has pursued research on Arabic models, their training and specialisation, and the production of data and evaluation benchmarks—

essential components for any sector that wants to possess its own research capability.

In April 2025, Misraj researchers published a paper titled *Kuwain 1.5B: An Arabic SLM via Language Injection*.

The idea was research at its core.

Could a new language be injected into a small language model (SLM) while preserving its prior knowledge and avoiding the cost of retraining it in full?

The experiment began with the open-source TinyLlama model. The researchers then tested layer extension, the addition of Arabic vocabulary to the tokenizer, the freezing of parts of the original model, and the training of new layers to achieve more efficient language adaptation.

They added approximately 26,000 Arabic tokens to the vocabulary and studied the effects of the positions of the new layers, the distribution of training data, and architectural alternatives.

Their results reported an average improvement of around 8 percent across the Arabic tests they used, while preserving the model's previous capabilities.

What matters to me here is not celebrating the percentage in itself.

What matters is that we are looking at a team addressing a problem in learning and cross-lingual knowledge transfer, intervening in the model architecture and training mechanisms to arrive at a solution.

These are the kinds of questions that open the doors to a technology industry.

Other work has followed from the Misraj and Kawn ecosystem, including **Baseer**, a vision-language model (VLM) designed to understand Arabic documents and extract their structure.

In its 2025 paper, *Baseer: A Vision-Language Model for Arabic Document-to-Markdown OCR*, the researchers addressed the challenges of Arabic optical character recognition (OCR), from joined letters and varied typefaces to diacritics and complex page layouts.

The work started from a pretrained multimodal model, which was then specialised using Arabic data, alongside the development of an evaluation benchmark for document understanding called Misraj-DocOCR. The model achieved a published word error rate (WER) of 0.25, outperforming the comparison systems on that measure within the evaluation set used.

The ecosystem also has work in Arabic–English translation through **Mutarjim**, automatic diacritisation through **Sadeed**, and dialect models through **Lahjawi**, alongside work on embeddings, multimodal datasets, and research evaluations. [12](#) [13](#) [14](#)

None of this means that Misraj invented every foundation model it used. That is not the standard of originality I am defending in the first place.

The difference is that it does not merely repackage rented capability under a new name. It intervenes in the model's structure, training method, specialisation, or evaluation tools to produce knowledge and technology on which others can build.

Anyone who confuses the two does not understand the nature of modern research.

FlashAttention did not invent neural networks. PagedAttention did not invent virtual memory. ASML did not invent optical physics. Yet each made a technical contribution of profound value.

Invention is not the birth certificate of every component.

It is the contribution that makes a combination of knowledge produce a new capability.

This is the path that, in my view, Misraj and Kawn Lab model locally, and one around which similar research laboratories and companies deserve to multiply.

The issue is larger than the success of one Arabic model or a particular product.

It is about **building a research layer capable of producing the technology on which the companies of the future will stand.**

Why We Need Market Makers, Not Just More Applications

The economist Joseph Schumpeter wrote of innovation as a force that reshapes the economy through “new combinations,” and his name became associated with the concept of creative destruction. [16](#)

In his account of economic development, entrepreneurship is not confined to entering an existing market and competing for a share of it. Innovation can create new products, methods of production, and markets, changing the rules of competition themselves.

I believe this distinction is critically important for Saudi Arabia.

We have an opportunity to build a large number of artificial intelligence applications, and this is already happening across different sectors. But multiplying applications is not enough on its own to build a deep knowledge industry.

We could end up with hundreds of companies competing to provide services that all depend on the same global models, differing only in their interfaces, pricing, and target sectors, while the foundational knowledge and the centres that develop it remain outside those companies.

We might even reach a stranger paradox: a market crowded with companies presenting themselves as artificial intelligence companies, while all of them must look abroad when they want to develop the core of the technology they use.

A market with a thousand interfaces to the mind, and too few laboratories to develop it.

The importance of deep research companies therefore extends beyond their immediate products.

When a laboratory such as Kawn develops a specialised model, a training method, a dataset, or an evaluation benchmark, it is not creating a technical asset for itself alone.

It is helping to create an infrastructure around which an entire sector can emerge.

A good evaluation benchmark gives companies a clear technical target on which to compete. A specialised dataset opens the way for new experiments. A highly efficient small model can become the foundation for dozens of products. A faster inference algorithm can lower costs for companies that played no part in inventing it.

This is how knowledge markets form and knowledge spillovers arise.

The real value of research is not confined to the revenue of the company that conducted it. It includes everything that the research enables other companies and institutions to build upon it.

When a discovery becomes a tool, technology, or model that can be licensed, knowledge itself becomes a high-value export commodity.

This is precisely why companies such as Misraj and laboratories such as Kawn matter.

They can play a part in creating a market and a sector by developing technical assets, cultivating research talent, opening new scientific problems, and supplying application companies with technologies that give them genuine differentiation.

We need institutions that resemble, in their effects, the industrial laboratories whose discoveries created entire industries—not merely the funding of an ever-growing number of interfaces.

This brings me back to the economist Friedrich Hayek and his idea of dispersed knowledge, and to the philosopher Michael Polanyi and his concept of tacit knowledge. [15](#) [17](#)

Technical knowledge does not exist entirely in research papers or software code. Part of it lives in researchers' experience, engineers' intuition, failed experiments, and the ability to diagnose problems before their causes become clear.

These faculties cannot be purchased with a monthly subscription to a cloud platform.

They accumulate inside research institutions that treat the unknown as their daily work.

If we want to build an industry capable of renewing itself, we need institutions that preserve and develop this accumulation—not companies whose technical identity changes whenever a new commercial model appears.

The Investor Who Funds Appearances Will Get a Market of Appearances

I do not believe the problem is a shortage of money allocated to artificial intelligence. It lies in the questions we ask when deciding where that money should go.

When a company enters the market claiming that it has developed original artificial intelligence technology, the first question should be: Where does your technical contribution lie?

Did you develop a model? Did you change its architecture? Did you invent a training method, an algorithm, or an inference system? Do you own model weights, specialised training data, or an independent technical infrastructure? Can you identify the difference between the technology you made and the capability you bought?

These are not trick questions. They are questions about the identity of the company itself.

I see no reason why answering them should be embarrassing unless the marketing story is larger than the actual achievement.

If the company merely talks about how many artificial intelligence agents it has, its enterprise integrations, its feature lists, and the model's ability to write and analyse, then it has described its product. It still has not told us what it invented.

It is amusing that some presentations explain in detail what GPT can do, then attribute all of those abilities to the company giving the presentation.

It is as if the model passed through an API from one institution to another and lost its scientific lineage along the way.

I believe we need a different kind of investment awareness toward deep technology.

A research company should not be treated as a software company that is late to launch its product, nor should it be asked to achieve within months what some discoveries require years to accomplish.

Conversely, a model-wrapper company should not receive a high research valuation simply because it knows how to use scientific vocabulary in its presentations.

We need capital capable of bearing scientific risk, understanding research and development cycles, distinguishing intellectual property from software integration, and recognising that some of the most valuable technical assets are not ready to become mass-market products at birth.

More important still, we must give researchers a clear path from the laboratory to the market.

If a Saudi researcher develops an algorithm that accelerates inference by a meaningful percentage, an architecture that uses memory more efficiently, or a method that makes small Arabic models competitive at a specialised task, their only options should not be to seek a job at a global company or leave the discovery in a repository of academic papers with no one to turn it into an industry.

That researcher should find the investor who understands the work, the industrial company that needs it, and the market that rewards it.

We do not need more founders who excel at persuading investors that they invented artificial intelligence. We need more researchers who can invent something that makes such persuasion unnecessary.

What I See Coming by 2030

I believe the future of the artificial intelligence industry will carry a major economic paradox.

As general-purpose foundation models become more capable, those capabilities will become available to more developers and companies, and many functions now sold as exceptional advantages will become less scarce.

Building a multi-tool agent, a system that retrieves enterprise knowledge, or an assistant that carries out a chain of tasks may come to resemble the use of standard components from a cloud platform.

In that case, many companies will lose part of the advantage they built by assembling those components.

I expect the current wave of wrapper companies to undergo a sharp sorting: some will become successful vertical-software companies, some will turn into service and integration companies, and some will lose their advantage when the major platforms provide the same function.

Claiming to own original artificial intelligence will not protect any company from this shift.

At the same time, four areas of knowledge will become more strategically important: model architectures; training and inference efficiency; precise domain specialisation in data and models; and the computing infrastructure that makes artificial intelligence less costly and more reliable to run.

I do not think the competition will be merely a race to increase parameter counts. It will be a race over the relationship among capability, cost, energy, latency, and reliability.

A technology that reduces memory consumption may have greater economic impact than an application that adds a new feature to a chat interface.

A specialised small model may acquire more strategic value than a larger one when it performs a precise industrial task at lower cost and can be deployed on the organisation's own premises.

An algorithm that gives thousands of companies a modest improvement in inference cost may be more valuable over the long term than dozens of applications that depend on a single model.

This is the future for which I believe Saudi Arabia should prepare.

We have a large industrial economy; accumulated expertise in energy, petrochemicals, water, transport, and financial services; and operating environments that contain problems a general-purpose conversational model will struggle to solve.

There are opportunities in industrial artificial intelligence, time-series models, learning from sensor data, low-latency inference, Arabic document understanding, multimodal systems, and efficient deployment on edge devices.

These fields can become sources of exportable discoveries if we treat them as research problems, not merely as opportunities to create new applications.

This is why I believe companies and laboratories such as Misraj and Kawn can form an important nucleus for a broader transformation.

Not because they will monopolise innovation, and not because one model will transform an entire industry, but because they are adopting the direction we need to multiply: the move from consuming model capabilities to developing the knowledge that determines what models can do.

At that point, the nature of ambition changes.

Instead of asking: How many Saudi companies provide GPT solutions to enterprises?

We ask: How many technologies developed by Saudi teams are used by institutions and companies outside Saudi Arabia?

How many models, algorithms, evaluation tools, or inference architectures have travelled from a Saudi laboratory into the global technology ecosystem?

How many startups were built around a scientific discovery rather than a commercial subscription to a language model?

These are not cosmetic indicators.

They measure the position we occupy in the knowledge value chain.

We could become one of the region's largest artificial intelligence markets without becoming a producer of its influential technologies.

We could invest billions in computing infrastructure without possessing new engineering knowledge that others need.

Scale gives us presence. Original knowledge gives us influence.

Scientific Insight Is Not Born from a Proof Form

There is an idea that, in my view, contemporary technology discourse sometimes misunderstands: that every vision of the future must begin with complete empirical proof.

That is a sound principle for verifying a specific claim, but it is not a complete account of how discoveries are born or strategies are built.

Science itself has recognised the role of scientific intuition, insight, abductive reasoning, and the ability to see relationships before they have been formulated as a complete theoretical model.

The scientist and philosopher Michael Polanyi held that part of knowledge extends beyond what we can state in rules and instructions, and that experience creates a capacity for perception that cannot easily be reduced to a complete mechanical description. [17](#)

Industrial history is full of decisions whose makers began not with statistical certainty about the future, but with an engineering or scientific vision whose value was not clear to most of their contemporaries.

If the pioneers of every technology had demanded that the future present all its evidence before they invested in it, pioneering would have lost its meaning.

My wager on the rising value of deep technology is therefore not a claim that the future has already been written. It is a reading of an economic direction that I see forming through the declining scarcity of general capabilities, the expanding control of platforms, and the rising value of knowledge that is difficult to imitate.

This is a strategic reading, and a good strategic reading does not wait until a trend has become old news before daring to name it.

There remains, however, an unmistakable difference between anticipating the future of an industry and claiming that you invented a technology which already exists elsewhere.

Intuition is an inalienable right of the scientist, the leader, and the entrepreneur.

Inventing a fictional scientific biography for a product built on GPT is not intuition. It is relabelling the country of origin.

To Make What the World Needs

I return to ASML, the company that began its work in a modest Dutch shed.

It did not become a global power because it produced the applications that run on chips, or because it bought the best products made by others and remarketed them under a different name.

It chose a harder and deeper position in the value chain.

It chose to produce cumulative engineering knowledge and turn it into machines needed by companies that make products used by billions of people.

This was not the fastest route to fame or revenue.

But it created a kind of influence that the proliferation of applications alone cannot confer.

This is what I want for Saudi Arabia's artificial intelligence industry.

I do not want us to measure our success by the number of Arabic-speaking interfaces built over GPT, the number of agents with local names, or the number of companies declaring themselves "leaders in artificial intelligence" while the extent of their knowledge asset is a commercial subscription to a model developer.

I want to see Saudi research companies produce models, algorithms, architectures, training methods, and evaluation tools that make the world look to us as a source of technical knowledge, not merely as a large market that consumes it.

This is why I see Misraj AI and Kawn Lab as an important example of the spirit we need to cultivate: the spirit of a laboratory that asks the question before designing the interface, investigates the structure of the technology before giving it a commercial name, and aspires to produce knowledge on which other companies can build their products.

An application company reaching thousands of customers may be a significant commercial achievement.

But a laboratory that develops a technology needed by thousands of companies can open an entire market.

That is the distance between participating in an existing market and helping to create a new one.

We do not need to oppose the use of global models. We need to end the ridiculous masquerade in which a technology importer wears the coat of its inventor.

Whoever builds an application on Claude should say that they built an application. Whoever offers an agent using GPT should take pride in the product. Whoever succeeds in developing a profitable SaaS business should celebrate the customers and revenue.

But any company that wants the title of deep technology should show its depth in the technology, not in the vocabulary it uses to describe its interface.

Rented minds can build successful businesses, but they do not give their renters the right to claim they made the mind.

When I look at the decade ahead, I am not preoccupied with our becoming the country with the greatest number of applications calling on the world's intelligence.

I am preoccupied with our becoming the country that develops some of the knowledge that intelligence needs in order to become better.

Nations do not rise to the position of knowledge makers through the number of names they place on other people's inventions, but when their own inventions earn a name in what others make.

The difference between the two positions may look small on an application screen.

In the balance of industrial history, it is the difference between remaining a permanent customer of the future and becoming one of those who make it.

References and sources

1. [Our history](#) — ASML.
2. [Light & lasers — Lithography principles](#) — ASML.
3. [Lenses & mirrors — Lithography principles](#) — ASML.
4. [Cymer](#) — ASML.
5. [The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism](#) — The Quarterly Journal of Economics / Oxford Academic.
6. [Platform envelopment](#) — Strategic Management Journal / Wiley.
7. [New tools for building agents](#) — OpenAI.
8. [FlashAttention: Fast and Memory-Efficient Exact Attention with IO-Awareness](#) — NeurIPS.
9. [Efficient Memory Management for Large Language Model Serving with PagedAttention](#) — ACM SOSP 2023.
10. [Kuwait 1.5B: An Arabic SLM via Language Injection](#) — arXiv.
11. [Baseer: A Vision-Language Model for Arabic Document-to-Markdown OCR](#) — arXiv.
12. [Kawn Lab](#) — Misraj AI.
13. [Research and Publications](#) — Kawn.
14. [Lahjawi: Arabic Cross-Dialect Translator](#) — Association for Computational Linguistics.
15. [The Use of Knowledge in Society](#) — Online Library of Liberty.
16. [Capitalism, Socialism, and Democracy](#) — University of Connecticut course text.
17. [Primary Resources: Transcendence And Self-Transcendence](#) — The Polanyi Society.