Center for Applied Complexity & Intelligence
Zentrum für Angewandte Komplexität & Intelligenz

Welcome to the Future?

Why we confuse optimization with progress

An introduction to the work of the
Center for Applied Complexity & Intelligence

January 2026

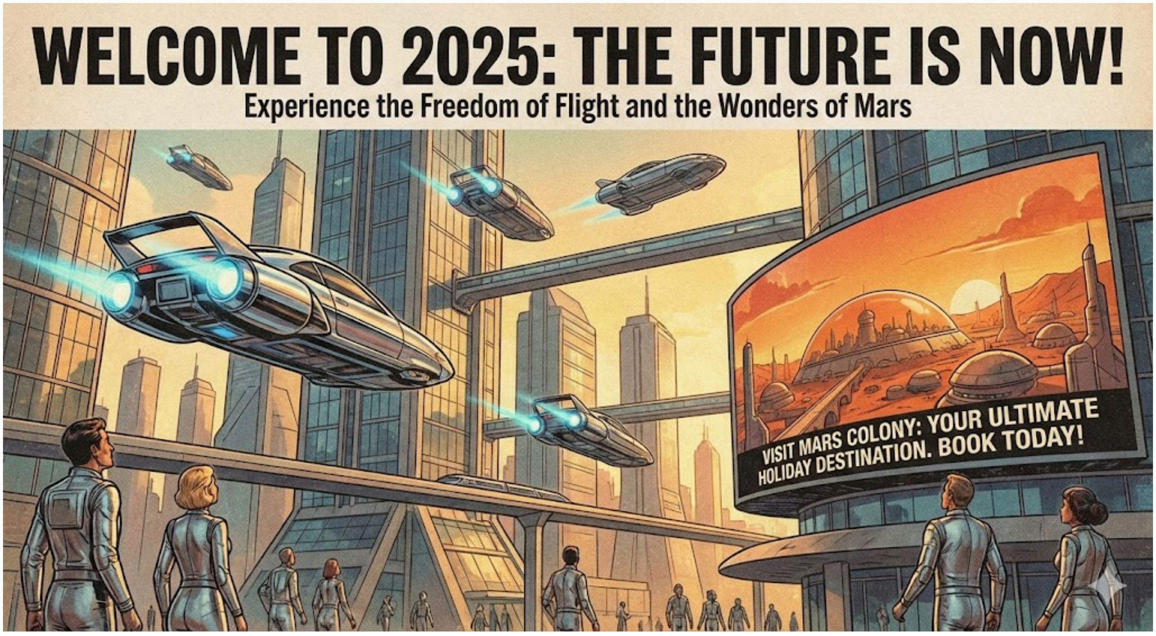

Figure 1: The future of the 20th century: flying cars, Mars tourism. The picture shows 2025. Things turned out differently.

The promised future

The image on this page comes from a vision that shaped the entire 20th century. It shows flying cars hovering elegantly between skyscrapers, billboards advertising vacation trips to Mars colonies, and people in silver suits looking toward a bright future. This vision was not a fringe phenomenon, but rather the mainstream of technological optimism. From the world's fairs of the 1930s to the science fiction of the 1950s to the visions of the future in the 1970s, there was a consensus: the year 2025 would be a completely different world.

It is now 2026. Flying cars have not arrived. Mars colonies do not exist. The subways we take to work date back to the 19th century, as Peter Thiel once remarked, while we send pictures of cats on our smartphones. This is not an exaggeration, but a sober assessment of the situation. The question that begs to be asked is not why individual technologies have failed. The question is more fundamental: Why have we stopped thinking in truly innovative ways?

The answer proposed by the Center for Applied Complexity & Intelligence is uncomfortable but necessary. We haven't stopped working. We haven't stopped optimizing. We have not stopped investing money in research and development. What we have stopped doing is questioning the fundamental architectures on which our technologies, our organizations, and our societies are built. Instead, we have confused optimization with progress.

The illusion of progress

At first glance, the diagnosis of stagnation seems absurd. We have smartphones with more computing power than the Apollo missions. We have artificial intelligence that writes texts and generates images. We have electric cars and solar panels. Every day, new products, new apps, and new services appear. How can anyone talk of stagnation?

The answer lies in distinguishing between two fundamentally different types of change. The first type is optimization within an existing paradigm. You take an existing architecture and make it faster, smaller, cheaper, more efficient. The second type is paradigm shift, the introduction of a fundamentally new architecture that solves old problems in completely new ways or opens up possibilities that were previously unthinkable. The 20th century was rich in paradigm shifts. The 21st century has not been so far.

The figures speak for themselves

Economist Tyler Cowen systematically examined this observation in his influential 2011 book "The Great Stagnation." His thesis is that since around 1973, the American economy, and with it the Western world as a whole, has harvested the "low-hanging fruit" of technological progress. The major breakthroughs that radically changed life between 1870 and 1970 have largely failed to materialize.

Cowen illustrates this with a simple thought experiment. If you compare life in 1870 with life in 1920, the differences are revolutionary. Railroads and automobiles fundamentally changed what distance means. Telegraphs and telephones fundamentally changed what communication means. Electric light changed our relationship to night. Indoor plumbing changed hygiene and life expectancy. In contrast, when comparing life in 1970 with life in 2020, Cowen says the differences are more superficial. We have better televisions and interactive screens. But the fundamental structures of everyday life, such as how we live, how we get around, how we work, have changed far less.

The United States Bureau of Labor Statistics has attempted to quantify this intuition. According to its findings, the slowdown in productivity since 2005 has led to a cumulative loss of $10.9 trillion in economic output in the American corporate sector. This corresponds to a loss of $95,000 per employee. These figures are abstract, but they point to a real phenomenon.

Paradigms that refuse to die

The stagnation becomes particularly apparent when you consider the age of the fundamental technologies that underpin our digital infrastructure. The Von Neumann architecture, on which virtually all computers operate, was developed in 1945. That makes it 80 years old. The Unix operating system model, on which Linux, macOS, Android, and large parts of the internet are based, dates back to 1969, making it 56 years old. The C programming language, in which operating systems and critical infrastructure are written, was developed in 1972. SQL, the language for database queries, dates back to 1974. The TCP/IP protocol, the foundation of the internet, also dates back to 1974.

These technologies are not only old. They are dominant. COBOL, a programming language dating back to 1959, is still estimated to process the majority of all financial transactions worldwide. Banks, insurance companies, and government agencies run on code that is older than most of their employees. This is often celebrated as a sign of quality and stability. The argument is that these technologies have proven themselves, that they are robust, that they work.

But this interpretation overlooks a crucial point. These technologies do not survive because they are optimal. They survive because the switching costs have become astronomical. Billions of lines of code, trillions of dollars in infrastructure investments, decades of accumulated knowledge and habit have created a system in which real alternatives are hardly conceivable. This is not a sign of strength. It is lock-in, a captivity in path dependencies that does not promote innovation, but rather prevents it.

Peter Thiel, the investor and technology critic, summed up this situation succinctly: “We wanted flying cars, but instead we got 140 characters.” The statement referred to Twitter and was intended as a criticism of Silicon Valley’s superficiality. But it hits a deeper point. The computing power of a modern smartphone surpasses anything that was available during the Apollo missions. But what is it used for? For social media, for games, for sending photos. The hardware is revolutionary. The applications are not.

Stagnation beyond computer science

Criticism of technological stagnation is often understood as a problem of computer science or Silicon Valley. However, the patterns observed by the Center for Applied Complexity & Intelligence are evident in virtually all areas of modern society. Stagnation is not limited to software. It pervades energy systems, mobility, medicine, construction, and education. The same structures can be found everywhere: fundamental architectures that have remained unchanged for decades, while optimizations are made on the surface.

Energy grids: An architecture from the 19th century

The electrical grids that supply modern societies with power are based on an architecture that has hardly changed since its inception. The basic principle is simple: large, centralized power plants generate electricity, which is transported to consumers via high-voltage lines. This model was designed for a world in which coal and later nuclear power plants produced constant electricity and consumption was predictable.

In its report "Electricity Grids and Secure Energy Transitions," the International Energy Agency found that 70 percent of all power lines and transformers in the US are more than 25 years old. Much of the infrastructure was built in the 1960s and 1970s and is approaching the end of its 50- to 80-year lifespan. The Council on Foreign Relations summarizes it this way: The basic structure of the grid has remained largely unchanged for decades, even though it has grown from Edison's original 59 customers to hundreds of millions of users.

The problem is not that the grids don't work. The problem is that they were built for a world that no longer exists. They were designed for a one-way flow of electricity from central power plants to passive consumers. Today, they have to cope with bidirectional flows, integrate variable generation from solar and wind farms, and coordinate millions of decentralized resources. The architecture is no longer fit for purpose. But the costs of change are so high that we continue to optimize instead of rebuilding.

The figures are impressive. By 2040, more than 80 million kilometers of electrical lines would need to be added or modernized worldwide. That is roughly equivalent to the length of the entire existing grid. At the same time, more than 3,000 gigawatts of renewable projects are waiting for grid connections because the infrastructure is not growing fast enough. After more than a decade of stagnation, investment in grids has only recently begun to increase again on a global scale, but it would need to double to over $600 billion per year by 2030 to meet climate targets.

Aviation: Slower than 50 years ago

Commercial aviation provides a particularly vivid example of stagnation. It may sound surprising, but modern passenger aircraft actually fly slower than their predecessors in the 1960s and 1970s. A flight from New York to Los Angeles, which took five hours and 43 minutes in 1967, now takes 40 minutes longer. The stated cruising speeds of modern commercial aircraft range from about 480 to 510 knots, compared to 525 knots for the Boeing 707, a standard aircraft of the 1960s.

The reasons for this are understandable. Aerodynamic drag increases significantly as you approach the speed of sound. Flying faster consumes considerably more fuel per passenger mile. After the oil crisis of the 1970s, when fuel prices skyrocketed, airlines began systematically reducing cruising speeds to save costs. This was a rational decision under the circumstances.

But the result is remarkable. In half a century of technological progress, we have not become faster, but slower. The Concorde, which flew at twice the speed of sound between 1969 and 2003, was discontinued and has not been replaced. Since 2003, there has been no commercial supersonic air travel. Airplanes have become safer, more comfortable, and more efficient. But the fundamental experience of flying has not improved. The paradigms have remained the same.

The explanation does not lie in the limits of physics. It lies in the limits of our willingness to develop new architectures. Boom Supersonic and other companies are working on new supersonic aircraft. But the fact that we are still waiting for them 20 years after the end of Concorde shows how slow real progress has become. The industry is focused on efficiency improvements within the existing paradigm. The next Boeing 777X will travel at Mach 0.85, just like its predecessors.

Construction: The productivity-resistant industry

If there is one sector that particularly embodies stagnation, it is construction. While productivity in the American economy has tripled since 1948, and this positive trend is visible in almost all sectors, construction is a notable exception. Labor productivity in the construction sector has not increased. It has declined.

A 2025 study by Harvard Business School analyzes this phenomenon under the title “Why Has Construction Productivity Stagnated?” The authors note that large buildings today are constructed more or less the same way they were 100 years ago: assembled by hand by an army of workers using steel, concrete, and masonry. The industry has benefited virtually nothing from advances in manufacturing or logistics that make it possible to produce something as complex as an iPhone at an affordable price.

McKinsey & Company has calculated that global labor productivity in construction has grown by only about one percent per year over two decades. In comparison, growth in manufacturing was 3.6 percent. While factories transformed into highly productive machines using automation, robotics, and digital twins, construction remained stuck with methods that transport materials to muddy job sites and start from scratch every time.

What is interesting about this case is that the stagnation is not limited to the US. The recent trend in labor productivity in the US, which has largely stagnated, is fairly consistent with trends in other large, wealthy countries. Sweden, for example, is often praised for its high proportion of prefabricated construction, but has also seen largely flat productivity growth since the 1990s. Japan also makes more use of prefabrication than the US and has even experimented with automated high-rise construction, but has seen virtually no improvement in construction productivity since the 1970s.

There are many reasons for this. Increasing land use regulations could be a plausible reason, as stricter regulations discourage construction companies from pursuing larger projects, keeping them relatively small and reducing incentives for technological innovation and economies of scale. Data on patent activity in the construction industry also suggests that as companies and projects shrink, innovation declines.

Pharmaceutical research: The forgotten crisis of antibiotics

The last new class of antibiotics to make it to market was discovered in 1987. Since then, there has been what scientists call a "40-year innovation gap." This stagnation is no minor issue. Antibiotics are among the most important medical achievements of the 20th century. They have saved countless lives and made modern surgery, cancer therapy, and intensive care medicine possible in the first place. Without effective antibiotics, routine operations become life-threatening procedures.

The Wellcome Trust Foundation has investigated why the development of new antibiotics has become so difficult. Part of the explanation lies in science itself. The "low-hanging fruit" was harvested between the 1960s and 1980s, when systematic screening programs of natural compounds from soil samples discovered a wealth of new active substances. Promising compounds are harder to find today.

But scientific difficulties only explain part of the story. As the pharmaceutical industry became better equipped than ever before with advances in medicinal chemistry, molecular biology, and genomic tools, expectations for productivity rose. The industry turned away from the laborious search for naturally occurring compounds and instead focused on target-based high-throughput screening of synthetic compounds. But these high-tech platforms delivered only disappointing results.

The underlying reason for the stagnation is economic. Developing new antibiotics is expensive, with research and development costs often exceeding a billion dollars. Unlike drugs for chronic conditions, antibiotics are typically prescribed for short courses of treatment, making them less profitable. As a result, many pharmaceutical companies have shifted their focus to more lucrative therapeutic areas such as oncology. Between 2010 and 2014, the FDA approved only six new antibiotics, a stark contrast to the 19 approvals between 1980 and 1984.

The combination of poor financial incentives and technical challenges has led many large pharmaceutical companies to reduce or abandon their antibiotic development programs altogether. In 2019, small and medium-sized companies dominated the field, accounting for about 90 percent of new antibiotics in development. However, these companies do not have the resources for the lengthy and expensive clinical trials required for approval. The result is a pipeline that is nowhere near sufficient to address the growing threat of antibiotic-resistant pathogens.

Nuclear energy: Generation II forever?

The nuclear industry offers another revealing example. The vast majority of nuclear reactors in operation worldwide today belong to what is known as Generation II. These reactors were commissioned in the late 1960s and are based on light water reactor technology developed during that period. The reactors built in the 1980s and early 1990s are essentially of the same type and make up the majority of the more than 400 commercial reactors worldwide.

Generation III reactors have existed as a concept for decades. They are referred to as “advanced light water reactors” and offer improvements in fuel technology, thermal efficiency, modular design, and safety systems, particularly through the use of passive rather than active safety systems. However, due to the long period of stagnation in the construction of new reactors and the continued popularity of Generation II designs in new builds, relatively few third-generation reactors have actually been built.

The situation is even more remarkable when one considers that many “new” reactor designs are based on concepts that are also decades old. Most designs for liquid salt reactors are derived from the Molten Salt Reactor Experiment of the 1960s. Nuclear energy is considered by some to be a mature technology that has been in a state of stagnation for many years. It could be the first contemporary, established energy generation technology to face global decline if new approaches are not found.

Education: The classroom of 1920

If the grandparents of today’s students were to walk into a typical lecture hall in 2026, they would immediately know what to do: sit down, be quiet, and listen to the lecturer. The basic structure of teaching has hardly changed in 100 years.

A survey of 275 economics professors revealed that they spend 70 percent of class time lecturing, 20 percent on discussions, and 10 percent on student activities. Studies from the 1990s report similar figures. Lectures remain the most common form of teaching in higher education and account for the largest share of teaching time.

The 1920s laid the foundation for many educational practices, and our current system still relies heavily on ideas from that era. The structure of the traditional school day, the emphasis on standardized testing, and even the physical layout of classrooms can be traced back to a time when education was designed for an industrialized society. Technology has changed, overhead projectors have been replaced by PowerPoint, and books by PDFs. But the basic architecture of learning, an expert imparting knowledge to passive recipients, has remained the same.

Why are we stagnating?

The examples from various sectors show a common pattern. It is not about individual technologies reaching their limits. It is about a systematic phenomenon that affects virtually all areas of modern society. The question is: Why? Why do we endlessly optimize within existing paradigms instead of creating new ones?

The logic of lock-in

Economist Brian Arthur has extensively studied the phenomenon of path dependency and technological lock-in. Once a technology has become established, network effects and switching costs arise that cement its position. Every new installation, every training course, every regulation based on the existing technology increases the cost of switching. At some point, switching becomes so expensive that it becomes practically impossible, even if better alternatives were theoretically available.

This explains why COBOL still processes financial transactions. It explains why power grids are based on 19th-century architecture. It explains why we build buildings the same way we did 100 years ago. The switching costs are astronomical. But that still doesn't explain why we stopped looking for alternatives. Lock-in is an obstacle, but not an absolute one.

Recognizing warning signs

One legitimate criticism of analyses such as this concerns their retrospective nature. In hindsight, it is always clear what was optimization and what was genuine progress. The Concorde was a mistake, the internet a breakthrough. COBOL is lock-in, Unix is a foundation. But no one can reliably recognize paradigm shifts as they are happening. How can you distinguish in real time between a promising innovation and a dead end?

The Center does not claim to be able to answer this question definitively. There is no surefire method for recognizing genuine progress in advance. What does exist are heuristics, warning signs that indicate potential dead ends. These signals do not guarantee anything, but they do increase the likelihood of recognizing misdirections earlier on.

The first warning sign is extreme centralization. When a technology can only be developed by a handful of players because the resource requirements exclude everyone else, a monoculture emerges. Alternatives are not refuted, but stifled. In the energy sector, this pattern is evident in the dominance of large centralized power plants, which structurally disadvantage decentralized approaches. In the semiconductor industry, a few companies control the entire production chain. In AI research, the cost of training large models is so high that only a few companies can compete at the top.

The second warning sign is skyrocketing switching costs. When every new investment deepens dependence on the existing architecture, a self-reinforcing system emerges, regardless of whether it is on the right track. COBOL-based banking systems are not being replaced because every new interface deepens dependence. Over decades, combustion engine infrastructure has created an ecosystem that systematically makes alternatives more expensive. Every new investment in existing paradigms makes change more expensive, not cheaper.

The third warning sign is the suppression of criticism by success. When a technology delivers impressive results, the voices questioning its fundamental limitations fall silent. The success of the combustion engine has prevented its systemic costs from being seriously assessed for over a century. The efficiency gains of industrial agriculture have long overshadowed questions about soil health and biodiversity. And the fact that large language models produce impressive texts says nothing about whether they are on the path to true intelligence or stuck in a dead end.

The fourth warning sign is the self-reinforcing spiral of attention. As soon as a paradigm gains visibility, a dynamic sets in that increasingly detaches itself from the technical substance. Investments flow, chairs are established, companies are created whose business models are often based more on promise than substance. Young professionals orient their education toward what is in demand, not what might be right. The media amplifies the attention, regulatory authorities adapt laws, educational programs are reorganized, conferences and journals specialize. Each of these steps increases the legitimacy of the dominant approach and deprives alternatives of further resources. Technologies that could offer valid alternatives are relegated to the shadows, not because they have been disproved, but because no one talks about them anymore. The concentration intensifies until the lock-in is virtually complete.

None of these warning signs is pathological in itself. What matters is their accumulation and their intensification over time. If several of these signs occur simultaneously and reinforce each other, caution is advised. These heuristics are no substitute for certainty. But they shift the focus from the question “What is right?” to the question “Where is caution required?” This is more scientifically modest and more useful in practice.

The institutional dimension

Peter Thiel and other proponents of the “social theory” of stagnation argue that there are no physical or scientific reasons why we cannot increase the pace of innovation. Stagnation is ultimately a social problem, they say. Societies have prioritized security, regulation, and risk avoidance over the potential rewards of groundbreaking innovation.

There is some truth to this criticism. Regulatory density has increased in many areas. Approval procedures for new technologies have become more lengthy and expensive. The liability risk for companies that try something truly new has risen. Institutional incentives favor incremental improvements over radical innovations.

But the institutional explanation also falls short. It does not explain why stagnation is occurring simultaneously in such diverse areas, in private and public sectors, in heavily and lightly regulated industries, in different countries with different institutional arrangements. There must be something deeper at work.

The limits of mechanistic thinking

The Center for Applied Complexity & Intelligence argues that stagnation has its deepest roots in a model of thinking that has reached its limits. This model treats systems like machines: as collections of components that interact according to fixed rules and whose behavior can be understood by analyzing the parts. It is the model that made the industrial revolution possible and made the natural sciences so successful.

But complex systems do not function like machines. Energy grids, cities, ecosystems, organizations, brains—they all exhibit properties that cannot be derived from the properties of their parts. They exhibit emergence, the occurrence of behaviors at the system level that do not exist at the component level. A flock of birds is not choreographed. It arises from simple local rules that each individual bird follows. An ecosystem is resilient without there being a designer. A brain thinks, even though no single neuron knows what thinking is.

The mechanistic model cannot grasp emergence. It cannot understand systems whose behavior arises from interaction, not from instruction. It cannot develop architectures that organize themselves, that learn, that adapt. And that is why we endlessly optimize within the old paradigms. We don't know how to think differently.

This comparison should not be misunderstood as a clear dichotomy. Modern engineering sciences have long integrated elements of complexity thinking. Control theory, cybernetics, distributed systems, site reliability engineering—all these disciplines work with feedback, emergence, and nonlinearity. Many supposedly mechanistic systems have long been complex adaptive systems in practice. The problem is not that complexity thinking is unknown. The problem lies in where mechanistic assumptions are applied beyond their domain of validity, where systems are treated like machines even though they have long since exhibited the characteristics of living networks. The criticism is not directed against mechanistic thinking as such, but against its unreflective universalization.

> When a paradigm reaches its limits, faster optimization is not enough. New architectures are needed. And new architectures require a new understanding of how systems work.
>
> — Our guiding principle

What is actually new here

A skeptical reader might object at this point that none of what is described here is really new. Emergence, path dependency, lock-in, complexity—these terms have been present in scientific discourse for decades. The Santa Fe Institute has been researching complex adaptive systems since the 1980s. Donella Meadows has written about systems and leverage points. Stafford Beer has developed cybernetic models for organizations. W. Ross Ashby formulated the fundamentals of cybernetics before most of today's researchers were born.

This objection is justified, and it would be intellectually dishonest to ignore it. The Center for Applied Complexity & Intelligence does not claim to have invented a new theory. What we claim is something else: we consistently apply existing systemic thinking to areas where it has been ignored until now. The theories have been around for decades, but they are not being applied where they are most urgently needed.

Let’s look at reality. Complexity science is a recognized discipline with university chairs, conferences, and journals. At the same time, energy grids run on 19th-century architectures. At the same time, banks process their transactions using COBOL from 1959. At the same time, AI research invests billions in a single architecture without seriously looking for alternatives. The knowledge exists, but it does not reach the systems it could transform.

We are not the first to attempt this transfer. But we are among the few who pursue it systematically, across domains, and with a view to real architectures. Our achievement therefore does not lie in inventing new theories. It lies in the systematic transfer of known complexity principles to entrenched real architectures. We are not developing new mathematics or discovering new laws of nature. What we are developing are new architectural perspectives on existing systems and a new intervention logic for dealing with stagnant paradigms.

The tension: When progress goes backwards

The analysis so far could give the impression that more innovation is always better, that stagnation is the only problem, and that we just need to move faster. But this view would be a dangerous simplification. Reality is more complex, and the Center for Applied Complexity & Intelligence is aware of this complexity.

There are areas where we may already have reached an optimum. Further changes worsen the situation rather than improving it. In some cases, we are not experiencing stagnation, but active regression disguised as innovation. The examples are numerous and sobering.

The car that no longer knows its owner

Modern automobiles have become rolling computers. They are networked, software-controlled, and constantly updated. That sounds like progress. But the reality is often different.

In 2022, BMW began offering seat heating as a monthly subscription. The hardware was already installed in the vehicle, but to use it, customers had to pay $18 per month or $415 for permanent activation. The outrage was enormous. People felt it was unreasonable to pay for a feature that already physically existed in their property. In September 2023, BMW withdrew the model. The board of directors explained: “User acceptance was not high. People felt like they were paying twice. That was not actually the case, but perception is reality.”

The subscription model for seat heating has disappeared, but the trend remains. Other manufacturers such as Audi, Hyundai, and Volvo have introduced similar models. The state of New Jersey has introduced a bill that would make it illegal to offer subscriptions for features whose hardware is already installed in the vehicle. The pendulum is swinging back.

Even more problematic is the growing fragility of software-controlled vehicles. In December 2024, Tesla Cybertruck owners reported that a software update had completely crippled their vehicles. The firmware had been damaged during installation, and the vehicle could no longer be started. In one documented case, the 48-volt battery even discharged, so that the owners could not even remove the charging cable.

A year earlier, Tesla had sent an update to over two million vehicles to introduce stricter driver monitoring. But the update did not go smoothly. Owners reported vehicles trying to download the update for days, draining the battery in the process and continuing to attempt to update even after Wi-Fi was disabled via internal cellular connections. A vehicle that no longer obeys its owner.

The tractor that no longer belongs to the farmer

The problem goes far beyond luxury cars. In January 2025, the U.S. Federal Trade Commission, together with the states of Illinois and Minnesota, filed a lawsuit against John Deere. The allegation: the tractor manufacturer is illegally preventing farmers from repairing their own machines.

Modern agricultural machinery is highly dependent on software. Many spare parts must be electronically “paired” with the tractor, similar to a printer driver with a computer. These pairing files, called “payload files” by Deere, can only be installed using special software that is available exclusively to authorized dealers. Farmers who perform repairs must call and pay a dealer technician just to enter a digital code that authorizes the mechanical repair.

The consequences are serious. A report by the U.S. Public Interest Research Group estimates that American farmers lose $3 billion annually due to tractor downtime and pay $1.2 billion more for repairs because they are dependent on dealers. One farmer reported that a single instance of downtime while waiting for a dealer cost him between $30,000 and $60,000.

Particularly noteworthy is the legal position that John Deere has taken since 2015. The company’s lawyers argue that farmers do not actually own their tractors because the machines run on software. They merely have an “implicit license for the lifetime of the vehicle to operate the vehicle.” The machine in the field belongs to the farmer. The software, without which it cannot function, belongs to the company.

Colorado passed the first “right to repair” law for agricultural machinery in 2023. Canada has enacted laws allowing farmers to bypass digital locks. The pendulum is swinging back, but slowly.

The printer that sabotages its own DRM

The printer industry offers an almost absurd example of technological regression disguised as progress. Modern printer cartridges contain chips that communicate with the printer and determine when a cartridge is “empty.” In practice, this often means that printers refuse to work even though 20, 40, or in extreme cases 80 percent of the ink is still available. The EU consumer organization BEUC found in tests that 42 percent of the cartridges tested still contained significant amounts of ink when the printer declared them empty.

Manufacturers such as HP have sent out firmware updates that changed communication protocols without warning, so that only the manufacturer’s own cartridges were accepted. Customers who had already purchased third-party cartridges suddenly found their printers unusable. Lawsuits are piling up. HP settled for $1.5 million in 2018. Canon was sued in 2021 for a particularly brazen practice: the manufacturer had disabled the scanner function of multifunction devices if no authorized ink cartridge was inserted, even though a scanner does not require ink.

The most absurd moment came in 2022, when a global chip shortage prevented Canon from producing enough cartridges with the necessary DRM chips. The solution? Canon published official instructions explaining to customers how to bypass the company’s own DRM. In doing so, the company essentially admitted that the chips were completely irrelevant to print quality. Their sole purpose was to tie customers to more expensive original cartridges.

Devices that have a shorter lifespan than their predecessors

A long-term Norwegian study examined the lifespan of household appliances over several decades. The results are revealing. In the 1990s and early 2000s, the lifespan of washing machines plummeted from an average of 19.2 years to 10.6 years, a decline of 45 percent. Ovens fell from 23.6 years to 14.3 years, a decline of 39 percent.

The causes are complex. Part of it is due to changing usage habits. In 1960, the average Norwegian family of four did two loads of laundry per week; in 2000, that number had risen to eight. Part of it is due to the materials. Vintage appliances from before the 1960s were made of solid brass, cast iron, and steel. Modern appliances use more plastic and aluminum. Part of it is due to increasing complexity. More electronic components mean more potential sources of error.

Interestingly, the study shows that not all appliances have shorter lifespans. Refrigerators, freezers, and dishwashers have largely maintained their lifespan. If planned obsolescence or a single factor were responsible, one would expect to see a general downward trend. The picture is more nuanced.

The pendulum and the limits of judgment

These examples reveal a pattern that goes beyond individual missteps. Technological and economic systems seem to move in waves. They cross boundaries, provoke resistance, and are then corrected by regulation, consumer boycotts, or simply market failure. BMW withdraws its seat heating subscriptions. The FTC sues John Deere. Canon explains to its customers how they can circumvent its own DRM. The pendulum swings.

Cory Doctorow coined the term “enshittification” for this phenomenon, which was voted word of the year in 2023. It describes a cycle in which platforms initially offer high-quality services to attract users, then lower the quality to serve business customers, and finally squeeze both sides to maximize short-term profits. Google Search is his prime example. The service became big thanks to relevant search results and minimal advertising. Today, it is littered with ads, manipulated by search engine optimization, and often less useful than it was a decade ago.

These observations lead to profound epistemic modesty. Who can judge in real time what constitutes genuine progress and what does not? The stagnation thesis developed in this paper is itself not immune to this question. Perhaps the stability of some paradigms is not a sign of lock-in, but a sign that an optimum has been reached. Perhaps the washing machine of 1990 really was better than that of 2020, and perhaps that of 2030 will be better again when the pendulum swings back.

What we can say with certainty is that progress is not a straight line. It is not even an upward curve. It is a complex, often chaotic phenomenon that moves in waves, crosses boundaries, is corrected, and takes new directions. The demand for “more innovation” without direction is just as naive as the assumption that everything that exists should be preserved.

The Center for Applied Complexity & Intelligence does not see itself as an advocate of progress at any price. It sees itself as a place where the patterns according to which complex systems function, stagnate, deteriorate, or improve are examined. The answer to the question of whether a particular change is progress can often only be given in retrospect. What we can do in the present is to understand the structures that enable or prevent change and, on this basis, make smarter decisions.

Real progress and its pitfalls

Despite all the criticism, it would be wrong to succumb to pessimism—the last few years have brought some remarkable breakthroughs. The mRNA vaccines developed during the COVID pandemic are based on decades of fundamental research and have opened up a whole new paradigm in vaccine development. Reusable rockets have dramatically reduced the cost of access to space.

Tyler Cowen, whose "Great Stagnation" thesis this paper takes up, has also nuanced his position in the meantime. He sees signs that we are "probably coming out of the great stagnation now," pointing to mRNA vaccines, generative AI, and advances in green energy. The question is not whether progress is possible. The question is whether we recognize the structures that distinguish real progress from self-reinforcing dead ends.

The trap of self-reinforcing loops

In practice, the warning signs described above combine to form a mechanism that is difficult to break: the self-reinforcing loop. The pattern is always the same. Once a paradigm gains enough momentum, it begins to create the conditions that cement its own dominance, regardless of whether it is the best approach.

The mechanism works in several stages. Initial successes attract investment. Investment finances infrastructure tailored to the dominant paradigm. This infrastructure makes it cheaper to work within the paradigm and more expensive to pursue alternatives. Training programs become specialized, career paths narrow, and regulations are adapted. Each of these steps increases the legitimacy of the existing paradigm and deprives alternatives of resources. In the end, the paradigm is dominant not because it is the best, but because it is the only one that anyone can still afford.

This pattern can be seen in AI research, where a single architecture absorbs virtually all resources, while alternatives languish due to lack of funding. It can be seen in energy systems, where the infrastructure of central power plants structurally disadvantages decentralized approaches. It can be seen in pharmaceutical research, where the approval system favors large molecules over new forms of therapy. It can be seen in construction, where standards and training are geared toward concrete and steel, even though other materials would be more environmentally friendly.

The danger does not lie in the dominant paradigm being bad. Perhaps it is even good. The danger lies in the fact that the system has no way of finding out. Alternatives are not refuted, but stifled. The question of whether the path taken is the right one is not answered; it is not even asked, because no one has the resources to pursue it seriously anymore.

Enabling the escape

The Center for Applied Complexity & Intelligence does not observe these dynamics from an academic distance. It sees itself as a catalyst for change, a place where self-reinforcing loops can be broken.

This also raises a fundamental question of attitude. We have become accustomed to viewing the systems we live in as given, as if they were laws of nature or fate. That's just how the market works. Technology is developing in this direction. Institutions have grown in this way. This attitude is convenient, but it is also wrong. The systems that surround us are man-made. They were designed by people, built by people, and maintained by people. And they can be changed by people.

The old saying goes: Everyone is the architect of their own fortune. That sounds like individualism, but it contains a deeper truth. We are not passengers on a ship whose course we cannot influence. We are the shipbuilders, the captains, the navigators. If the ship is heading in the wrong direction, we can change course. If the ship is poorly built, we can build a better one. The question is not whether change is possible. The question is whether we have the courage to try.

The Center's strategy is not to attack the existing system head-on. Such attacks bounce off established structures. The strategy is to create irritations, small disturbances that cause the system to vibrate. If a system is shaken long enough, cracks will appear. New possibilities can penetrate through these cracks.

Specifically, this means that the center connects people who think alike but work in isolation. There are researchers, engineers, and entrepreneurs everywhere who see the limitations of existing paradigms but have no community that shares their perspective. There are institutions and companies that have recognized that the status quo is unsustainable but do not know how to challenge it without jeopardizing their position. Finding, connecting, and supporting these actors is a central task of the Center.

It's not about providing ready-made answers. It's about creating the space in which better questions can be asked. Why does a neural network have to be structured the way it is? Why does an energy network have to be centrally controlled? Why does an organization have to be structured hierarchically? These questions may sound naive, but they are not. They are radical, in the original sense of the word: they go to the root of the problem.

A different perspective

The Center for Applied Complexity & Intelligence arose from an observation that crystallized over more than 20 years of practical experience in various fields. The observation was that technical systems always reflect the structures of the organization in which they are created. If communication within the company stagnates, it also stagnates in the system. If teams are fragmented, application landscapes become fragmented. If decision-making processes are complex, complex integration patterns emerge.

This observation was initially a practical tool. But over time, it became clear that the same patterns appear everywhere, in energy networks, in mobility concepts, in scientific paradigms, in social structures. The question became bigger: Why do the same problems repeat themselves in completely different areas? Why do the same patterns arise regardless of context?

The answer led to a change in perspective. The problems in different areas are not different. They are expressions of the same basic structure. They all follow mechanistic models that treat systems like machines, not like living networks. And they all encounter the same limitations.

Intelligence as an emergent phenomenon

The core idea developed by the center is simple to formulate, but far-reaching in its consequences. Intelligence is not an algorithm. It is not a program that can be written, a function that can be optimized, or a mechanism that can be constructed. Intelligence is an emergent phenomenon of complex systems.

What does this mean? It means that intelligence is a system property that arises from interaction, not from the properties of individual components. A neural network exhibits behaviors that no single neuron exhibits. An organization can solve problems that no single member could solve. A market aggregates information that no single participant possesses. In all these cases, something new arises at the system level that does not exist at the component level.

> Intelligence is the ability of a system to act contextually, adapt, and learn from interaction without this behavior having been explicitly programmed. It emerges from complexity, or it does not emerge at all.

This perspective has immediate practical consequences. When intelligence emerges, it cannot be constructed. One can only create the conditions under which it can arise. This requires a completely different approach than programming algorithms or optimizing parameters. It requires the design of architectures.

The limits of self-organization

This perspective should not be misunderstood. Emergence is not a value in itself. Self-organized systems can be resilient, but they can just as easily be inefficient, unfair, and unstable. Markets, social networks, and organizations regularly demonstrate that emergent behavior can just as easily lead to monopolies, coordination failures, or concentrations of power as it can to innovation and adaptability.

History is full of examples of emergent dynamics that nobody wanted. Financial crises arise from the interaction of rational individual decisions. Cities develop traffic jams even though every driver chooses the fastest route. Platforms that started out as open marketplaces turn into monopolies that exploit their own users. Emergence produces order, but it guarantees neither fairness nor efficiency nor stability.

The alternative to mechanistic thinking is therefore not to abandon control, but to adopt a different form of control. This control is context-sensitive rather than universal, adaptive rather than rigid, and deliberately limited rather than all-encompassing. It does not intervene in every interaction, but shapes the framework conditions under which interactions take place. It does not attempt to determine outcomes, but to channel dynamics.

The Center for Applied Complexity & Intelligence does not take a naive position that systems should be left to their own devices. The key question is not whether systems should be controlled, but how, where, and at what level. The answer to mechanistic overcontrol is not laissez-faire, but intelligent architecture.

Four principles

The Center’s work is based on four epistemic principles that serve as guidelines for research and practice. These principles have not been chosen arbitrarily. They arise from the insight that complex systems function differently from machines and therefore must also be understood differently.

PRINCIPLE 01

Patterns before mechanisms

The first principle states that we look for recurring structures before we formulate solutions. The focus is on forms of behavior, not on their implementation. Structures are stable, mechanisms are interchangeable. The same pattern can occur in completely different areas. Those who understand the patterns understand something deeper.

PRINCIPLE 02

Systems before components

The second principle states that an element is meaningless without its environment. We consider systems on several levels simultaneously: the micro level of agents, the meso level of subsystems, and the macro level of the overall organization. Intelligence arises from the relationships between these levels, not from the properties of a single level.

PRINCIPLE 03

Emergence before instruction

The third principle concerns system design: intelligent behavior cannot be imposed from above. Instead, conditions must be created in which it can emerge. Frameworks instead of regulations, structures instead of commands. Systems are empowered, not programmed.

PRINCIPLE 04

Context before content

The fourth principle states that meaning does not arise from data, but from its embedding. The same information has different effects in different contexts. The Center explicitly models context as a structuring factor that constitutes the meaning of information.

What we do

The Center for Applied Complexity & Intelligence is not a traditional research institute. It has no permanent team, no hierarchy, and no central laboratories. It is itself an example of what it researches: an emergent network of people, companies, associations, and initiatives that support each other and work together to overcome stagnant paradigms.

The work is divided into several areas. On a theoretical level, the Center investigates how complexity generates intelligence. It identifies universal patterns that occur in different areas, abstracts them, and formalizes them to make them transferable. On a practical level, the Center supports projects that validate these patterns in concrete applications. It connects like-minded people across disciplines and transfers knowledge between research and practice.

The network operates according to the principles it researches. There is no central control that determines what work is done. There is no fixed hierarchy that makes decisions. Coordination arises from the connections between the participants, just as in the systems studied by the Center. This is not an organizational weakness. It is a deliberate experiment: the question of whether a network organized according to the principles of emergent intelligence can actually be more intelligent than traditional structures.

What we do specifically

The abstract description as a “network” and “school of thought” raises a legitimate question: What does this mean in practice? Why would someone support the Center or collaborate with it? The answer lies not in a roadmap with milestones and deliverables, but in a series of intervention types that result from the approach described.

The first form of intervention is architectural analysis of stagnating systems. The Center examines specific areas where optimization has reached its limits, identifies the underlying architectural assumptions, and asks which alternatives are systematically not being pursued. Such analyses can be carried out for companies, public authorities, or research institutions that feel that their systems are not improving despite continuous investment.

The second form of intervention is the development of experimental parallel architectures. This does not mean replacing existing systems, but rather designing small, limited experiments that test alternative approaches under real conditions. The FLUID, AionCore, and Cognex projects are examples of such parallel architectures in the field of computing models and cognitive systems.

The third form of intervention is context modeling. Many systems fail not because of a lack of functionality, but because they do not understand context. The Center develops methods to explicitly model context instead of treating it as background noise. This applies to software as well as organizations.

The fourth form of intervention is networking among like-minded individuals. There are people everywhere who are working at the boundaries of existing paradigms but are isolated. Researchers at universities who are pursuing alternative AI architectures. Engineers in companies who know that their systems are reaching structural limits. Founders who are trying something that doesn't fit into any existing category. Finding, connecting, and supporting these people is a key task.

The Center does not promise ready-made solutions. It provides a framework for systematic reflection on problems that cannot be solved by conventional means.

How we recognize better architectures

The demand for "new architectures" remains abstract as long as it is unclear what distinguishes a better architecture from a worse one. The Center works with a set of orientation criteria that offer no guarantees but point the way forward.

The first criterion is reduced switching costs. Architecture that makes it difficult to switch is suspicious. Good architectures enable evolution; they do not force lifelong commitment. This does not mean that stability is unimportant, but stability should come from quality, not dependency.

The second criterion is decentralizability. Systems that only function when a central authority coordinates everything are fragile and susceptible to concentration of power. This does not mean that decentralization is always better, but an architecture that excludes decentralization in principle has a structural problem.

The third criterion is context sensitivity. Architectures that ignore context and apply the same solution everywhere regularly reach their limits. Good architectures adapt to their environment instead of adapting their environment to themselves.

The fourth criterion is learnability. Systems that remain static after completion become obsolete. Good architectures have built-in mechanisms for adaptation and evolution without requiring a complete rebuild for every change.

These criteria are not a checklist for guaranteed success. They are heuristics that help distinguish promising directions from likely dead ends.

Why this center?

A legitimate question is why the Center for Applied Complexity & Intelligence is needed when there are already established institutions dealing with complexity. The Santa Fe Institute has been researching complex systems for decades. IIASA works on systemic issues at the global level. Universities have chairs for systems theory and cybernetics.

The difference lies in the focus. Existing institutions primarily produce basic research and publications. They analyze complexity, but rarely intervene in stagnant systems. The Center does not see itself as an alternative to these institutions, but rather as a bridge between theory and practice. The question is not whether complexity science exists. The question is why it does not reach the systems it could transform.

For companies that feel their systems are not improving despite continuous investment, the center offers a perspective beyond the usual optimization logic. For researchers working at the boundaries of existing paradigms and seeking resonance, it offers a network of like-minded people. For institutions that know that the status quo is unsustainable but don't know how to challenge it, it offers a safe space for unconventional questions.

> All truth passes through three stages: First, it is ridiculed. Then it is fought against. And finally, it is accepted as self-evident.
>
> — Arthur Schopenhauer

If we are wrong

It is possible that this analysis is incorrect. The Center for Applied Complexity & Intelligence does not claim infallibility, and it would be intellectually dishonest to present its own position as unassailable.

Perhaps some of the paradigms described here as stagnant are actually close to optimal. Perhaps Von Neumann architecture is dominant not because the switching costs are too high, but because it is simply the best solution for most applications. Perhaps centralized architectures are unavoidable in many areas because coordination without hierarchy does not scale. Perhaps the limitations of transformers are not architectural dead ends, but temporary hurdles that can be overcome with more data, more computing power, and clever engineering solutions.

The Center does not view these possibilities as a threat, but rather as a necessary part of serious research. Our work is not aimed at proving ourselves right. It aims to reveal systematic blind spots that the current consensus overlooks or suppresses. If existing paradigms prove to be sustainable in the long term, then this sustainability should be the result of critical examination, not tacit acceptance.

What we can say with certainty is that the questions we are asking should be asked. Whether the answers confirm what we suspect or prove us wrong remains to be seen. Both outcomes would be a gain in knowledge. The worst thing would be not to ask the questions in the first place.

The real future

The image at the beginning of this paper shows a future that never came to pass. Flying cars, colonies on Mars, a world full of technological wonders. One could interpret this as a failure, as the collapse of overly ambitious promises. But perhaps that is not the right frame of reference.

The visions of the 20th century were extrapolations of what already existed. Flying cars are optimized cars. Mars colonies are optimized colonies. Bigger rockets, faster planes, more powerful computers—all of these things fall within existing paradigms. The idea was that the future would be like the present, only more of it. Faster, bigger, further.

Perhaps that was the wrong idea. Perhaps the real future lies not in optimized versions of what already exists, but in fundamentally different architectures. Not in cars that fly, but in mobility systems that organize themselves. Not in power plants that are bigger, but in energy grids that balance themselves decentrally. Not in AI models that have more parameters, but in systems that actually learn.

The Center for Applied Complexity & Intelligence is working on this different future. Not by designing utopian visions, but by understanding the principles that enable such systems. The work is slow because real paradigm shifts are slow. It requires patience because new architectures do not emerge overnight. But it is necessary because optimization within existing paradigms has reached its limits.

We work where progress has stood still for decades
and develop new principles for real breakthroughs.

complexity-intelligence.org · contact@complexity-intelligence.org

Sources and further reading

The arguments and data presented in this paper are based on a range of sources from academia, business, and journalism. The following compilation provides references for further reading.

On the theory of technological stagnation

Tyler Cowen developed the thesis of the "Great Stagnation" in his book of the same name from 2011. The work argues that since around 1973, the American economy has largely harvested the easily reachable fruits of technological progress. The book is available through the Mercatus Center at George Mason University.

Peter Thiel formulated his critique of technological development in lectures at the Yale School of Management and in the manifesto "What Happened to the Future?" by the Founders Fund. The often-quoted statement about flying cars and 140 characters comes from this context.

The American Enterprise Institute has addressed the debate on technological stagnation in several articles, including "Social and Physical Theories of Technological Stagnation" and interviews with Tyler Cowen about the current state of the discussion.

On energy networks and infrastructure

The International Energy Agency regularly publishes reports on global energy infrastructure. The report "Electricity Grids and Secure Energy Transitions" contains detailed analyses of the age and condition of global grid infrastructure as well as the necessary investments.

The Council on Foreign Relations offers an accessible introduction to the structure and history of the American power grid with the Backgrounder "How Does the U.S. Power Grid Work?"

The World Economic Forum publishes the annual report "Fostering Effective Energy Transition," which analyzes the state of the global energy transition and the role of grid infrastructure.

On stagnation in aviation

Simple Flying and other aviation publications have published several articles on the question of why commercial aircraft today fly slower than in the 1960s. The MIT Department of Aeronautics and Astronautics has provided technical explanations of the relationships between speed, drag, and fuel consumption.

On productivity in the construction industry

The National Bureau of Economic Research has published several working papers on the stagnation of construction productivity, including "The Strange and Awful Path of Productivity in the U.S. Construction Sector." The Harvard Business School has presented a current analysis with "Why Has Construction Productivity Stagnated?" (Working Paper 25-027).

The Federal Reserve Bank of Richmond offers an economic-historical perspective with "Five Decades of Decline: U.S. Construction Sector Productivity" (Economic Brief 25-31).

On the antibiotic crisis

The Wellcome Trust Foundation has published extensive materials on the question of why the development of new antibiotics has become so difficult. The ReAct Group, an international network for combating antibiotic resistance, also offers detailed analyses.

The American Medical Association published an article on the social history of the antibiotic era in the Journal of Ethics, which provides historical context for the "golden age" of antibiotic discovery.

On nuclear energy

The Wikipedia articles on Generation II, III, and IV reactors provide an overview of the technological development of the nuclear industry. The Generation IV International Forum publishes information on future reactor concepts.

An academic analysis of stagnation in the nuclear industry can be found in the article "Destined for decline? Examining nuclear energy from a technological innovation systems perspective" in the journal Science Direct.

On stagnation in the education system

The blog Education Rickshaw published a pointed critique of the persistence of traditional teaching methods with "After 100 Years of the Same Teaching Model It's Time to Throw Out the Playbook."

Kappan Online offers an analysis of structural inertia in the education system with "Educating students for an outdated world."

Regression in the guise of progress

Coverage of BMW's subscription model for seat heating is extensively documented. Edmunds, The Drive, and Carbuzz have thoroughly analyzed the controversy and BMW's subsequent reversal. The Consumer Rights Wiki maintains a summary of events.

Issues with Tesla software updates have been documented by The Drive and the Tesla Motors Club. Reports of "bricked" Cybertrucks date from December 2024 and were picked up by Top Speed and other automotive publications.

The Federal Trade Commission lawsuit against John Deere was filed in January 2025. The FTC has published a press release summarizing the allegations. The U.S. Public Interest Research Group maintains a chronological documentation of the "Right to Repair" conflict with John Deere since 2015. NPR and NBC News have reported on the lawsuit and its background.

On the printer industry and DRM, Techdirt, Vice, and The Register have reported extensively. The story of Canon's DRM bypass instructions during the 2022 chip shortage was documented by Gizmodo and others. The EU's Cool Products Initiative has published a report on the environmental impacts of DRM practices in the printer industry.

The Norwegian long-term study on household appliance lifespans was presented by Norwegian SciTech News in 2025. The original appeared in an academic publication and distinguishes between different types of appliances, nuancing the undifferentiated thesis of "planned obsolescence."

Cory Doctorow's concept of "Enshittification" was developed in his book of the same name, published by Verso Books in 2025. The American Dialect Society chose the term as word of the year 2023. Washington Monthly, Current Affairs, and other publications have extensively discussed Doctorow's theses.

Theoretical foundations

Donella Meadows' "Thinking in Systems" offers an introduction to systems thinking and the concept of leverage points. Stafford Beer developed the idea of the organization as a cybernetic system in "Brain of the Firm." Nassim Taleb's "Antifragile" analyzes systems that become stronger under stress. W. Ross Ashby's "Introduction to Cybernetics" lays the theoretical foundations for understanding variety, regulation, and adaptivity.
