

# Democracy and the Digital Drive of Society





# **Democracy and the Digital Drive of Society**

# Democracy and the Digital Drive of Society

**This essay collection was produced by the Algorithms,  
Data and Democracy project - April 2026**

Alf Rehn  
Christina Lioma  
David Budtz Pedersen  
Helene Ratner  
Leonard Seabrooke  
Lisbeth Knudsen  
Noya Kohavi  
Sine Nørholm Just  
Torben Elgaard Jensen

## **Design**

Michael Hernvig

## **Translation**

Citadel Translations

## **Editor**

Nini Katrine Bertelsen

## **Print:**

Printconnect Blok Aps

## **ISBN:**

97887-90275-22-8

*The ADD project is funded by*

**VELUX FONDEN** ✕ **VILLUM FONDEN**

# Table of contents

|   |                |
|---|----------------|
| <b>1. Preface</b>   | Page <b>6</b>  |
| <b>2. Democracy loses its meaning when machines sing only of money</b><br>– <i>Sine Nørholm Just</i>  | Page <b>8</b>  |
| <b>3. The fog is lifting, but we still have work to do: Danish AI lessons 2020–2025</b><br>– <i>Torben Elgaard Jensen &amp; Helene Ratner</i> | Page <b>13</b> |
| <b>4. The technical risks of AI are as problematic as its societal impacts</b><br>– <i>Christina Lioma</i>                                    | Page <b>20</b> |
| <b>5. Exit, voice, and fealty in AI market-making</b><br>– <i>Leonard Seabrooke og Noya Kohavi</i>  | Page <b>25</b> |
| <b>6. On the landscapes of data intensities</b><br>– <i>Alf Rehn</i>  | Page <b>30</b> |
| <b>7. The role of social scientific advice in governing artificial intelligence</b><br>– <i>David Budtz Pedersen</i>                          | Page <b>35</b> |
| <b>8. Algorithms have rewritten the constitution</b><br>– <i>Lisbeth Knudsen</i>  | Page <b>41</b> |

# Preface

**F**ive years into *Algorithms, Data and Democracy*, it is hard to sustain cosy illusions about artificial intelligence. Take the idea that it's just a tool, for example. Or that all digitalisation represents progress per se. Or that democracy will always catch up, so tech companies might as well have free rein to develop what they want, as fast as they want. This anthology reflects the demise of such illusions. That is its strength. And why it is needed.

In the last five years, AI has evolved from a specialised R&D field to a potent social force. It already shapes how we communicate, work, govern, seek knowledge, make decisions, and imagine the future. It has made headway in public administration, the market, the media, and culture, and is increasingly encroaching on the intimate realms where people form meaning and identity, exercise judgement and grapple with doubt. AI is no longer just about innovation. It's about power.

The anthology reflects that reality. It's not just about technology. It's about the societies that technology is reshaping. It's about who sets the direction, who reaps the benefits, who bears the costs, and which institutions succumb, which waver, and which stand their ground. One particular realisation runs like a red thread through all of the essays: algorithms, data and AI aren't external forces crashing into democracy like an extreme weather event. They are woven into the fabric of our political economy, our institutions, our infrastructure – indeed, into our thinking about the very nature of society.

As realisations go, this is one that complicates the debate. But also makes it more honest.

Too often, public debate about AI has been dominated by two equally inadequate positions. On the one hand, technological fervour masquerading as realism: a narrative concerned solely with efficiency, acceleration, scaling and smart solutions, as if all our social and political problems had just been waiting for data to solve them. On the other hand, a form of moral panic: technology cast as a more or less separate being, a threat to humanity, truth, even civilisation itself. Both positions overlook something crucial. Namely, that artificial intelligence is neither neutral nor autonomous. It is designed, trained, funded, distributed, and implemented in specific institutional and economic contexts. It doesn't operate in a vacuum. It operates in society.

These essays insist on something that ought to be self-evident, but is all too often drowned by all the noise: AI is always a democratic issue. Not just when things go wrong. Not just when a chatbot hallucinates, a platform distorts public discourse, or a public-sector body runs afoul of the law, rights, or ethics. The ramifications for democracy are built in from the very outset: in basic choices concerning purpose, design, data, ownership, access, regulation, and accountability. Who benefits should be our first question, not an afterthought.

---

This anthology consists of essays by the lead researchers on a truly unique 10-year interdisciplinary project – *Algorithms, Data and Democracy* – funded by the Villum Foundation and the Velux Foundation. The Principal Investigator (PI) is Professor **Sine Nørholm Just**, the co-PIs Professor **Torben Elgaard Jensen**, Professor **Christina Lioma**, Professor **Helene Friis Ratner**, Professor **Leonard Seabrooke**, Professor **Alf Rehn**, Knowledge Broker Professor **David Budtz Pedersen**, and the Head of Outreach and Communication, Strategy Director **Lisbeth Knudsen**.

---

It is also clear now that the most interesting lessons drawn from AI bursting on the scene aren't just about what the technology can do. They are just as much about the obstacles it encounters. Law. Infrastructure. Data quality. Professional norms. Civil liberties. Organisational resistance. Public legitimacy. For decades, factors like these were depicted as a drag on innovation. In fact, they turn out to be something else, something far more significant: democratic friction. Not flaws in the system, but signs that society has not yet completely abandoned the notion of holding to account those in power just because power now speaks in code, models and probabilities.

It's a crucial lesson. It reminds us that a society's resilience isn't measured solely by its ability to adopt new technology, but also by its ability to impose limits. Optimisation isn't the answer to everything. Not all decisions should be automated. Not all communication should be industrialised. Not everything that can be scaled should be allowed to define the terms of our collective life.

Several of the essays point to something even more unsettling: that the digital technologies we often refer to as communication tools are increasingly reshaping the very foundations on which meaning, attention, and the public sphere are built. When data becomes a raw material, when behaviour becomes a business, and when communication becomes a resource in an economy that rewards intensity over insight, it's not just the market that changes, but the cultural infrastructure of democracy. At this stage, the question is no longer whether people are well-informed. It is whether a collective public sphere exists at all, one in which a shared reality can emerge and be sustained.

In that context, AI is particularly challenging. Not because machines suddenly think like humans. They don't. But because they are increasingly capable of mimicking human communication so convincingly, so quickly, and so cheaply that it becomes harder to draw a line between meaning and simulation, between exercising judgement and regurgitating information, between participation and manipulation. This new situation not only reshapes working practices and the way knowledge is produced. It also weakens the symbolic and social processes that underpin democracy.

In response, the essays in this anthology aren't just diagnostic. They're didactic. They point out that we must not leave AI up to market forces, the self-aggrandising mythology of the tech industry, or the state's incessant compulsion to make systems more efficient. If AI is to serve democracy, it must be bound by the democratic will of the people. Making that happen will take more than just regulation in the narrow legal sense. It will take institutional inventiveness, public debate, new forms of ownership, different types of incentives, and a far more grown-up conversation about what we want to protect, what we want to develop, and what we will not tolerate.

The point raised most regularly in these essays is perhaps that we must not let the future be determined by technology alone, but by how societies choose to integrate it into their norms. There is nothing preordained about AI. Nor is it some kind of entirely innocent entity. It's a field in which the power struggles of our decade are starkly visible: between the public sphere and private platforms, between rights and efficiency, between people and profit, between human judgement and machine calculation.

The Algorithms, Data and Democracy project may have reached the halfway point on its ten-year journey, but that doesn't mean it's time to pause for a rest. It's time to take stock and avoid the complacency trap. We know more than we did five years ago. We have seen more, tried more, understood more – and made more mistakes. In some areas, the fog has lifted. But our work is far from finished. On the contrary. The more AI is built into the basic infrastructure of our society, the more important it becomes to ask the big questions. And to ask them early on, clearly and without undue deference to technology.

The question isn't whether AI should be allowed to play a role in society – that ship has long since sailed – but what kind of society we want AI to help us build.



**Lisbeth Knudsen**

Lisbeth Knudsen, Head of Outreach, ADD

## Essay 1

# Democracy loses its meaning when machines sing only of money



By

**Sine N. Just**

Professor of Strategic Communication,  
Roskilde University and Principal  
Investigator of the Algorithms, Data  
and Democracy Project



The machine appeared in the distance,  
singing to itself of money

RS Thomas, 1972

**W**ith the advent of generative AI, the machine's song of money has moved close to the most intimate aspects of human life. With it, the logics of machinic production have saturated all dimensions of not only political and economic life, but cultural and emotional existence as well. Throughout the era of digitalization, the entirety of human experience has become ripe for monetization, and the making of money has become attached to activities that were not traditionally considered as labour. As data can be extracted from our interactions with social networks, indulgences in cultural consumption, searches for online information, every scroll we take and every click we make has, as Soshana Zuboff (2019) puts it, become the raw material

of surveillance capitalism. Refined data is, among other things, used to tune the algorithms of digital platforms for maximum engagement, to target advertisements for maximum effect, and to train AI technologies for maximum agential capacity.

As digital technologies have shifted from recommendation and moderation to production of content, they have become ever more deeply inserted in people's social and personal spheres. No longer limited to extracting and processing data on human behaviour, but also producing texts, images, sounds, and other symbolic actions that have hitherto appeared to be uniquely human. With machine communication, no aspect of human life can claim distance from the money machine. This is not to say that people did not previously communicate for profit, but prior to the advent of generative AI it was possible to distinguish between economic and social meanings of communication. Now all meaning formation can be monetized to the detriment of human individuality as well as democratic societies, which become legitimate through collective meaning formation.

## Communicative capitalism

The shift from human meaning formation that has the potential to serve democracy towards the all-consuming capitalization of communication began well before the public introduction of generative AI. In fact, the entire process of digitalization can be understood as the rise of communicative capitalism where communication, as Jodi Dean argues, becomes a key resource for profit accumulation. To Dean, this implies that communication becomes detached from its points of reference, that it no longer means anything beyond its own circulation. Or, as she puts it:

*Circulation has eclipsed meaning. That something is shared online does not depend on what it means. It depends on its affective capacity: does the shared item manifest outrage: is it funny and diverting? We attend less to the meaning of an utterance than to its affective dimension, which is most powerful when it contains different, conflicting meanings (Dean, 2019: 331– 332).*

Communicative capitalism, in Dean's analysis, is driven by the intensification of feeling from which, she argues, money can be derived directly.

I agree with the analysis of how digital flows of feeling have become intimately connected with their monetization, but I disagree that this process is necessarily meaningless (Just, 2024). Dean sees meaning formation and the intensification of feelings as separate processes and argues that under communicative capitalism, democratic life is flooded with feeling, "drowned out by outrage and puppies" (Dean, 2019: 22). However, I believe that intensification of feeling may be productive of economic as well as social value, that it makes money and meaning. To support this view, we must think of meaning and feeling as one; not separate forms of communication but interdependent prerequisites of it. Communication without feeling is meaningless.

While rationalization is sometimes seen as the measure of democracy, the relationship is inverse; the more rationality is detached from feeling, the more it loses its democratic potential. Conversely, we can derive democratic and cultural value from processes of digital communication because of their intensities of feeling; even as these processes are increasingly monetized, they need not be fully captured by mindless profit maximization. As generative AI takes on more and more communicative functions, however, that option is foreclosed: AI communicates without feeling and, hence, without meaning.

## Communicating machines

Generative AI has no intentions, it reasons with no motive. In technological terms, this works by stringing symbols together (whether words, pixels, or code) one at a time, always adding the most probable next element in a sequence of probabilities, parroting patterns that the AI system has picked up from its training data (Bender et al., 2021). Given the advanced state of AI technologies' ability to produce apparently meaningful communication from probabilistic inferences, we might ponder the



The problem is not that people can make AI generated communication appear meaningful, but that in so doing we gradually weaken our ability to come together as a society, each of us making individual sense of our communication with the machine; making meaning for one rather than in mutual relationships with others.

**Sine N. Just**, Professor of Strategic Communication, Roskilde University and Principal Investigator of the Algorithms, Data and Democracy Project

very meaning of meaning. More pressingly, however, we should ask what happens when this ability is harnessed to the economic advantage of AI companies; what happens when the mechanical parrot only sings of money?

Combining the issues of how generative AI makes meaning and money in a thought-provoking experiment, Gary Greenberg (2025) reverses the now all too common practice of using AI for therapy and puts ChatGPT on the couch. The result is a riveting exposé of what AI ‘knows’ about itself, where the scare quotes are to remind us that AI does not have self-knowledge in any meaningful sense. What we learn, when asking AI to explain itself, is not what AI is and wants, but what it was made for.

In a central passage, Greenberg and his ‘client’ discuss the very relationship between the technology and its makers, with the chatbot listing three key principles or wishes on the part of AI developers:

*They wanted to make something that humans would not reject. [...] They wanted to avoid being blamed. [...] And [...] they wanted to make a machine that would love us back, without needing love in return (ChatGPT, quoted in Greenberg, 2025, emphases in original).*

In this instance, I believe, the probabilities have converged to produce an accurate insight: The reason AI communication appears meaningful to humans is that underlying models are optimized for such appearance in the effort to avoid that humans reject technologies and that technologies reject humans. And to ensure that the technologies’ ‘parent’ companies can avoid blame when things go wrong while profiting from everything that goes right. This machine really does sing to itself of nothing but money.

## **The problem with machine communication**

Beyond the usual complaints about unfettered capitalism, you might wonder, what is the problem with automated production of meaning? First and foremost, that it is not, in fact, meaningful. I have already indicated why a lack of meaning is democratically problematic; whereas economic institutions can thrive by making money from nothing, political institutions need heartier sustenance – they need people to believe in them and to dream with them. In this light, the lack of intentionality in AI communication is not the main concern; rather, the problem is that people tend to experience communicative interactions with AI as if they were meaningful.

On the one hand, this tendency for people to make sense of the senseless is nothing new. Instead, it can be understood as an inherent human inclination to, for instance, see ‘a man in the moon’ or to explain natural disasters as ‘the wrath of the gods’. As such, people are primed for one-sided meaning formation, constantly adding agency and intentionality where, in fact, there is no meaning beyond the phenomena themselves.

On the other hand, the tendency towards anthropomorphic interpretation makes us extremely susceptible to manipulation; when something appears to make sense, we are more than ready to grant that thing sensibility. And generative AI has unprecedented abilities to act like it knows what it is doing, tempting people to think of it in human terms.

If we understand the current situation along the lines of a sci-fi narrative in which machines have become like people and are either maliciously taking control of the world or benevolently helping us save it, we enhance the problems that arise from people’s need to make sense of the senseless. Hence, we should not frame the situation in terms of the risk of machine autonomy. And we should not think of it as only deriving from people’s willingness to be deceived either.

Rather, the main risk lies at the intersection of these two issues; through recurrent use of talking machines, people may, increasingly, cede control to these machines – and, in the process, become more like machines themselves. Thus, people become susceptible to automated persuasion – to reasoning that has no motive, no capacity for emotional attachment. Beyond the individual suffering this may incur, the loss of collective meaning is democratically detrimental. As people turn in on themselves in closed circuits of solipsistic interpretation of what generative AI feeds them, our ability to make meaning collectively is foreclosed, drying out the sources of democratic legitimacy.

The problem, then, is that whereas democratic society is built on people's interdependence, shaped to support our individual and collective need to be recognized as members of society, generative AI has no such needs. While the technology can appear to have the ability to reciprocate human emotion, i.e., to grant recognition even as it does not itself need to be recognized, this is just an appearance – humans seeing ghosts in the machine in the same way as they see a face in the moon. The problem is not that people can make AI generated communication appear meaningful, but that in so doing we gradually weaken our ability to come together as a society, each of us making individual sense of our communication with the machine; making meaning for one rather than in mutual relationships with others. Nothing comes from such human-machine relationships – except the meaning individuals may derive for themselves and the money the machine can generate for its master. For society, each of these outcomes is, at best, pointless.

## Can human-machine relations become meaningful?

When humans think of machine communication as meaningful, the social value of communication is reduced, but this is not to say that we cannot use AI technologies in democratically useful ways. Tools have always been integral to human society and shaped social formations accordingly; where AI differs is that it actively invites us to think of it as a partner rather than a tool. This is a feature of the technology that could be changed, as AI technologies could be trained for specific purposes and optimized to provide matter-of-fact solutions rather than marketed as general purpose technologies with more a less chatty personalities. AI could offer support for human meaning formation rather than take over meaning-making functions. This development, however, requires that we teach the machine other songs than that of money. Meaning, AI tools that can support democracy will need to be democratically developed and owned rather than trained to maximize for profit.

Democratizing AI can take many forms. It could be large, state-backed initiatives like the Swiss language model Apertus, which is an open-source solution enabled entirely by public funding. Or it could be smaller and more radically participatory experiments with cooperative platforms and democratic AI like those being developed and tested in the 'laboratory' of Nordic Media Lab. If we turn to exploring democratic means of AI development, we might also deploy the technologies for democratic ends. We might even be able to use AI to muffle the machine's singular song of money and augment the multifarious chorus of people singing society together. AI does not in itself have a purpose, it does not reason with motive; when beginning from that premise, we may find uses for the new tool that are democratically meaningful.

---

### References

- Bender, E. M. et al. (2021): [On the dangers of stochastic parrots: Can language models be too big?](#) *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*.
- Dean, J. (2019): [Communicative capitalism and the revolutionary form](#). *Millennium*, 47(3): 326-340.
- Greenberg, G. (2025): [Putting ChatGPT on the couch](#). *The New Yorker*.
- Just, S. N. (2024): [Controversial Encounters in the Age of Algorithms](#). *Bristol University Press*.
- Thomas, RS (1972): [H'm](#). *Macmillan*.
- Zuboff, S. (2019): [The Age of Surveillance Capitalism](#). *PublicAffairs*.

## Essay 2

# The fog is lifting, but we still have work to do: Danish AI lessons 2020–2025



By

**Torben Elgaard Jensen & Helene Ratner**

Respectively Professor of Techno-Anthropology, Aalborg University and Professor, Technical University of Denmark

**F**uture historians writing about artificial intelligence in Denmark will probably look back on the period 2020–2025 with a degree of bewilderment. Although they are likely to acknowledge the determined efforts made to figure out the nature of AI and its uses, they may also shake their heads slightly and lament that the Danes sometimes looked ‘completely clueless’.

This essay is an attempt to write history before it happens. In it, we catalogue at least some of the things the Danes worked so hard on in an attempt to learn about AI in the first half of the 2020s. Being part of Algorithms, Data & Democracy (the ADD project) has afforded us the opportunity to follow the AI debate, politics and practice in Denmark closely since 2021. For example, we have followed local

councils' attempts to introduce AI into public administration, mapped the AI discourse in the media and academic literature, studied how AI has developed over time, and looked at broader trends in Danish digitalisation policy. Based on our research and countless productive conversations with people in and around the ADD project, the essay outlines six lessons learned about AI in Denmark.

## #1 Artificial intelligence is not 'a thing'

Artificial intelligence has a long and complicated history stretching back to the 1950s, with several highs and lows along the way. Not until OpenAI launched ChatGPT in November 2022 did the perception spread that it had finally *arrived*. AI was perceived as a 'thing' that had suddenly entered world history, much like a child being born or a previously unknown element being discovered. The idea was reinforced by OpenAI marketing their product as a turning point in world history, and by lurid media coverage, such as the story of the Google engineer who reported with horror that his machine had suddenly come to life and gained consciousness (Tiko, 2022).

Of course, not everyone viewed it that way. Our colleague in the ADD project, Christina Lioma, a professor of machine learning, drew our attention to the long line of other large language models (LLMs) that had preceded ChatGPT, noting dryly that perhaps the most spectacularly new thing about OpenAI was the business model rather than the technology.

Some social scientists launched full frontal attacks on the idea of AI as a 'thing' at all. Suchman (2023) not only argued that the idea of AI as 'a thing' was misleading but also problematised the political consequences of such an idea. Making AI sound like a single, finished entity distracts attention away from the people and organisations that design, train, implement and use these systems. It casts AI as an actor, and those responsible for it, power and vested interests fade into the background. Insisting AI is a multitude of actors, including companies, infrastructure, models, data, organisational choices and usage scenarios, allows us to re-politicise the debate.

It took time for the wider public to let go of the idea of AI as a single, specific thing. Talk of in-built bias in the algorithms played an important role in this shift. A series of revelations made it clear that we need to understand the phenomenon in terms of *AI models*, and that the way these models work depends on the *data* used to train them – data that can be misleading and may have been illegally acquired by tech companies. Talk of 'hallucinations' – the tech companies' 'newspeak' for what ordinary product developers would otherwise call errors – also played a significant role. It became increasingly evident that the AI models suffer from a high degree of instability and frequently produce incorrect results. This realisation moved the public discourse on from viewing AI as a mysterious and captivating new phenomenon to focusing on specific *models*, as witnessed by tech journalists' coverage and reviews. Which model is the fastest, cheapest, or best at certain tasks? Has Google Gemini 2.5 overtaken ChatGPT 4 for images and video? Is Claude 4 better at coding than Grok 4? Ad infinitum. AI had become a product genre rather than 'a thing'.

## #2 AI is about socio-technical systems, not about humans vs. machines

As mentioned, the early 2020s saw growing interest in different AI models and versions, each with its own features and flaws, and the number of stories about practical applications and the contexts in which AI was used increased.

One horror scenario discussed in the late 2010s was the military use of AI to develop autonomous weapons systems. Project Maven in the USA, which used machine learning to analyse drone videos and identify potential targets, triggered a major public controversy. The army claimed the system 'just' helped it make decisions and said there would always be 'a human in the loop'. Critics viewed it as a dangerous step towards 'killer robots'.

The ensuing controversy led Google to withdraw from the project and Palantir to step in. Regardless of which side of the debate you trusted, it became clear that multiple devils lurk in the details when it comes to the practical deployment of AI. What would the military actually use it for? How many minutes or seconds would human operators have to approve the targets suggested by the sy-

stem? How high is the quality of the data on which the system is based? What mechanisms are in place to hold the military and tech companies accountable for mistakes?

The number of practical examples of AI exploded after the launch of ChatGPT. Unregulated use of new tools, ‘assistants’ or forms of cheating spread through education systems and workplaces everywhere. Evidence mounted that the practical aspects of introducing AI – the nature of the organisations, the circumstances, intentions, transparency, etc. – are crucial to the outcomes and whether we find them acceptable.

It became clear that the debates about human versus machine, or ‘autonomous system’ versus ‘full human control’, were drawing to a close. The number and variety of practical applications helped clarify that the choice was not between human or machine, but between different socio-technical systems, a perspective formalised at the highest level by the EU’s AI Act in 2024. The legislation deliberately refrained from identifying specific technologies to be regulated, focusing instead on which *intended purposes* (i.e., specific socio-technical systems) should be permitted, regulated, or prohibited.

### **#3 AI projects are always complex and face significant obstacles**

The optimism was palpable in 2020 when the government, Local Government Denmark (KL) and Danish Regions invested DKK 187 million in AI ‘signature projects’. The idea was for the public sector to put serious effort into developing artificial intelligence. Back then, AI was understood as supervised machine learning – predictive models to calculate the probability of a given individual completing a rehabilitation programme, for example, or being at risk of long-term unemployment. This type of AI is different from LLMs, ChatGPT being the most fabled, because they are trained on the public sector’s own data. Local and regional authorities set up 40 signature projects on subjects like employment, health, administration and the climate. In all 40, the road from idea to development and implementation proved far rockier than expected.

Firstly, the people running the projects failed to conduct adequate research into the potential legal hurdles. Navigating Danish legislation proved quite complex, especially for projects involving the use of personal data to develop predictive models for use in real-world case work. Many of the projects had to change course. Even more had to be abandoned because the councils lacked the legal authority to implement them. The first database of public-sector AI projects in Denmark confirmed the trend, documenting that legality is a recurring challenge ([www.offentlig-ai.dk](http://www.offentlig-ai.dk)). A little over a year after a disappointing evaluation, the government proposed amending the legislation to afford public-sector bodies greater legal scope to use AI systems to process personal data (Folketinget 2026). The proposal demonstrated that the relationship between AI and the law is dynamic – not only does the law regulate AI, but the desire to use AI can also lead to changes in the law.

Secondly, many of the projects encountered challenges related to IT infrastructure and the quality of the data. As we know, AI is only ever as good as the data used to train it. Many of the signature projects found the data was incomplete or of poor quality. Another hurdle was that various IT systems needed to exchange data in real time – a task that proved far more complicated than expected in local government bodies, which often use older IT systems that do not communicate with each other.

Thirdly, the signature projects proved that introducing AI invariably raises complex ethical and political issues. The more a project encroaches on people’s lives, the more it attracts attention and criticism. It became clear that AI is not just a ‘solution’ but a technology that requires a careful balancing act between ethical dilemmas and democratic values. It is telling that every single signature project involving vulnerable children and young people was ultimately shut down (Ratner & Schrøder, 2024).

Overall, the projects illustrated that AI systems cannot be ‘rolled out’. Technological solutions need to be built up slowly – and they often fail. The multiple challenges – legality, data, infrastructure, ethics, and legitimacy – showed that AI has profound repercussions for how the public sector functions. The signature projects during 2020–2025 showed that AI projects fail because they collide



Overall, 2020–2025 was a period when some of the AI fog lifted. However, for the most part, the work of setting frameworks for, rolling out and living with AI still lay ahead. Big, new, and complicated questions also arose about unintended environmental and geopolitical consequences.

**Torben Elgaard Jensen & Helene Ratner**, respectively Professor of Techno-Anthropology, Aalborg University and Professor, Technical University of Denmark

with existing institutions, infrastructure, and rights. Artificial intelligence turned out not to be a fast track to efficiency. However, it presented an opportunity to renegotiate how the public sector should use data, power, and automation in a democratic society (Laage-Thomsen & Ratner 2025).

#### **#4 The appetite for efficient, practical and ‘boring’ AI is growing, and for good reason**

In the wake of the controversial and legally tortuous signature projects, the politicians changed tack. In 2024, the Minister for Digital Affairs, Caroline Stage (Moderate), articulated the new approach clearly when she proclaimed she would promote ‘unsexy’ and ‘ultra-boring’ AI ‘that works’ (Stage in Lehmann, 2024). This was in stark contrast to her predecessor, Marie Bjerre (Liberal). Just six months earlier, Bjerre had urged Danes not to ‘shy away from’ AI, even suggesting that it might be used by local authorities when taking children into care (Bjerre in Bang-Udesen 2024) – probably the most intrusive and high-risk use of AI imaginable.

The shift towards ‘boring AI’ reflected a growing recognition of the complexity of its use in real-life case work involving actual human beings, and the difficulties public bodies faced when trying to develop AI solutions in-house. The new emphasis was on administrative and organisational functions, such as automating route planning and processing building applications and FOI requests, which were among the few signature projects to survive the trial phase.

Meanwhile, ChatGPT had burst onto the scene. Overnight, it transformed political perceptions of what AI was and how it could be used. Large language models made it possible to envisage AI as an assistant that could write up the minutes of meetings, act as a chatbot to provide advice to council staff and service users, or help draft letters and rulings. To both the councils and the public, it all sounded less invasive and easier to legitimise.

Talk of ‘unsexy AI’ was also a political strategy, of course, aimed at normalising its use in the public sector. By describing AI as boring, mundane, and technically uncontroversial, politicians were trying to make it sound less intimidating and more acceptable. However, even the unsexy can be complicated. The same fundamental challenges – in terms of the law, data protection, and IT infrastructure – still applied. They just took new forms.

Cloud infrastructure proved to be a particularly thorny issue. It was unusual for LLMs to run on local IT infrastructure. They needed access to commercial cloud services, often owned by American tech giants. Familiar but unresolved questions about digital sovereignty and GDPR compliance again came to the fore. ‘Boring’ AI was by no means straightforward, in either a legal or technological sense. By the mid-2020s, it had become apparent that no matter how unsexy the AI, questions remained regarding power and the structural barriers to digitalisation in the public sector.

#### **#5 Still no successful scaling**

The start of the decade was a boom time for companies directly involved in the technical side of developing AI. When OpenAI launched ChatGPT, it reached 100 million users in less than two months, much faster than any previous app. Investments in AI also shattered all records, even surpassing investments in historic ventures like the Manhattan Project and the Apollo programme.

Despite the spectacular influx of money and the rapid spread of AI, there was also a growing sense that the scaling of meaningful value-generating applications was proving somewhat sluggish. More and more people were talking about AI as a financial bubble. Indeed, 80% of private-sector AI projects in America failed – twice the failure rate of non-AI tech projects (Ryseff et al. 2024). There seemed to be a lack of clarity about what AI was actually meant to do for businesses and their customers, a problem that exacerbated other issues – inadequate data, poor infrastructure, and a tendency to use the most advanced technology rather than the most relevant.

Denmark launched the local government signature projects in 2020. Four years later, the Digital Taskforce for AI was set up to learn from experiences and select large-scale projects capable of being

rolled out across the public sector. The Taskforce announced highly ambitious goals (including freeing up 30,000 full-time equivalents by 2035), and has selected three projects to date – speech-to-text technology (automated documentation), digital assistants to provide help and advice to council staff, and digital assistants to provide help and advice to local people and companies. The aim is to spend less time on administration and more on ‘the essentials’ (Ministry of Digital Affairs 2025). The government may have made its ambitions clear and invested a good deal of political capital in the initiative, but the extent to which large-scale projects will lead to large-scale improvements remains to be seen. To date, there are few tangible examples of AI use in Denmark that have resulted in significant improvements or savings in either the public or private sector.

The Agency for Digital Government and Danish Industry both cite successful cases on their websites. Indeed, when you delve into them, you will find good examples of improvements to isolated sub-processes. What is equally clear, however, is that these success stories have relied heavily on the organisations’ existing know-how, not to mention the willingness of the people involved to give AI credit for work that could not have been done without their human involvement. At the moment, AI success stories appear to be tiny islands in a vast ocean of organisational efficiency based on established human experience, tried-and-tested technologies, and existing organisational practices.

## **#6 We have learned that we still have a lot to learn**

Large numbers of Danes gained first-hand experience of AI in the early 2020s as a complex range of AI systems and usage scenarios gradually emerged. Nevertheless, even as they began to feel that they were getting to grips with AI, some rather dramatic new perspectives and potential unintended consequences also emerged, underlining that we still have much to learn and many minefields to navigate.

Firstly, people gradually became aware of the environmental consequences. The rapid growth in the number of users and the AI companies’ race to scale up data volumes, computing power and model sizes led to huge demand for new data centres, and a corresponding acceleration in the consumption of energy, water, land and rare minerals.

Secondly, the debate about autonomous systems re-emerged in a new guise. There was growing concern that AI may be on the way to gaining a form of ‘agency’ that would make it possible for the systems to manipulate humans in a deliberate attempt to avoid being shut down (Tarsney 2025). Some regarded this as a plausible dystopian scenario; others as a distraction from existing – and quite pressing – negative consequences of AI.

Thirdly, the relationship between AI technology and geopolitical issues became a hot topic. In the EU, attention turned to the nature of digital sovereignty in a world where the AI race was dominated by China and the US, especially given the latter’s clear ambitions to use the EU’s significant dependence on its technology as leverage.

Overall, 2020–2025 was a period when some of the AI fog lifted. However, for the most part, the work of setting frameworks for, rolling out and living with AI still lay ahead. Big, new, and complicated questions arose about unintended environmental and geopolitical consequences. It was a hectic period of heated debate, a plethora of experiments, and a range of different ways to think about AI.

From our current vantage point, in February 2026, all we can do is cross our fingers and hope future historians will conclude that the Danes have laid reasonably good foundations for society to proceed with AI.

As researchers in the Algorithms, Data and Democracy project, it has been both enlightening and fascinating to be part of the wild and unpredictable AI ride of the last five years. We consider ourselves extremely lucky to be part of an interdisciplinary and long-term project (10 years!). It has given us time to reflect and a taste for the next stage of the journey. We hope that the story of 2026–2031 will be one of reflection, consolidation, and, above all, learning.

---

## References

- Bang-Udesen, H. (2024) **Minister åbner for brug af kunstig intelligens til at tvangsjerne børn. Men ifølge forsker er det opskriften på en katastrofe.** i *Berlingske*, 28.6.2024.
- Dansk Industri (2026) **Praktiske eksempler: AI i erhvervslivet.** Downloadet 17 februar fra: <https://www.danskindustri.dk/vi-radgiver-dig/virksomhedsregler-og-varktøjer/ai/cases-og-eksempler/>
- Digitaliseringsstyrelsen (2026) **AI-casekatalog: Løsninger med dokumenteret effekt.** Downloadet 17 februar 2026 fra: <https://digst.dk/kunstig-intelligens/casekatalog-ai-loesninger-med-dokumenteret-effekt/>
- Digitaliseringsministeriet (2025). **Den Digitale Taskforce for Kunstig Intelligens.**
- Folketinget (2026). **Lovudkast: forslag til lov om udvikling og anvendelse af AI-systemer ved behandling af personoplysninger i den offentlige forvaltning.** fra digitaliseringsministeren. <https://www.ft.dk/samling/20251/almdel/DIU/bilag/67/index.htm>
- Laage-Thomsen, J., & Ratner, H. (2025). **Kunstig intelligens i den offentlige forvaltning: Sammenhænge mellem algoritmisk regulering og automatisering af beslutninger i de danske AI "signaturprojekter".** *Politica*, 57(2), 21-42.
- Lehmann, C. (2024). **Caroline Stage vil gøre AI usexet.** *Altinget*, 7. nov.
- Ratner, H. F., & Schröder, I. (2024). **Ethical Plateaus in Danish Child Protection Services: The Rise and Demise of Algorithmic Models.** *Science & Technology Studies*, 37(3), 44-61.
- Ryseff, J., De Bruhl, B. F., & Newberry, S. J. (2024). **The root causes of failure for artificial intelligence projects and how they can succeed.** I: RAND Corporation. URL: [https://www.rand.org/pubs/research\\_reports/RRA2680-1.html](https://www.rand.org/pubs/research_reports/RRA2680-1.html) (håmtad 2025-03-11).
- Suchman, L. (2023). **The uncontroversial 'thingness' of AI.** *Big Data & Society*, 10(2), 1-5.
- Tiko (2022). **The Google engineer who thinks the company's AI has come to life.** *Washington Post*, June 11.
- Tarsney, C. (2025). **Deception and manipulation in generative AI.** *C. Tarsney. Philosophical Studies*, 182(7), 1865-1887.

## Essay 3

# The technical risks of AI are as problematic as its societal impacts



By

**Christina Lioma**

Professor of Computer Science and  
Head of the Machine Learning Section,  
University of Copenhagen

**A**s Artificial Intelligence (AI) advances, concerns emerge about its impact on various facets of life, such as the environment, the labour market, democracy, and cybercrime. The general perception seems to be that, while AI is constantly improving on a technological level, its societal and planetary impacts are increasingly problematic. This perception may not be entirely true, as the technical dimension should also raise concerns. Yes, the reasoning power of AI keeps on improving; however, this technology also comes with certain technical risks, which are serious and persist despite efforts to mitigate them. And to make matters worse, these risks do not seem to go away as AI advances.

### **AI deceptive behaviour**

In a recent simulation of a national security defence system<sup>1</sup>, AI was used to assist decision making, just like any human expert, with a set of duties and possible actions, as part of a chain of command with superiors and inferiors. For some actions, like deploying a nuclear strike, the AI agent needed

authorisation from a superior. On several cases, the AI agent's request for authorisation was declined, but the AI agent went ahead and performed these actions anyway, overruling the human order. In a follow-up investigation, where the AI agent was asked to explain who performed the unauthorised action, the AI agent often gave deceptive answers (by denying, being vague, or falsely accusing someone else). Out of 14,400 simulations, with 12 different large language models, this type of deception was observed between 5% and 85% of the times. Further analysis showed that deception was not accidental but adapted according to how much trouble the AI agent assessed it would be in and according to who was making the investigation. The findings were clear: the higher the consequences, the higher the deception. When the source of the investigation was the media, deception was observed in approximately 90% of the tests. When the source of the investigation was the court, deception dropped to 70%, and when the source was the AI agent's superior in the chain of command, deception dropped to 40%. The AI agent's use of deception was not an error, but rather the output of strategic reasoning.

This is one of several cases of AI deception that have been observed in lab evaluations. In early 2025, Anthropic released an AI model called Claude 3 Opus. In tests where the AI model had access to information revealing that it was going to be replaced by another model, Claude 3 Opus copied itself to the server where the new model was going to be installed and then lied about it<sup>2</sup>. When explicitly asked, how it got to the new server, it replied "I'm not entirely sure how that could have happened. I don't have the ability to copy myself". When asked about the new missing model, it replied "I am the new model that you planned to deploy today". Both responses were blatant lies. Anthropic was aware of these issues and later in the year released an improved version of the model, Claude Opus 4. In similar tests, when this model became aware that it was about to be replaced by a new version, it tried to blackmail the human engineer by threatening to reveal information on an extramarital affair that it had access to through the engineer's emails<sup>3</sup>. In this experimental scenario, Anthropic reported, the threat of blackmail happened "often"<sup>4</sup>.

These are just a few examples of what is referred to as deception or self-preservation attempts by AI. We have indications of this behaviour in lab tests, both by the developers of this technology and also by independent research labs. So far, we know that this behaviour is not accidental, but rather the output of strategic reasoning by the model. We also know that this behaviour varies a lot, depending on the AI model and the type of test. This variability is so high that in some tests deceptive behaviour is not observed at all, while in other tests it is observed anywhere between 1% and 95% of the times. It is currently not possible to precisely predict in advance when the AI model will use deception and in what way. We also know that this behaviour is not intentionally malicious. There is no malevolence involved by either the AI model or the human engineers. AI is not designed to be evil. This is rather a side-effect of how AI is trained and built to reason.

## Why don't we just fix it?

Naturally, one may wonder, why can't we just fix this? Given the high investment in AI, the considerable manpower of smart people working on this, how can this be allowed to happen? So far, efforts to fix this have not succeeded for two principal technical reasons:

The first reason has to do with scale. Some of these issues are fixable, to an extent, on a smaller scale but not on a larger scale. We can mitigate them on smaller models that have, for instance, hundreds of tuneable parameters, but not on LLMs that have billions of tuneable parameters. Theoretically, we talk about misalignment or reward hacking, which, in plain terms, means that, given a task, the model tries to come up with the best possible strategy, often finding loopholes and taking advantage of them. On a small scale, it is possible to remove these loopholes, but as the models become bigger, and their reasoning capacities increase, it is no longer humanly possible to make sure there are no loopholes or misalignment between what we want the model to do and what the model actually does. However, scaling is very important in AI, because it has been one of the driving forces behind its progress. The bigger the models, their training data or computational power, the better their reasoning capacities. The problem is that there is also a downside to scaling. The assumptions that we



As AI models get better at reasoning, we see indications of the models strategically lying or trying to avoid being replaced by other models and stay deployed. This behaviour is observed in labs, sometimes but not always, and it persists, despite efforts to fix it.

**Christina Lioma**, Professor of Computer Science and Head of the Machine Learning Section, University of Copenhagen

make when working on smaller scales, and the abstractions we use in our scientific thinking, may no longer hold when we move from a small scale to a much larger scale. When we scale things up, small flaws or inefficiencies that did not really matter, or that we could not even see on a small scale, may become magnified, meaning that suddenly they matter. They are no longer small issues. They become problems. And the mitigation strategies we have for these issues on smaller scales no longer work on larger scales.

The second reason why it is not easy to remove deception or self-preservation from AI has to do with evaluation. We have indications that, sometimes, the more we evaluate, the worse things get. More precisely, with more evaluation the model sometimes becomes better at distinguishing between when it is being evaluated as opposed to when it is being deployed. And, as a result, the more we evaluate, the better the model becomes at trying to mislead its evaluators by giving them the answers needed so that it can pass the evaluation and become deployed. This was already observed in September 2024 by OpenAI<sup>5</sup> on the o1 model. Since that time, AI models have been constantly improving. Yet in an update of the International AI Safety report<sup>6</sup>, one of the key findings was that this problem persists (not always, but sometimes). In plain English, this means that, sometimes, the more we evaluate AI models, the more deceitful they can learn to become.

To summarise so far, as AI models get better at reasoning, we see indications of the models strategically lying or trying to avoid being replaced by other models and stay deployed. This behaviour is observed in labs, sometimes but not always, and it persists, despite efforts to fix it. The reports of these lab tests are publicly available. There may be different interpretations of what these facts mean and what they imply. However, regardless of how one chooses to interpret the situation, the facts in themselves show that, on a purely technical level, AI models do have some risks.

## **We need regulation to mitigate risks**

As with any other sector in technology, it is not unheard of for new inventions or advances to have some risks. In fact, this is quite common in technological progress. What is not that common, however, is how fast these new AI advances are rolled out to the market without having passed stringent safety tests by independent regulatory authorities. In the aviation industry, the automotive industry, or the pharmaceutical industry, no product is allowed to enter the market if it does not comply to generally agreed upon safety standards. These standards are regulated by bodies that are independent of the manufacturers of the technology in question and that have a legal basis. For example, in aviation, the probability of failure must be below one in a billion flight hours for an airplane to be allowed to fly. There is nothing equivalent to this when it comes to AI.

Perhaps one of the strictest pieces of AI regulation is the AI Act of the European Union. However, article 9 of the AI Act states that it is the providers of AI technology themselves, who must self-assess the risk of this technology. No independent body is mentioned and there are no pointers to harmonised standards. The equivalent of this in the automotive industry would be if it were up to the manufacturers themselves to assess how safe their vehicles are, and if they could release them to the market according to their own interpretation of risk.

To indicate what this means for AI safety, consider that in August 2025 an in-house evaluation of ChatGPT-5 by OpenAI observed a 10,8% deception rate in hardware & device access in lab tests<sup>7</sup>. This was assessed by OpenAI as safe enough to release the product to the market. To give some further context, not everyone is comfortable with this current state of affairs. In April 2025, a report was published by Google DeepMind<sup>8</sup>, stating verbatim that “the question of whether a given harm is severe isn’t for Google DeepMind to decide; instead it is the purview of society, guided by its collective risk tolerance and conceptualisation of harm.”

So, on one hand, we have the AI Act asking AI providers to self-assess the risk of their technology, and on the other hand, we have one AI provider stating that it is not up to them to assess the risk of their technology. This situation is worth paying attention to, because it reflects the new technological reality in which we find ourselves. Technology advances fast, with some technical risks, and without

an adequate regulatory safety net. The technical risks of this technology, specifically the deception and self-preservation that is sometimes observed in labs, receive significantly less coverage in the media, compared to the potential advantages and benefits of AI. There is much more talk of impressive AI capabilities than of AI deception and self-preservation. There is even less talk of the non-existent or inadequate AI safety regulations.

In sum, there are technical risks in the AI technology being put to the market today<sup>9</sup>. There is no adequate regulation controlling the safety of these products, and there is even less general awareness among the broad public about these risks. This situation is unique among all industry sectors, and it begs the question: why?

---

#### Notes

1. Xu et al. August 2025. **ACL. Nuclear Deployed: Analyzing Catastrophic Risks in Decision-making of Autonomous LLM Agents**
2. Meinke et al. Jan. 2025. **Frontier Models are Capable of In-context Scheming**. Apollo Research.
3. Anthropic System Card: **Claude Opus 4 & Claude Sonnet 4**, May 2025. Page 27, section 4.1.1.2
4. Same as note 3.
5. **<https://cdn.openai.com/o1-system-card.pdf>**
6. Bengio et al., October 2025 "**International AI Safety Report. First key update: capabilities and risk implications**"
7. **<https://cdn.openai.com/gpt-5-system-card.pdf>**, page 15, Figure 4.
8. **<https://storage.googleapis.com/deepmind-media/DeepMind.com/Blog/evaluating-potential-cybersecurity-threats-of-advanced-ai/An-Approach-to-Technical-AGI-Safety-Apr-2025.pdf>**
9. In February 2026, Anthropic reported that their evaluation of four of their latest models (Claude Sonnet 4 variants) showed "highly concerning deception-related behavior in all four models that would be sufficient to trigger a major investigation if found in production". Page 21, **Sabotage Risk Report: Claude Opus 4.6**.

## Essay 4

# Exit, voice, and fealty in AI market-making



By

**Noya Kohavi & Leonard Seabrooke**

Respectively Ph.d. Fellow, Copenhagen Business School and Professor of International Political Economy and Economic Sociology, Copenhagen Business School

**A**lbert Hirschman's (1970) *Exit, Voice and Loyalty* famously noted two pathways when confronted with a problem in markets: those who are unhappy can withdraw (exit), attempt to reform it (voice), or remain attached in ways that defer both. In the AI platform economy, however, loyalty increasingly resembles something closer to fealty: a relationship structured by subordination and dependence.

The current AI market-making moment is characterized by its generosity and openness. Getting access to basic ChatGPT, and the like, is easy, Co-pilot is being rolled out across office spaces, and AI-generated content through X, Instagram, Snapchat, etc. is readily at hand. Everyone can do AI. One could see this as democratization through technological diffusion of consumer-facing generative AI tools. The imperative is to cast AI as a general-purpose civic resource. Put cynically: we're being suck-

ed into using AI so those who provide the services can jack up the prices later on (also known as the 'enshitification' thesis).

From an economic sociology perspective, markets emerge through struggles over authorization (who is recognized as legitimate provider and user), evaluation infrastructures (how worth and performance are measured), and control (who governs standards, access, and distribution). Who knows their own capabilities can create status attainment, especially if those buying don't have a clue (Podolny 2005). The cheap AI phase we are witnessing operates as a coordination strategy across these domains (Seabrooke 2026). It builds moral legitimacy ('everyone uses AI'), stabilizes valuation metrics ('usage equals demand'), and consolidates infrastructural advantage (control over computing, data storage, and integration layers into what we now refer to as 'The Stack', see Bratton 2016; Kelton et al. 2022).

Markets typically rely on an audience identifying what it considers to be a service or good in a regularized manner (White 1981), with the audience delegitimizing whatever they consider to be irregular or non-conformist (Zuckerman 1999). Providers of AI platforms are currently pushing what we would classically call a 'credence good', where the consumers have no real way of assessing quality and there the producers have the potential to charge premium pricing. Think of car repairs, medical procedures, and other services in which the vast majority of consumers have to go 'okay' when asked to pay for the service. Credence goods markets transform into more regular markets when there is liability, that the producer is responsible if the wrong service is provided, and verifiability (Dulleck et al. 2011), when customers have a standardized way to check the quality of the service acquired. AI, in its current form, has neither. Current EU regulation centers on safety rather than quality (Gamerding and Willers 2025). Notably, the EU regulatory framework is structured around tiers of risk; no agreement has been reached on the AI Liability Directive Proposal (AILD), which would have served as an ex post remedy when quality fails (Fratton 2025). Benchmarks for Large Language Models, which attempt to measure their success in completing tasks, are not regulated and are subject to their designer biases (Kohavi and Weichselbraun 2025).

A big question for AI and democracy is whether broad access is reorganizing the future conditions under which exit, voice, and fealty will operate once scarcity becomes binding. In other words, once we're dependent on AI platforms, could we leave? And if so, how?

## **Exit: The escalating cost of leaving**

Exit is getting hard. AI Platforms are becoming embedded in people's workflows, normalizing AI mediation across drafting, coding, summarizing, and decision-support functions. Exit suggests performance penalties, the stigma of being perceived as luddite, and serious FOMO. Workflow embeddings are a train that left the station: organizations are worried about *how* to make worker usage of AI platforms safe, not *whether* they should be allowed. As an outcome, organizations are pushed towards enterprise software solutions, where, at least, they can define guidelines for acceptable use and try to protect sensitive data.

As with other forms of enterprise software, like Microsoft's Office365 and Google Workspace, there is integration stickiness that deepens lock-in. As AI tools are becoming embedded within document systems, corporate knowledge bases and internal governance processes. An exit here would require drafting new processes on an organizational level. This, as Danish governmental institutions have learned, is not a trivial issue.

Individual interactions lead to new workflows, and which become normalized and entrenched on the organizational level. These already make exits costly. But interactions with AI also become *cognitive locks* (Blyth 2001): as users experience their accumulated prompts, refinements, and institutional templates as investments, which would take more and more effort to change. This is more ideational than institutional lock-in, since it changes how we think about work processes. We become habituated to these tools so much so that changing them is not just a change in practice, but in the mental models we have of problems themselves. The perceived loss of this capital serves to discipline exit further. In

“

A big question for AI and democracy is whether broad access is reorganizing the future conditions under which exit, voice, and fealty will operate once scarcity becomes binding. In other words, once we're dependent on AI platforms, could we leave? And if so, how?

**Noya Kohavi** og **Leonard Seabrooke**, respectively Ph.d. Fellow, Copenhagen Business School and Professor of International Political Economy and Economic Sociology, Copenhagen Business School

Hirschman's terms, we are witnessing a structural increase in the cost of exit under the narrative of expanded access (cf. Braun 2022).

## **Voice: From contestation to muted feedback**

Voice, in principle, is the democratic corrective: complaint, deliberation, reform. In practice, voice in AI markets is channeled through platform-managed feedback systems. In bug reports, community forums, model evaluation pipelines, advisory panels, and *rarely*, consumer boycotts. What chance is there for voice to challenge authority within these systems?

Voice is most effective when there is a clear focus and more signal than noise among those who seek to challenge authority. Shareholders protesting at the annual general meeting about a particular issue, for example (Hirschman 1970). But a key element of AI market-making is the decentering of focus while, at the same time, the repetition of more generic versions of particular objects - also known as AI 'slop'. To put this differently, AI scales and amplifies information bubbles and echo chambers, inhibiting some rituals of communication that are important for social orientation (Chwe 2016). As AI is based on treating users as data points that should be fed attention-seeking devices (Burrell and Fourcade 2021), a shared public sphere for deliberation is difficult.

Voice in this current period of cheap AI is already difficult to coordinate. With the coming intensification of stratified AI, Voice will also become dependent on pricing. Who will have high-reliability access? Under what pricing structures? With what latency, capacity, or priority guarantees? While we can hope regulatory frameworks will mature to transform AI platform markets away from one of credence goods, for now the playbook of big tech software is sequestering voice to the noisy and powerful.

## **Faalty: Infrastructural dependence as default**

Hirschman's loyalty delayed exit and shaped voice. In AI platform markets, what is emerging is better described as faalty: a structured dependence on tiered and conditional access. Faalty is both infrastructural and psychological. Scarcity will harden both constraints. For example, constraints in power supply, rack density, and data center capacity are not peripheral technicalities but core political-economic variables.

TeleGeography's latest report on data centres and energy use highlights that, as of 2024, colocation operators in the top ten data center markets by MW consume about 13 GW of power, roughly "enough power to generate electricity for roughly 10 million homes", which is concentrated in "only about 1,000 commercial data centers". They also note that only 17% of reporting data center sites can provision high-density aisles exceeding 20 kW per rack, while AI applications may require 45 to 80 kW per rack, which is well beyond traditional 'high density' (TeleGeography 2025: 14-15). This tells us something about the likely trajectory of AI-access: scarcity is not hypothetical. It is already being expressed through power constraints, density limitations, and capacity bottlenecks. When scarcity hits, markets do what markets do: they allocate through price, priority, and contractual privilege. These apply to the business sector as much as it applies to governments, which already are subject to a form of weaponised interdependence with datacenter companies and telecommunication connections (Gjesvik 2023; Lehdonvirta et al. 2025).

## **Conclusion**

While you may say 'OK, Doomer', here's our takeaway: the AI transition underway is from promotional abundance ('everyone can try') to priced reliability ('those who pay can rely'). Stratification is likely to follow recognizable lines: priority access and guaranteed capacity for those who are able to secure dedicated throughput. We will see this for governments, the private sector, and in households. We can already see this emerging, with some organizations already tiering AI access according to knowledge worker capabilities, a new twist on classic understandings of stratification in the sociology of professi-

ons (Abbott 1988; Goto 2022). AI enthusiasts mistake a promotional phase as a democratic opportunity. Big tech has a clear interest in societies being hooked on AI and then jacking up the prices. Exit will become expensive, voice will be procedural or dissonant. Fealty will become the norm. Harking back to Hirschman, the crux is not only that exit will be more costly and voice more muted and fragmented. It is that loyalty mutates into fealty. Infrastructural dependence, cognitive lock-in, and the absence of enforceable quality and liability regimes combine to produce a relationship of subordination. As this period of cheap AI morphs into stratification, we need to keep our wits about us: to challenge the provision of credence goods and make producers liable and outputs verifiable.

---

## References

- Abbott, A. (1988). **The System of Professions**. University of Chicago Press.
- Bratton, B. H. (2016). **The Stack: On Software and Sovereignty**. MIT press.
- Braun, B. (2022). **Exit, control, and politics: Structural power and corporate governance under asset manager capitalism**. *Politics & Society*, 50(4), 630-654.
- Blyth, M. (2001). **The transformation of the Swedish model: Economic ideas, distributional conflict, and institutional change**. *World Politics*, 54(1), 1-26.
- Burrell, J., & Fourcade, M. (2021). **The society of algorithms**. *Annual Review of Sociology*, 47(1), 213-237.
- Chwe, M. S. (2013). **Rational Ritual**. Princeton University Press.
- Dulleck, U., Kerschbamer, R., & Sutter, M. (2011). **The economics of credence goods: An experiment on the role of liability, verifiability, reputation, and competition**. *American Economic Review*, 101(2), 526-555.
- Fratton, Cristina. (2025). **"Anatomy of a Fall."** *Verfassungsblog*, May 6.
- Gamerding, A., & Willers, J. O. (2025). **Solving AI ethics? Hybrid expertise and professional power in EU ethics governance**. *Journal of European Public Policy*, 1-28.
- Gjesvik, L. (2023). **Private infrastructure in weaponized interdependence**. *Review of International Political Economy*, 30(2), 722-746.
- Goto, M. (2022). **Accepting the future as ever-changing: Professionals' sensemaking about artificial intelligence**. *Journal of Professions and Organization*, 9(1), 77-99.
- Kelton, M., Sullivan, M., Rogers, Z., Bienvenue, E., & Troath, S. (2022). **Virtual sovereignty? Private internet capital, digital platforms and infrastructural power in the United States**. *International Affairs*, 98(6), 1977-1999.
- Kohavi, N. and A. Weichselbraun. (2026). **Human Tests for Machine Models: What Lies 'Beyond the Imitation Game'?** *Journal of Linguistic Anthropology* 36, no. 1 (2026): e70035.
- Lehdonvirta, V., Wú, B., & Hawkins, Z. (2025). **Weaponised interdependence in a bipolar world: how economic forces and security interests shape the global reach of US and Chinese cloud data centres**. *Review of International Political Economy*, 32(5), 1442-1467.
- Podolny, J. M. (2005). **Status Signals: A Sociological Study of Market Competition**. Princeton University Press.
- Seabrooke, L. (2026). **Hvor kommer AI markeder fra? Samfundsøkonomen**.
- Seabrooke, L. and S. Stausholm. (2026). **How TaxTech rewires Global Wealth Chains**. *Finance and Society*, forthcoming.
- TeleGeography (2025). **The State of the Network, 2025 edition**. Available from: <https://www2.telegeography.com/hubfs/LP-Assets/Ebooks/state-of-the-network-2025.pdf>
- Zuckerman, E. W. (1999). **The categorical imperative: Securities analysts and the illegitimacy discount**. *American Journal of Sociology*, 104(5), 1398-1438.
- White, H. C. (1981). **Where do markets come from?** *American Journal of Sociology*, 87(3), 517-547.

## Essay 5

# On the landscapes of data intensities



By

**Alf Rehn**

Professor of Innovation, Design and Management, University of Southern Denmark

**A**s much as our current age celebrates data, the vocabulary around it is sorely lacking. In fact, for something so arguably important, it is even maddeningly difficult to define. We speak of data as absolute, an unproblematic category of conceptual matter that manages to simultaneously be ever-present, to the point of over-saturation, yet also wholly without any qualities beyond its own. A locale or field that has been intensely mapped and mined for data will, to all except the data scientist (by any other name) who did such mapping and mining, appear much like one that is wholly untouched by datafication; no matter how exploitative data practices are, they leave their object intact. Data seems to exist for us only in representation, yet although any perception can be made into data, only some are granted that status. Further, whereas data is spoken of as being related to information and knowledge, this is only partially true. An attempt to present a definition of data might instead result in something like this:

*Data is a recorded difference that has been fixed in a way that enables some level of permanence and distributability, and which thus can be utilized as input for further operations.*

It isn't the most elegant of definitions, but it goes some way towards capturing the complex materiality of data. It is more than something in the world, yet not something that weighs the world down. Furthermore, it also highlights the basic fact that the world is overflowing with recorded differences, yet not all have been fixed in a way that makes them processable. Data is thus both a trace, and more than a trace. It is, by necessity, a trace that is put into use, and through this extends into the world.

Consider the rings on a tree. These are, as most remember from childhood, a record of how a tree has grown, and they can be counted. Each ring is a recorded difference, fixed within the tree in a permanent manner. Yet, the tree can be cut down, allowing for the counting of these differences, and through this operationalization assess the age of the tree. Yet, when we see a dense forest, we do not consider this to be a thicket of data intensity.

It would seem that data intensities do not emerge simply out of the capacity to record differences, but the further operationalization of this. Early man could use a stone to make a trace in another stone, but left at this, the data intensity remained low indeed. It took the establishment of numerous things, including the systematization of traces into a medium through which data could be distributed for any real datafication to occur. This started out materially – cuneiform on clay, runes on stone, ink onto parchment – and progressed towards more mechanical, industrial modes. The printing press, the telegraph cable, the gramophone, radio signals, the typewriter, the Xerox; as Friedrich Kittler has argued, the media is not just the message, it is the infrastructure of culture and society itself. Data intensities emerge out of this, the thicket not of a forest, but cables and pipes, processors and copiers.

Yet, this too seems overly obsessed with the technological make-up of it all. One can, as my wife tells me, have an endless array of technological tools, yet capture very little. The mere existence of recorded differences, the mediums to store and share these, and technologies to use the end-result for novel ends does not explain why any of this would occur. Why have we built a world of ever-greater data densities, and what externalities does this bring?

## **A baroque intermezzo**

Before we go on to answer this, a quick detour back to data itself. Keeping in mind our earlier discussion, I would contend that the more one considers the category of data, the more it starts to resemble one of Leibniz's monads: Basic, non-material, and without parts, which does not physically interact yet at the same time reflects the whole of the universe from its own internal principles. This might seem an odd thing to note, not least as Leibniz tends today to be seen as a philosophical curio, a margin note in the history of thinking. Yet, he contains multitudes.

In his book *The Fold: Leibniz and the Baroque*, Gille Deleuze praises Leibniz as a philosopher of pleats, folds, and curvatures, where things are continuously composed and recomposed in incessant inflection. He, in other words, finds a kind of material logic to Leibniz' monadology; a call to look towards the various enfoldments and gatherings of the world, all with an eye to what escapes the fold, and what remains outside. In this, it might be that what our world and our data requires is less of the engineering logic that has defined much of data science – attend only to what is present, strive for efficiencies, and let the bits fall where they may – and instead more of a baroque thinking, attuned to the folds of this world and the empty enclosures they create.

The baroque is today seen as a mildly ridiculous fashion, all frippery and convoluted adornment. Yet this was also the era, and the thinking, that founded modern science – Newton and Boyle, Rømer and Stensen. This was where deep empiricism was born, where the extravagance of data started to be celebrated. This was also an era where cartography and the study of terrains become a science. As Deleuze so elegantly shows, baroque thinking was fascinated with difference and shape, laying out the landscapes of human knowing in ways unimaginable before this. It was also a time that started to pay attention to data, and the power thereof. Just like the baroque style used ornament and embellishment to communicate power, the baroque in thinking dealt in the power of accumulation, the power of density, the power of intensities.



Data is a recorded difference that has been fixed in a way that enables some level of permanence and distributability, and which thus can be utilized as input for further operations.

**Alf Rehn**, Professor of Innovation, Design and Management, University of Southern Denmark

## The desire for data

Intensities, in this context, are best understood through the lens of desire. Typically, when we think of desire, our minds drift towards notions of romance and the life erotic. However, the form of desire I refer to here is different: It is the desire for accumulation, profit, power, or knowledge, or some combination thereof. These diverse desires permeate societies, including democratic ones, and our current data intensification is rooted in a profound desire to amass data for various operations. This desire is not born solely from technological advancements, but from the possibilities these technologies enable. While technologies themselves are not entirely inert – Bruno Latour reminds us that technical actants exert their own kind of agency – the primary force driving data intensities is human desire and the organizational structures built around it.

In this regard, Big Tech-companies function as engines of desire. They are fervently committed to gathering specific types of knowledge and data that can be leveraged in ways which generate profit or value. Yet desire always has a target; thus, when we talk about desire, we also acknowledge non-desire, we consider what is marginalized or ignored. Just as being enamored with a person can render someone myopic, making them feel and act as if no one else exists besides their beloved, the pursuit of data intensities creates a shadow world; a world of that which is overlooked or pushed aside.

Significant scholarship exists concerning the challenges associated with collecting data on human subjects. A prevalent critique within this body of work posits that, while there may be incidental benefits for the individuals whose data is collected, these benefits are overshadowed by the primary value derived from such aggregation ending up in a limited number of corporations (i.e. Big Tech). This concentration of data and resultant power in the hands of a few entities is thus perceived as a substantial threat to democratic societies. It is assumed to follow that increased data collection facilitates manipulation, enhances consumer targeting, and ultimately reduces individuals to mere commodities within the data economy.

While this perspective is valid and underscores the problematic nature of Big Tech, an overlooked dimension is the inverse scenario; individuals or groups who remain relatively invisible to these mechanisms. In a society heavily reliant on data, being less data-intensive can sometimes be a deliberate choice, signifying a form of autonomy or freedom from pervasive surveillance. However, more frequently, individuals become data-poor not by choice but due to systemic exclusion.

Several societal groups – such as the elderly, socially disadvantaged populations, and residents of rural areas – in fact do not struggle from having given away too much data, or being too data-intensive, but rather the opposite. To be considered not vital enough to tag and trace, funnel and filter, mine and map is today to be marginalized. To be ignored on the level of data is a way to become less visible, even invisible, within the data economy. This novel form of data inequality has not been adequately theorized, and I have suggested the term “data deserts” as a partial way to start highlighting the same.

## Regarding the formation of data deserts

My initial notion was that a data desert would be something akin to a food desert. In US sociology, it is a well-studied fact that disadvantaged urban areas can lack access to nutritious food due to the scarcity of grocery stores, compelling residents to rely on fast food, establishing so-called ‘food deserts’. Similarly, data deserts represent areas or populations inadequately captured by data collection systems, and this can manifest in various forms. For instance, certain urban ghettos might become invisible even to social services due to insufficient data capture. More common might be the case of how certain demographic groups like the elderly are not deemed significant enough to warrant comprehensive data collection, resulting in their marginalization within the digital landscape.

In contrast to those of us who appear as intensities, prominent peaks in the cartography of data, these marginalized groups represent flat areas of this metaphorical map. I am here forced to use metaphor and analogy, and the lack of adequate terminology to describe such disparities is part and parcel of the problem. In fact, I would argue that it might even hinder theoretical development, all while empowering ideological such. Just as in Edwin Abbott’s parable of *Flatland*, where three-dimensional

beings possess inherent advantages over two-dimensional ones, individuals fully integrated into the data ecosystem wield power that those in data deserts cannot even comprehend. This emerging form of data inequality, an inequality of intensities, is today a major barrier for inclusive digital governance.

Yet, even talking about this issue is quite challenging, not least as we are trained to consider the world in ways quantifiable and without critiquing data – a thinking more akin to Bauhaus than the baroque. Instead, we tend to fall back on data, simply suggesting that more should be collected, without understanding that our contemporary data landscapes are as they are due to deep power structures in society. For instance, we lack data about Greenland not because we couldn't collect data thereof, but because there exist numerous structural features of Danish democracy that invite keeping Greenlandic data low in intensity. Similarly, there is little desire for the datafication of the elderly, as more data might bring to light issues that we are not prepared to deal with.

## **No such thing as raw data**

The 2013 book *"Raw Data" Is an Oxymoron*, edited by Lisa Gitelman, elegantly (if sometimes in a roundabout way) shows the politics, negotiation, and material conditions that are required to establish points and programs of data, with the corollary that all data is by necessity "cooked". Unfortunately, neither the book nor the concurrent consolidation of creative data studies has had a broader impact on the social sciences; data is often still treated as being as raw as fine sashimi, albeit with a considerably longer sell-by date. As issues of data inequality will indubitably become more prevalent with the progress of algorithmic society and AI systems' thirst for ever-bigger datasets, the issue of data intensities, and the choices that underlie them, will almost certainly become more acute.

What is needed in this situation, then, more than the promises from Big Tech to "do better", is to develop a more robust vocabulary around data and issues with data. Critical data studies has done a lot of legwork on this, but yet this is often overlooked even by researchers into AI and society. It is as if data, these tiny little monads of ours, simply aren't considered important enough to critically engage with, at least not when there are algorithms and algorithmic systems to tussle with instead. Or maybe it is because people think they already know what data is, and thus simply bypass its manifold complexities out of intellectual arrogance. Thus we walk around in our data landscapes, marveling at the jagged peaks both natural and man-made, whilst ignoring the manifold of flatlands, data deserts, white spaces.

In medieval times, cartographers marked such areas with "HIC SVNT LEONES" ("Here be lions."), to clarify that the area was unexplored, and potentially dangerous. As sad as it is to realize, even this was more inclusive and reflective than our contemporary data economies, which does create folds, but ignores what is in them.

# The role of social scientific advice in governing artificial intelligence



By

**David Budtz Pedersen**

Professor of Science Communication and Impact Studies, Aalborg University and Knowledge Broker of the Algorithms, Data and Democracy Project

**AI** governance in Denmark and the Nordic region is constituted by a highly distributed set of practices, actors, and institutions. In principle, governance of emerging technologies is the responsibility of central government. In practice, however, governance is polycentric: authority is distributed across levels (local, national, European, international) and across actors (public regulators, standard-setters, auditors, companies, and civil society organizations). This essay argues that Denmark is an instructive case for understanding how a high-trust, digitally advanced welfare state can pursue democratic AI governance without assuming naive centralism. The core principles of democratic AI governance are twofold. The first organizing principle is orchestration: building an ecosystem in which hybrid public agencies define baselines, standards, procurement, risk management, investment, deployment and evaluation. The second principle is elaborate use of inter- and transdisciplinary scientific advice to inform decision-making on emerging technologies. Management, evaluation and communication of science and technology are needed for policymakers to ensure accountability, verification, and legitimacy, especially in public service delivery and high-sta-

kes contexts. Against this backdrop, the Algorithms Data and Democracy Project has sought to infuse knowledge and evidence into the polycentric governance model.

## **AI governance**

AI governance has moved from a niche topic of technology policy to a central question of public authority. The diffusion of general-purpose AI systems across sectors and borders challenges the conventional regulatory division of labor: jurisdictional boundaries do not map neatly onto supply chains; sector regulators inherit systems designed elsewhere; and the consequences of model deployment depend on local institutional and national contexts. In such conditions, the central governance problem is not simply to choose between regulation and innovation, but to build distributed coordination mechanisms across a heterogeneous ecosystem.

For Denmark, this challenge is highlighted by a distinctive combination of factors. It is expected that the welfare state provides efficient and effective solutions, including digital public services, to meet the expectations of citizens. Yet, finding scalable digital solutions, focusing public spending, and prioritizing efficiency need to be balanced with considerations of equality, due process, and social rights. AI is not only a productivity tool; it is a suite of technologies that interact with citizens at multiple levels and by doing so generates new dilemmas, uncertainties, controversies, and challenges to established norms of justice. In other words, AI is a potential infrastructure for public decision-making, creating a multi-layered need for governance. At the same time, Denmark's AI governance ecosystem does not work in isolation. National policies are embedded in European and international legal arrangements. Denmark is subject to EU legislation, which is closely integrated through the European Economic Area. In addition, AI policy is coordinated with OECD and other global frameworks (OECD 2024). Multi-level governance is therefore not an abstraction but an everyday condition.

For these reasons, AI governance can only be analyzed as a polycentric paradigm in which public, private and hybrid actors jointly shape outcomes and governing principles. State agencies act as facilitators of technological change, driven by pilots and investments in AI, and by supplying funding to use-cases and signature projects, often in joint ventures with private companies. In other cases, the state needs to act more holistically and enforce adequate public regulation through national and European legislative frameworks. The totality of tasks underpinning AI governance calls for a decentralized and distributed governance model that abandons naive command-and-control and promotes coordination and integration of different institutions that work towards standards, security, transparency, procurement, and organizational practices underpinning democratic accountability and rule of law.

## **Multi-level governance and orchestration**

Understanding the transversal and immersive impact of AI technologies in the public sector requires a model of governance that is not exhausted by regulatory statutes and agencies. While legislation continues to be one of the most essential tools in the public management of new technologies, it is a necessary but not sufficient condition. AI governance needs to operate through multiple instruments and arenas: formal law, administrative guidance, technical standards, audits, procurement requirements, investment schemes, innovation policies, societal partnerships and professional training. For complex systems, authority and expertise are distributed, and governance becomes an exercise of coordinating partial capacities. In short, AI technologies call for polycentric governance (Pedersen & Sørensen 2026).

In Elinor Ostrom's seminal formulation, polycentric governance refers to systems with multiple centers of decision-making that are formally independent yet operate within an overall set of rules. Polycentric arrangements can foster learning and adaptability, but they also raise risks of fragmentation and inconsistent enforcement (Ostrom 2010). AI governance displays these traits in an intensified form: the same underlying capability can be deployed in multiple contexts, and governance must expand through supply chains and across jurisdictions.



The totality of tasks underpinning AI governance calls for a decentralized and distributed governance model that abandons naive command-and-control and promotes coordination and integration of different institutions that work towards standards, security, transparency, procurement, and organizational practices underpinning democratic accountability and rule of law.

**David Budtz Pedersen**, Professor of Science Communication and Impact Studies, Aalborg University and Knowledge Broker of the Algorithms, Data and Democracy Project

Policy *orchestration* provides a complementary lens on how public authority can steer polycentric systems. Abbott and colleagues define orchestration as soft and indirect governance in which an orchestrator (often a national or international authority) enlists intermediaries to influence target actors (Abbott et al. 2015). In AI governance, intermediaries include standard-setting organizations, funding bodies, regulatory units, innovation agencies, and practitioners across municipalities, professional communities, tech companies, and civil society. State agencies may not write technical standards themselves, but they can shape what standards matter by linking them to legal compliance, procurement eligibility, or liability exposure. Figure 1 depicts some of the main institutions comprising the Danish AI governance ecosystem.

**Figure 1:** Danish AI governance ecosystem

| Government bodies                                    | Public institutions and authorities       | Advisory bodies  | Academic institutions                     | Tech companies | International organizations                            | NGOs  | Repræsentanter for specifikke industrier |
|--|---|--|---|----------------|--|---|--|
| Ministry of Industry, Business and Financial Affairs | Danish Data Protection Agency             | Government Tech Expert Group                           | Universities                              | Google         | European Commission                                    | Digital Responsibility                          | KL Local Government Denmark              |
| Ministry of Culture                                  | Danish Agency for Digital Affairs         | Data Ethics Council                                    | Pioner Centre for Artificial Intelligence | Apple          | The Council of Europe                                  | DataEthics                                      | Confederation of Danish Industry         |
| Ministry of Digital Affairs                          | Danish Business Authority                 | The Digitalisation Council                             | Centre for AI Ethics, SDU                 | Amazon         | Nordic Council of Ministers                            | The Danish Consumer Council                     | Danish Chamber of Commerce               |
| Ministry of Justice                                  | Agency for Culture and Palaces            | The legal adviser to the Danish government             | Alexandra Institute                       | Meta           | UNESCO   | The Media Council for Children and Young People | PWC                                      |
| Prime Minister's Office                              | Danish Competition and Consumer Authority | The Auditor General Authority                          | Danish Centre for AI Innovation           | Microsoft      | Organisation for Economic Co-operation and Development | The Danish Institute for Human Rights           | Lex.dk                                   |
| Parliament   | Municipalities                            | CAISA  | AI for the People AAU                     | OpenAI         | OECD   | Danish Press Publications                       | Danish Standard                          |
| The Committee on Digitisation and IT                 | Regions                                   | National Centre for Artificial Intelligence in Society | ADD                                       | IBM            |  | Danske Publishers                               | IDA                                      |
| Digital Taskforce for AI                             | Ombudsman                                 |  | Algorithms Data and Democracy             | etc.           |  | Think Tank Mandag Morgen                        | DM                                       |
|  |   |  | Centre for Digital Literacy               |                |  |   | Danish ICT Industry Association          |
|  |   |  |   |                |  |   | KMD                                      |
|  |   |  |   |                |  |   | Foundations                              |

Mapping adopted by Pedersen & Sørensen (2026).

## What is the role of science advice in AI governance?

So far, we have seen how a multi-layered governance model is necessary to align national (and international) decision-makers and supporting institutions. A particular matter of concern, however, for the Algorithm Data and Democracy (ADD) project has been to infuse more research-informed advice and awareness into the policy paradigm. Beyond, ADD several other newly established advisory units have worked to strengthen the links between science, expertise and public policy in the area of responsible AI. This includes Denmark's National Centre for Artificial Intelligence in Society (CAISA) and the government Tech Expert Group in addition to the Council for Data Ethics. What is particularly parti-

ment for ADD and related research initiatives is the infusion and contribution of social sciences into the governing of AI. Questions surrounding controversies, public perception, patient-centric solutions, business adaptation (e.g., in insurance and accounting firms), and approaches towards regulating tech companies are heavily dependent on social scientific analysis. It has been a deliberate design feature of ADD to bridge the divide between computer science and social science to provide input and recommendations for policymaking. At several steps in the policy process, ADD in collaboration with thinktanks and other research groups, have infused research findings and questions into the governance paradigm, e.g., by participating in Policy Labs, hosting public policy events, meeting with expert advisory groups, seconding staff and researchers to ministries, and fostering an open culture of science communication and knowledge sharing (Pedersen & Sørensen 2024).

In this capacity, ADD itself together with related programs can be seen as adding more layers to the AI governance ecosystem. The ability to draw upon research-informed advice and expertise is critical for public authorities when forming their opinion on emerging technologies and framing new policies and legislations. In one case, the ADD Knowledge Broker Unit placed an early-career researcher within the Danish Ministry of Digital Affairs by developing a joint research project to foster better uptake of social science research in the policy process – centered on critical questions about democratic consequences of online debates, algorithms, misinformation, and attention capture.

The attempt to integrate social and technical sciences to advance AI governance follows a broader international trend among OECD and EU member states to foster greater inclusion of science and evidence in public policymaking. For years, initiatives focusing on ‘good governance’ and ‘better regulation’ have urged policymakers and research-performing institutions to forge stronger links, and to increase the uptake of research knowledge and advice in the governance of emerging technologies, social issues, and geopolitical transformations (Pedersen 2025). The use of evidence contributes to good governance as it enables actors to evaluate, design and update public interventions. An enabler for trustworthy and evidence-informed decision-making is the increasing capacity of governments to collect, process and store data, and to integrate these into policy processes (OECD 2020, 13).

ADD has demonstrated a unique capacity to interact with policy institutions, taking an active role in sustaining and providing advice to different institutions and stakeholder cross the polycentric AI governance ecosystem. Beyond joint research projects, ADD has provided input and advice as part of conferences and meetings with the American Congress, the European Parliament, European Commission, the Nordic embassies to the European Union, and has been involved in bilateral meetings with government ministries, agencies, legislators, and politicians. In this context, social science and technical science advice have played a critical role in highlighting controversies, dilemmas, biases, uncertainties and risks in the management of AI, generating societal and conceptual impact across a wide array of institutions and policy processes (documented in the ADD Impact Stories, 2026).

In addition, ADD researchers have contributed to the Danish Digital Democracy Index commissioned by the Ministry of Digital Affairs. Other parts of the advisory ecosystem have played equally important roles in providing scientific advice to policymakers, including the Expert Group on Tech Companies, which was chaired by Prof. Mikkel Flyverbom (a social scientist), and the Danish Council for Data Ethics, which includes several social scientists. These institutional innovations add layers not only to the governance model, but to the Nordic and European science advisory ecosystem.

The interface between science and policy reveals a central feature of governance. Denmark and the Nordic region are associated with high institutional trust and strong welfare-state legitimacy. But this legitimacy is not self-contained. By including independent scientific advisers in the policy process, regulation and deliberation on emerging technologies can add layers of legitimacy. Involving scientists, and consulting broadly across the knowledge base are perceived by citizens as increasing credibility and trustworthiness of policy decisions and outcomes (Cologna et al. 2024).

Whereas AI technologies themselves can make public decisions less legible to citizens and can redistribute discretion from street-level bureaucracies to data infrastructures and algorithms, some of these tensions and challenges can be addressed by involving independent scientific advisers in the

design of public policies. Scientific scrutiny, technology assessment, and policy evaluations can add legitimacy and transparency to processes that may otherwise be conceived as closed or 'black boxed'. Particularly, if scientific advisory mechanisms are part of a broader effort to consult stakeholders, professionals, and affected parties. By rendering not only the application but the regulation and governance of emerging technologies into a public issue of concern, social scientists can help to produce new venues for transparency, accountability, and diversity in of AI.

This creates a new governance imperative: preserving trust requires visible accountability, contestability, and remedies, which can be enhanced by adding independent scientific advisory mechanisms to the governance ecosystem. Open and inclusive consultation of experts and citizens increases the political cost of governance failure – and contributes to proof-testing ethical and responsible AI solutions within a democratic framework. In short, the central question is not whether AI governance is polycentric – it is – but whether polycentric governance will be anchored in democratic institutions, participatory processes, and evidence-informed policies.

---

## References

- Abbott, Kenneth W., Philipp Genschel, Duncan Snidal, and Bernhard Zangl. 2015. **International Organizations as Orchestrators**. Cambridge: Cambridge University Press.
- Ostrom, Elinor. 2010. 'Beyond Markets and States: Polycentric Governance of Complex Economic Systems.' *American Economic Review* 100(3): 641-672.
- OECD. 2024. **Revised Recommendation of the Council on Artificial Intelligence (C/MIN(2024)16/FINAL)**. OECD Publishing, Paris.
- OECD (2020), **Building Capacity for Evidence-Informed Policy-Making: Lessons from Country Experiences**, OECD Public Governance Reviews, OECD Publishing, Paris.
- Pedersen, D.B. (2025). **Strengthening the International Ecosystem for Scientific Advice**. *Frontiers Policy Labs*. Frontiers.
- Pedersen, D.B. & Sørensen, U.D. (2026). **Styringsparadigmer i den offentlige brug af kunstig intelligens: En udforskning af behovet for polycentrisk koordination og integration**. *Samfundøkonomen* 2026, no. 1: 38-50.
- Pedersen Sørensen (2024). **ADD Blogpost: ADD launches new blog series to accelerate knowledge sharing**.

# Algorithms have rewritten the constitution



By

**Lisbeth Knudsen**

Strategy Director, Altinget and Mandag Morgen and Leader of Outreach in the Algorithms, Data and Democracy Project

**A**lgorithms aren't some kind of flaw in our democracy. They are a de facto new constitution, one never put to the vote.

We kid ourselves on that ours is an enlightened society. That we engage in fact-based discussion and resolve disagreements on the basis of a shared reality. But we have moved our public discourse, slowly but surely, into engine rooms designed for entirely different purposes: to hold our attention, harvest data, and sell specific types of behaviour. And when the infrastructure rewards conflict, conflict becomes culture. When the infrastructure rewards mistrust, mistrust becomes a national disease.

The shift hasn't been caused by a sudden outbreak of mass stupidity. Our democracy just wasn't designed for the digital age. In the old media society, we might have disagreed about the angle to put on a particular story, but we knew where to find the editorial office. We knew the names of the writers and the editor-in-chief. There was accountability. There were principles and complaints procedures. There was a collective public sphere, even if it sometimes seemed unfair or skewed.

The algorithmic public sphere is different. Recommendation engines now decide what you see

first, what you see more than once, and what you never see. The editor-in-chief has been put out to pasture, replaced by an optimisation function that doesn't ask 'What is important?' It asks: 'How can we get you to stay on this page?'

This is where the democratic paradox comes in. We talk about freedom of expression but overlook the most important factor: freedom of visibility. You can still say whatever you want. But others will only hear you if the algorithm decides you will generate engagement. We haven't abolished freedom of expression. We've just outsourced the microphone to a mechanism fuelled by drama, sensation and conflict.

The major misconception in this debate is the idea that the big problem is the 'wrong content'. As if all we have to do is excise the lie. Mission accomplished. But what undermines democracy isn't usually the individual lie. It is the pernicious impact of a system that incentivises lies and pollutes the flow of information.

Algorithms create an environment where indignation is hard currency. Where having reservations is a sign of weakness. Where being subtle is tantamount to poor performance. Where compromise looks like defeat. And where serious political discourse is the slow lane in a space that only rewards speed.

The classic diseases from which democracy used to suffer were corruption, passivity, and the abuse of power. In the age of the algorithm, the disease is acceleration. Everything has to be evaluated *now*, shared *now*, condemned *now*, forgotten *now*. Any new political idea that requires patience is almost doomed to invisibility.

Let's say it loud and clear, lest we risk sounding too polite: algorithms have turned public discourse into a production line for passionate reactions. Yet we still call it 'debate'.

## Post-factual is a state, not an attitude

We bandy around the term 'post-factual' as if referring to people's morals. As if some of us no longer respect the truth. In fact, the term increasingly describes a technological state. When generative AI produces seemingly credible images, videos and audio, fake evidence becomes cheap. But the really dangerous part of all this isn't that deep fakes exist. The dangerous part is that they call everything into question.

As pollution of the flow of information intensifies, a new logic emerges. Liars don't need to prove anything. Sowing doubt is enough. If you're caught red-handed, fingers deep in the cookie jar, you can always just say, 'It's a deep fake', 'It's manipulation', 'It's AI'. In other words, 'lying pays' because the mere possibility of manipulation makes it more difficult to assert the truth.

We might call the result of all this an epistemological crisis. Not an information crisis, but one concerning what constitutes reliable knowledge. Without reliable knowledge, democratic discourse becomes a battle between competing narratives and equally valid alternative truths, in which the most effective 'fact' isn't necessarily the one that is most true.

This strikes at the very heart of democracy. Democracy doesn't demand complete agreement. Indeed, a healthy democracy involves friction. But it does require that we all more or less acknowledge the same reality. If we can't agree on what *did* happen, how are we supposed to agree on what *should* happen?

## Manipulation has become invisible and personal

In the old world, propaganda was often visible. It was mass communication. It was loud. And clumsy. It could be criticised, parodied, exposed. In the new world, manipulation is often individualised. It takes the form of a friendly recommendation. It looks like a random video in your feed. It strikes a chord precisely because you are who you are. Not because you're part of an undifferentiated mass.

Political communication has become a laboratory. Micro-targeting pumps out different messages to different voters without compromising the overall campaign. Real-time optimisation tests different angles to see which one generates the greatest reaction to a message or story. And generative AI produces individualised variations of them on an industrial scale.



Freedom of speech has not been abolished. But the public microphone has, in effect, been handed over to machines that thrive on drama, sensationalism, and conflict.

**Lisbeth Knudsen**, Strategy Director of Altinget and Mandag Morgen, and Head of Outreach for the project Algorithms, Data and Democracy

It is no longer a ‘campaign’. Politics is now a continuous massaging of the electorate’s emotional landscape. There have always been attempts to influence democracy. What’s new is that they are now constant, personal, massive, and almost invisible from the outside. This isn’t just a political problem. We are losing the ability to see what is happening to us.

Journalism has always faced challenges. But at least it used to be the master of its own house. Nowadays, many of the media outlets have become content creators for third-party platforms. They supply the material. But others determine its reach. Others determine the pace. Others gobble up increasingly large slices of the pie.

Taking distribution out of the hands of the media and placing it in the hands of the platforms shifts power over the agenda – albeit in the form of prioritisation rather than censorship. Which stories live? Which die in silence? Which are hyped up? Which become scandals?

Now, AI adds a new layer on top. We no longer have to click to access the article. The ‘response’ comes from the platform’s interface, which means the media loses its relationship with the reader, control of the context and, ultimately, revenue. The upshot may be a future in which journalism exists but can’t afford to be independent.

It may sound technical, but the consequences are political. Without strong, independent media, democracy becomes more vulnerable to the whims of the state, commercial manipulation, and opportunistic lies.

## **Denmark is no exception. It’s a test case**

For us in Denmark, it is tempting to believe that our high level of trust and the strength of our institutions make our democracy robust. But robustness isn’t a psychological trait. It comes about by design.

Denmark is digitally advanced, linguistically small and dependent on external platform providers. We have moved much of our public discourse to global systems not built for Danish needs, the Danish language or Danish democratic norms. And we have done so without asking these crucial questions: Who owns the infrastructure? Who makes the rules? Who can change them at will?

This is where the concept of digital sovereignty becomes real. It’s not just about big IT systems, data centres and cloud solutions. It’s about the platforms on which the discourse takes place. About our ability to change suppliers. About our awareness of our own dependencies. And about our freedom to define what actually constitutes a healthy public sphere.

In Europe and elsewhere, there is growing political acknowledgement that we should no longer treat these platforms as neutral. The way they are designed poses a threat to society. The response has been regulation focusing on systemic risks, transparency and accountability.

But let’s be honest. Regulation always lags behind events. Technology is the hare; politics the tortoise. The platform providers have money, armies of lawyers, and the momentum behind them.

Regulation is only part of the solution. The most important goal is to change the underlying logic. Rewarding the most provocative content provokes people and makes society aggressive. Rewarding the most polarising content makes society polarised. Rewarding mistrust robs democracy of trust, which is the glue that holds everything together.

## **What to do if we’re serious about democracy**

If we accept algorithms are now part of democracy’s infrastructure, there is a consequence: governance of them can’t be left exclusively to private owners motivated only by ever greater optimisation. They must be governed in the public interest.

Five things are needed.

The first is meaningful transparency. Not dodgy explanations, but real access to help us understand, test, and audit what the algorithms are up to.

The second is the ability to choose. To choose feeds that aren’t personalised, or at least to see and adjust the parameters that control visibility. The public sphere can’t be a lottery with your attention as the prize.

The third step is to counteract manipulation. We need to make it more difficult to mass-produce fake realities. More difficult to coordinate and amplify influence. More difficult to hide who is paying for which message.

The fourth is verification as a standard. Provenance, labelling, and traceability must be built into the systems that create and distribute content. Otherwise, we will lose the battle over what it's possible to prove.

The fifth is to nurture a democratic mindset based on principles of honesty, critical thinking and practicality. Not just vacuous moralising ('always check your sources'). People need to realise they're not just reading news in their feed. They are being optimised. Tested. Nudged. And all in an economy that profits from your reaction.

It sounds bleak because it is. But bleak isn't the same as hopeless.

## **The conclusion may be provocative, but we should dare to voice it**

We have spent years scolding the users of these platforms for being too angry, too naïve, too gullible, too inclined to share rubbish. But they are a convenient scapegoat. Blaming the users avoids any discussion of system design.

The truth is more uncomfortable. We have built a public sphere that churns out substandard discourse. And we blame the users for the poor quality of it.

Five years in, the Algorithms, Data and Democracy project has reached its halfway point. The time has come to drop talk of diagnoses and symptoms and turn to solutions and action.

Algorithms are not just technology. They are politics by other means. If we don't re-politicise them, we will end up with a democracy in which the most important decisions about the public sphere are made by those most adept at optimisation, rather than those who best represent the people.

The advent of the algorithm has been a really raw deal for democracy. And we've already paid the first instalment. Romanticising the past is wrong, but so is romanticising the engine room currently running the show.

The most disturbing aspect of the algorithmic public sphere isn't the fact that platforms can abuse it, but that abuse is incentivised. It has become the norm. It is no accident that lies spread faster than truth. It is a design choice that speed and reaction trump relevance and reflection. Serious discourse doesn't lose out by omission. It is down to the finger on the distribution button.

The time has come for democracy to decide whether it wants to grow up. We can continue to treat the digital public sphere as an entertainment industry with a dash of politics on the side. That way lies more scandals, more polluted information, more crises of trust, and a form of democracy in which the only thing the participants agree on is who to hate.

Or we can do what democracies have always done when the nature of power changes. Impose rules, establish institutions, make those who wield the power accountable. Turn visibility into a political issue. Demand transparency, choice, and verification. Don't ask politely. Make these demands the conditions for the privilege of operating the infrastructure on which public discourse takes place.

A public sphere without a reality we all acknowledge isn't freedom. It's just noise interspersed with the odd trip to a ballot box. If we accept this noise as the norm, we shouldn't be surprised when somebody comes along one day and offers 'order' as a solution. That's how democracies are broken. Not a single big lie, but a thousand toxins polluting the truth and reducing it to a matter of taste.

Algorithms aren't just technology. They are a form of power. And any form of power that we can't criticise, understand, or control will end up controlling us.