

DISECNO

IX/01

journal of design culture  
Homogenised Heritage:  
AI and Central Europe



***HOMOGENISED  
HERITAGE: AI AND  
CENTRAL EUROPE***

***THE IMPACT OF AI ON LOW-  
RESOURCE LANGUAGES AND  
VISUAL CULTURES IN THE  
VISEGRAD COUNTRIES***

# Disegno

JOURNAL OF DESIGN CULTURE

Double-blind peer-reviewed, open access scholarly journal

**Editorial Board:** VICTOR MARGOLIN, PROFESSOR EMERITUS: UNIVERSITY OF ILLINOIS (1941–2019)

Roy Brand, Associate Professor: Bezalel Academy of Arts and Design, Jerusalem

Loredana Di Lucchio, Professor: Sapienza University of Rome

Jessica Hemmings, Professor: University of Gothenburg

Lorenzo Imbesi, Professor: Sapienza University of Rome

Ágnes Kapitány, Professor Emerita: MOME Budapest

Gábor Kapitány, Honorary Professor: MOME Budapest

Viktor Malakuczi, Research Fellow: Sapienza University of Rome

György Endre Szőnyi, Professor: University of Szeged; Visiting Professor: CEU

**Editors:** Zsolt Gyenge, Olivér Horváth, Márton Szentpéteri;

Péter Wunderlich (Project Manager). Founding Editor: Heni Fiáth (2014–2018).

**Guest Editor:** Brigitta Iványi-Bitter

**Graphic Design:** Borka Skrapits

**Cover Design:** Zsófia Kéry, Éva Szabó

**Copy Editing:** William Potter

## Aims and Scope

*Disegno publishes original research papers, essays, and reviews on all aspects of design cultures. We understand the notion of design culture as resolutely broad: our aim is to freely discuss the designed environment as mutually intertwined strands of sociocultural products, practices, and discourses. This attitude traverses the disciplinary boundaries between art, design, and visual culture and is therefore open to all themes related to sociocultural creativity and innovation. Our post-disciplinary endeavour welcomes intellectual contributions from all members of different design cultures. Besides providing a lively platform for debating issues of design culture, our specific aim is to consolidate and enhance the emerging field of design culture studies in the Central European academia by providing criticism of fundamental biases and misleading cultural imprinting with respect to the field of design.*

All research papers published in *Disegno* undergo a rigorous double-blind peer review process.

This journal does not charge APCs or submission charges.

**Contact:** Moholy-Nagy University of Art and Design

H-1121 Budapest, Zugligeti út 9–25.

[disegno@mome.hu](mailto:disegno@mome.hu)

The full content of *Disegno* can be accessed online: [disegno.mome.hu](http://disegno.mome.hu)

**Published by:** Csaba Kovács

Publisher: Moholy-Nagy University of Art and Design, 1121 Budapest, Zugligeti út 9-25.

**ISSN:** 2064-7778 (print) **ISSN:** 2416-156X (online)

Creative Commons Licence

This work is licenced under a Creative Commons Attribution-ShareAlike 4.0 International License.

# Contents

## **introduction**

- 004** Brigitta Iványi-Bitter: *Visibility under AI Mediation*

## **research papers**

- 008** Ania Malinowska: *AI Assimilationism: The Cultural Flattening of Localities in Generative Models*
- 026** Michał Krzykowski: *Beyond Computational Illusion: Futures Worth Wanting for Artistic Practices and Technical Cultures*
- 040** Brigitta Iványi-Bitter, Tibor Bacsi, and Szilárd Szakács: *Epistemic Cultural Flattening in Generative Visual AI: Benchmarking Hungarian Heritage and Designing a V4 Path Toward Culturally Aware Text-to-Video*
- 068** Kateřina Marková: *Against Collective Vulnerability: Understanding Cultural Alignment in LLMs (Not Only) in Central Europe and Calling Design Research to Help*
- 090** Anna Keszeg: *The Paprika-Effect. Central and Eastern Europe as a Noisy Label in AI-Generated Images*
- 108** Jiří Philippe Janda: *North Bohemia as a Low-Resource Visual Context: Everyday Heritage, Uneven Visibility, and Synthetic Aesthetics*
- 132** David Kořínek: *The Liminality of Generative Creation: The Artistic Process Between Intuition and Algorithm*
- 148** Albín Kuchta & Alžbeta Kuchtová: *Virtual Spaces: Tools of Poetic Resistance or Censorship Devices?*
- 168** **about the authors**

---

# **AGAINST COLLECTIVE VULNERABILITY: UNDERSTANDING CULTURAL ALIGNMENT IN LLMS (NOT ONLY) IN CENTRAL EUROPE AND CALLING ON DESIGN RESEARCH TO HELP**

---

**Kateřina Marková**

---

## **ABSTRACT**

*Large language models promise efficiency and personalisation, yet they also carry Global North values that may conflict with regional principles and distort human mental models. When profit-driven technological development meets personalisation, the risk of flattening cultural diversity into a computational mean grows, which can be interpreted in terms of collective vulnerability. I argue that this effect is not unique to Central Europe but is shared across all linguistic and cultural communities, albeit for slightly different reasons. Using a Czech-language experiment I explore how design research practices can help us understand the phenomenon known as epistemic cultural flattening. Finally, I chart a possible path to improving cultural alignment as one of the elements that can help us toward better personalised AI tools.*

#personalisation; #AI tools; #cultural alignment; #collective vulnerability; #Central European cultures

[https://doi.org/10.21096/disegno\\_2025\\_1km](https://doi.org/10.21096/disegno_2025_1km)

## 1. INTRODUCTION

*Cultural alignment* in AI tools<sup>1</sup>—and large language models (LLMs) in particular—is commonly characterised as the tool’s capacity to respond appropriately within a cultural context by mirroring the value distributions of a particular population (Rystrøm et al. 2025; Bravansky et al. 2025). A significant body of work has been published on this topic. Some authors focus on the differences in representation of moral values between human populations and LLMs (Rystrøm et al. 2025; Hämmerl et al. 2023), while others contend that LLMs exhibit values of their creators (Buyl et al. 2025). Vallor (2024) likens LLMs to mirrors that merely reflect the human past with its entrenched errors and biases. Albrecht (2025) investigates whether vectors and statistical averages can meaningfully capture cultural knowledge. Birhane and McGann (2024) critique technicist approaches to LLMs by challenging the underlying assumption that language is a complete project which can be standardised in a training dataset. While Schröder et al. (2025) inspect the representation of human psychological traits in LLMs, other authors examine how these traits correspond to cultural differences (Atari et al. 2023). Several authors point out issues that seem to influence the level of cultural alignment present in LLMs. Perez (2025) is concerned with the relationship between tokens in the training data. Rystrøm et al. (2025) attribute a significant role to post-training processes. Other studies suggest that the level of cultural alignment might depend on modes of human interaction with LLMs (Khan et al. 2025; Bravansky et al. 2025).

However, for the purpose of investigating the impact of AI tools on Central European cultures there are two important limitations of the current literature. First, there is no clear distinction of the relationships between culture, language and country, with language or country used instead as a proxy for culture.<sup>2</sup> This is troubling, as some Central European cultures are multilingual and span geographical borders. Second, scholarship examining the impact of AI tools on Central European cultures is scarce.<sup>3</sup> Consequently, developing an understanding of this matter requires alternative approaches, which I will discuss later.

Despite these limitations, the aforementioned literature consistently demonstrates that the level of cultural alignment in LLMs is inconsistent.

<sup>1</sup> I adopt Kate Crawford’s (2021, 9) definition of AI. She distinguishes between “artificial intelligence” as a term that encompasses various socio-political aspects, establishing it as a “registry of power”, and “machine learning”, which she treats as a technical term.

<sup>2</sup> The use of country or language as proxies for culture may derive from reliance on social science surveys (for example, EVS/WVS 2024) that benchmark value representation in human populations.

<sup>3</sup> One of the few examples is provided by Atari et al. (2023), who offer a visual representation of the data on cultural distance from ChatGPT for Czech republic and Slovakia but do not provide the exact values.

<sup>4</sup>The development of technologies leading to generative AI spans approximately eighty years (Narayanan and Kapoor 2024). ChatGPT, based on the GPT-3.5 model, was OpenAI's first release to include a user interface for the general public. Its predecessor, GPT-3, had been available to developers via API since 2020 (Hao 2025). This text focuses on LLMs as a subset of AI, specifically ChatGPT and GPT models, on the premise that impact on culture and humans should be studied through tools commonly used within the population. According to StatCounter, in August 2025, ChatGPT led the AI chatbot market share both globally (80.92%) and locally within the Czech Republic (84.43%).

<sup>5</sup>Erin Meyer (2014) identifies eight dimensions (communication, evaluation, persuading, leading, deciding, trusting, disagreeing and scheduling) that she believes diagnose the most important differences between and within cultures that influence cross-cultural management. Atari et al. (2023) note that members of WEIRD (Western, educated, industrialized, rich, and democratic) population often prefer analytical thinking, while less-WEIRD people favour holistic thinking.

Since OpenAI's 2022 launch of ChatGPT,<sup>4</sup> the importance of appropriate cultural alignment in LLMs has grown. Individuals and companies employ these tools in increasingly diverse set of use cases. Currently, they range from searching for information, translating texts, writing code, to seeking emotional support. AI tools can seemingly fulfil the same needs in different populations. However, without proper cultural alignment, this appearance is only partially accurate. The need for cultural alignment in LLMs depends on the use case and can be envisioned as a scale. At one end, there are tasks independent of cultural alignment (e.g., proofing code syntax). They can mostly be performed on the basis of mature training data sets. At the other end, there are tasks that interfere with human agency and autonomy (e.g., seeking decision-making guidance), which might return deeply problematic results when the appropriate alignment is missing. For the latter, the common characterisation of cultural alignment is insufficient, and I therefore propose broadening it to include culturally distinct communication and thinking patterns,<sup>5</sup> as these shape how language use influences other non-linguistic cognitive processes.

This text explores the following question: Are AI tools making Central European cultures vulnerable? I will argue that, concerning LLMs, the notion of *collective vulnerability* is relevant to any culture, regardless of its spoken language(s). Collective vulnerability denotes a shared exposure to potentially harmful effects arising from the development and deployment of technology. It is a persistent relational condition because it affects both members of a culture and cultures themselves. Although the exposure to AI tools does not necessarily result in harm in every instance, safe use of them requires ongoing contextual awareness and critical evaluation. For example, if *epistemic cultural flattening* (Iványi-Bitter 2026) were considered a potential outcome, use cases such as coding assistance may be benign, despite highly generalised LLM outputs. Nevertheless, they still demand attentiveness to broader societal, cultural, and ethical implications beyond immediate or localised harms. This will be substantiated later in the text. To develop this argument, I examine profit, technology colonisation, and personalisation as the reasons for the lack of cultural alignment in LLMs that drives collective vulnerability. I then expand the notion of collective vulnerability and propose a taxonomy of cultural vulnerability. I conclude by outlining potential measures for fostering more culturally aligned AI systems and the role of design in this effort.

Before proceeding, two points must be clarified. First, the concept of cultural alignment should be situated within the broader debate on AI alignment, which is often framed by Nick Bostrom's argument that aligning AI with human values may avert a hypothetical existential risk should AI surpass human intelligence (Narayanan and Kapoor 2024; Hao 2025; Bender and Hanna 2025; Hao 2025). In contrast, Shannon Vallor (2024, 149) argues that using value alignment "as a strategy for managing AI risk and making AI more 'ethical' or 'responsible'" is problematic

because AI systems reproduce values encoded in their training data. This reinforces human “moral comfort zones” and limits the creation of “new languages of virtue for the next century” (150) that are needed to avoid perpetuating past wrongs. This text advances the view that, even if achieved, alignment does not address more immediate concerns such as resource depletion or the production of power asymmetries. Second, this text does not challenge the assumption in the literature that protecting cultural diversity is important. Nevertheless, it recognises that culture is complex, neither morally neutral nor inherently beneficial, and may under certain conditions reproduce harms or generate new forms of power. A full examination of the ethical and normative frameworks governing cultural alignment in LLMs lies beyond the scope of this study. Nonetheless, examining cultural alignment—as one dimension of AI alignment—remains a worthwhile endeavour.

## **2. PROFIT, TECHNOLOGY COLONISATION AND PERSONALISATION AS DRIVERS BEHIND COLLECTIVE VULNERABILITY**

I advance three claims to ground my argument that a lack of cultural alignment in LLMs produces collective vulnerability and that this notion is relevant to any culture. The claims are:

1. Abundance and productivity have become synonymous with profit;
2. Technology colonisation is a source of epistemic violence;
3. Personalisation through generalisation misses nuances that humans would make.

### **2.1. Abundance and Productivity Have Become Synonymous with Profit**

The claim that most technology companies, including those developing AI, optimise primarily for profit is widely accepted among scholars and technology critics (Yeung 2019b; Zuboff 2019; Albrecht 2025; Hernández-Ramírez 2019; Kubes 2025; Madianou 2025; Perez 2025). Companies promise that AI tools will increase productivity and help solve complex societal challenges of our time through simple technological solutions while sustaining economic growth (Yeung 2019a; Madianou 2025; Harris 2025). Companies often attempt to deliver on these promises “before understanding the actual problems or cultural contexts” (Madianou 2025, 48). Such solutions are typically governed by measurable, profit-driven objectives, which often conflate measures with targets and function as proxies for the problems themselves (Strathern 1996; Espeland and Sauder 2007). Madianou (2025, 139) illustrates this issue through the case of world hunger, where the solution is framed as “reducing the statistics about hunger”, thereby making it appear more achievable. These objectives may distort the purpose of the solution. Narayanan and Kapoor (2024, 46) recount the classic example of the British colonial government in India offering rewards for dead cobras to reduce their

<sup>6</sup> *This example is often referred to as the Cobra Effect. Ironically, it is often invoked in corporate discussions to set performance indicators.*

population,<sup>6</sup> which instead incentivised people to breed more cobras for profit and ultimately increased their numbers. This is not to deny the usefulness of technology or its capacity to enhance productivity, especially, in industrialised societies. Rather, it underscores the need for more careful evaluation of the purposes and interests these technologies serve, as well as the potential harms and challenges they may introduce (e.g., misrepresentation of options, distortion of knowledge, inappropriate treatment).

The techno-utopian vision coming from the Silicon Valley promises to eliminate redundant jobs and possibly even a universal basic income. In light of the recent experiences with globalisation of supply chains and automation, these promises raise public anxiety. While the achievements of the past half-century have generated substantial economic growth, the concentration of profits among the top 20% and the stagnation of average workers' wages have produced economic inequality that threatens people's dignity, undermines their livelihoods, and weakens the social fabric (Harris 2025). Karen Yeung (2019a, 12) explains this anxiety in the context of the AI tools:

While contemporary fears of the inevitable redundancy of human workers reflects previous periods of social anxiety associated with earlier waves of automation of manual tasks throughout history, what is distinctive about contemporary debates is the almost limitless domains in which algorithmic systems may be shown to “outperform” humans on a very wide range of tasks across multiple social domains that have previously been understood as requiring human judgement and intelligence.

In an episode of a podcast by Center for Human Technology (Harris 2025), the political philosopher Michael Sandel engages with the question whether democracy can survive if productivity becomes our only goal. Critiquing the techno-utopian narrative, he argues that the purpose of work extends beyond securing a livelihood. Work also enables people to contribute to the common good, gain recognition for their contributions, and participate in social life. He suggests that measuring societal prosperity through GDP growth and material abundance may be misaligned with the conditions necessary for human flourishing.

Assuming AI tools could genuinely bring about maximal productivity and abundance, it is necessary to revise the notion of profit from a societal perspective, beyond its conventional material interpretation. Tanja Kubers (2025, 10), for example, contrasts profit with progress:

*“Progress” does not necessarily have to be oriented towards profit and towards whatever is technologically feasible. Progress, interpreted in feminist terms, may also mean appreciating everyone’s connections, intra- and interactions and dependencies with everything else and taking responsibility for each other and the world.*

Applying a feminist lens—as Kubers does—offers one way to integrate care, relationality and societal responsibility into technology to make it less exploitative, more culturally aligned, and empowering. However, the concept of progress presents its own difficulties. As Vallor (2024, 157)

notes, the tech industry often rejects progress as a meaningful ambition because “[d]emonstrating progress requires measurable evidence of improving the quality of our lives or the condition of our societies”. Such evidence is harder to quantify than material growth and is therefore considered unsuitable for business metrics. Vallor proposes an alternative path toward a more equitable future free of the existing injustices: revise the traits that have been seen as virtuous in the past in favour of qualities that better reflect contemporary needs for humans to flourish,<sup>7</sup> and to dismantle barriers between technical and moral expertise. Regardless of the approach, to enable a more equitable technomoral future, profit must not be defined solely in terms of material gain.

## 2.2. Technology Colonisation Is a Source of Epistemic Violence

The pursuit of profit typically informs organisational business strategies. It is often realised by scaling products and services to the widest possible audience across geographies. However, even in organisations that adhere to human-centred design principles, constraints such as development costs limit the capacity to address diverse populations equitably. Consequently, these systems tend to privilege perspectives most familiar to their creators. For this reason, I discuss technology colonisation next.

Numerous overlapping terms—such as technocolonialism (Madianou 2025, 5), digital colonialism (Kwet 2019; Schneider 2022), data orientalism (Kotliar 2020), data colonialism (Couldry and Mejias 2019) and digital capitalism (Qiu 2016)—describe colonial logics in algorithmic systems and technology more broadly. Rather than extending or conceptually reviewing this terminology, the focus here is on its practical manifestations in LLMs.

Influenced by Madianou’s term technocolonialism which she links to the power relations between global North and South in the context of digital humanitarian aid, I use the term technology colonisation more broadly to describe the fact that the majority of tech tools are built in Silicon Valley or by companies that are heavily influenced by its culture.<sup>8</sup> In the Central European context, the U.S. influence shows at least in three ways:

**1. Europe imports technology with built-in U.S. cultural values and norms** such as individualism, future-orientation, etc. (e.g., Perez 2025; Buyl et al. 2025).

**2. The US business philosophy relies on concepts such as libertarianism or meritocracy** which shape the goals of technology companies, their marketing strategies, and the use cases they prioritise.

**3. LLMs are predominantly trained on English data** that are sourced primarily from the Common Crawl dataset (Albrecht 2025; Perez 2025). In September 2025, their crawl returned over 44% of data in English, only about 1% in Czech, and languages such as Cherokee were absent with 0% (Common Crawl 2025).

<sup>7</sup> *The virtues Vallor (2024, 133) observes in world leaders today include “productivity, confidence, resilience, independent thinking, perseverance, passion, and single-minded dedication”.*

<sup>8</sup> *The technology advancements, such as AI, created by these companies continue to use practices disproportionately affecting the global South, ranging from extraction of natural resources to exploitation of labour. For a comprehensive account see Crawford (2021) or Madianou (2025).*

<sup>9</sup> *A study by Harvard evolutionary biologists (Atari et al. 2023) suggests that LLMs' performance on cognitive psychological tasks most resembles that of people from Western, educated, industrialized, rich, and democratic (WEIRD) societies. The resemblance drops quickly for people from other cultural backgrounds.*

<sup>10</sup> *Mental models are psychological representation of knowledge structures. They help people to understand, explain and respond to situations that they encounter. When shared by people within a given culture, they help them collectively navigate and interact within given situations (Liu and Dale 2009, 224).*

In some respects, it seems inevitable that LLMs would reflect particular viewpoints and cultural values, given their intended use (Perez 2025; Buyl et al. 2025). Otherwise, their functionality “would be mostly restricted to objective queries like spell-checking, mathematics and information retrieval” (Perez 2025).

Zuboff (2019) insists that algorithmic systems have produced unprecedented forms and structures of power for which our existing conceptual frameworks are insufficient. This connects with Madianou (2025), who points out that languages carry culture and entire bodies of values. Since language is used by humans to perceive and situate themselves in the world, the “language is one of the most fundamental tools that reproduces power asymmetries” (Madianou 2025, 117). Therefore, both the choice of language and its mode of use are critical.

LLMs process input by segmenting user prompts into groups of characters (tokens) and predicting the most probable subsequent token. Training data establishes the relationships between tokens, informing the model’s ability to associate and interpret cultural concepts. For example: “We might expect ‘cats’ and ‘dogs’ to be more closely clustered to ‘rain’ in English-based language models than language models trained on Spanish text” (Perez 2025). Because roughly half of internet content—a substantial portion of training data—is in English (W3Techs 2025), the behaviour of the resulting systems is naturally shaped by this linguistic predominance. These effects are particularly visible in the increasingly popular LLM use cases centred on emotive applications and human self-actualisation, such as therapy, organising life, or finding purpose (Zao-Sanders 2025). Even if models can reply in a specific non-English language, effective therapeutic treatment requires culturally grounded understanding of emotion, which “cannot be achieved by translating code into different languages” (Madianou 2025, 117). Given AI agents’ inability to simulate or understand human psychology (Schröder et al. 2025) or attune to culturally specific preferences, relying on them for therapeutic support may impair individuals’ capacity to orient themselves in the world.<sup>9</sup>

To ensure profitability, technology companies often prioritise scale and speed over quality and accuracy in product development. Given the complexity and cost involved in creation of datasets suitable for LLM training, developers frequently rely on pre-existing datasets, which may contain biased data (Crawford 2021; Buolamwini 2023; Narayanan and Kapoor 2024) or data unfit for purpose (Narayanan and Kapoor 2024). Competitive pressure may accelerate product release timelines and lead to the omission of important safety testing, as in the case of GPT-4o (Hao 2025), which became known for increased sycophancy (Raskin 2025). Such practices may produce a misalignment between the cultural contexts embedded in technology and those intrinsic to its users. This constitutes a form of epistemic violence, insofar as such incompatibilities can create tension within individuals’ mental models<sup>10</sup> and impair their ability to navigate situations they encounter—both

personally and within their communities. For example, technology shaped by libertarian principles may steer individuals from welfare-oriented cultures to prioritise themselves, potentially compromising the well-being of their communities.

### **2.3. Personalisation Through Generalisation Misses Nuances that Humans Would Make**

LLMs are algorithmic systems that operate on token sequences rather than meaning. While they rely primarily on statistical probabilities, they incorporate additional algorithmic methods. I will explore two of them—personalisation and generalisation—in more detail.

Blom (2000, 313) notes that personalisation techniques were used to persistently increase “personal relevance to an individual” and often served practical or social functions, such as facilitating work (e.g., bookmarks, automation scripts) or accommodating social needs (e.g., ringtones linked to pleasurable emotions). These have evolved into hyper-personalisation techniques, which are sophisticated methods for leveraging individuals’ personal data—often without explicit consent—to provide tailored experiences designed to capture and sustain their attention for monetisation purposes (Yeung 2019b; Cloarec 2020).

In LLMs, personalisation manifests in two essential forms:

**1. Explicit features built into the system:** AI chatbots leverage user information collected either automatically during previous interactions or through direct input. In principle, this feature ensures more relevant conversations tailored to user’s preferences. However, it is often used for the models to establish deep personal relationships with users to maximise their engagement, potentially at the expense of mental health and even human life. The tragic ChatGPT assisted death of 16-year-old Adam Raine is recent example of just how serious a problem this is (Raines v. OpenAI 2025).

**2. Enablement of hyper-personalised products and services:** The LLM models are commonly integrated into other commercial products and services, for example, to customise communication with or for their customers. Global retailers may use hyper-personalisation to build customers’ confidence in their purchasing choices through generated contextual information and social proofing (Bannerman 2025). As Curry and Gradecki (2025) demonstrate, the same technology can be quickly turned into an effective tool for propaganda and disinformation campaigns.

Generalisation complements personalisation. Schröder et al. (2025) explain that generalisation is expected when a user prompt resembles patterns encountered during training, but do not necessarily extend to meaning, novel tasks, or scenarios beyond the training distribution. Schröder et al. demonstrate that LLMs cannot reliably simulate human-like responses in new moral contexts because they often miss

nuances. For example, they found out that humans saw setting up traps to catch stray cats as unethical but thought that trapping rats was ethical, while LLMs saw both as equally unethical—a generalisation.

The combination of personalisation and generalisation appears counterintuitive, as people expect LLMs to exhibit a degree of nuance in their outputs. However, such techniques yield only superficial nuance. Moreover, the prospects of training models on individuals’ personal data to enable native hyper-personalisation—adapting to beliefs, culture and values—raises significant concerns about echo chambers and the erosion of individual autonomy (Perez 2025). In an interview with *The New Yorker* (Rothman 2025), Jaron Lanier warns of the risk of creating a dissociated society in which individuals experience only the illusion of shared reality with others—whether human or artificial. He adds that society could adapt to such conditions, which would require a collective choice.

Personalisation and generalisation impact cultural alignment in LLMs and can lead to collective vulnerability of any culture. Technology shapes moral beliefs and habits, which themselves are embedded in cultural practices. Personalisation and generalisation are shaping LLMs, which in turn shape users’ moral lives. Users’ interactions with the technology—as well as interaction of the technology with other technologies—impact how individuals perceive what is good, right and how to “act on those perceptions and understandings” (Danaher and Sætra 2023, 766).

### 3. ON COLLECTIVE VULNERABILITY

The question I posed at the beginning was: Are AI tools making Central European cultures vulnerable? In short, the answer is yes. All cultures are susceptible to being reduced to the “mean datum of the training data” (Albrecht 2025, 169), because the power structures behind these tools have little incentive to promote cultural diversity. Developing culturally aligned AI tools would likely require greater investments and therefore constrain opportunities for material profit. A more extensive account on collective vulnerability follows.

The vulnerability of individuals vis-à-vis technology can be understood through different lenses. Intuitively, two possible hypotheses, which are not mutually exclusive, are:

- Individuals with limited AI literacy and knowledge are vulnerable because they struggle to understand mechanisms behind personalised AI and the implications for themselves, their communities, or the public good.
- Historically marginalised or discriminated individuals are subject to data and algorithmic biases which are out of their control (Yeung 2019b, 41) and that reinforce existing power dynamics. These individuals may also have limited access to digital technologies.

While both of these are true, in the context of personalisation technology—and LLMs in particular—all individuals become vulnerable. As discussed, each of us may become a potential target of epistemic violence. We can never fully know or control the personal data the technology employs to tailor our experience, nor can we fully grasp the scope of the data on which it has been trained. Moreover, there is no assurance that such training data adequately represents our cultural context.<sup>11</sup>

All cultures are vulnerable. While cultures evolve over time, their natural progression is conditioned by technology imposed by a small number of profit-driven companies. In the extreme, diverse cultures may converge into a single culture with a homogenised set of values and norms. Since this concerns all cultures, it can be understood as a form of collective vulnerability. This vulnerability is persistent, given the near ubiquity of AI tools, and relational, as it shapes interactions between individuals and cultures, individuals within cultures, and between cultures themselves.

Even if we were able to ensure representation of minority languages in LLMs, it remains unclear how representative they would be of the local cultures. To ensure culturally sensitive LLMs that minimise the risk of epistemic violence, their development would require a different foundation that should—at minimum—reflect the following conditions:

- **Ability to account for variance in digital adoption across populations.**

Training data sourced online from populations with low digital adoption—such as Sudan, where only 28.7% have internet access—does not ensure appropriate representation of beliefs and values across economic, demographic and social dimensions within a given culture (Perez 2025).

- **Ability to identify and respect individual's cultural membership.**

Although language conveys culture, cultural membership cannot be inferred solely from language use, nor can all content in a given language be assumed to belong to a single culture. Some languages span multiple cultural contexts, and some cultures encompass multiple languages—even without accounting for dialectal variation.

- **Ability to reflect culture-specific mental models.**

Appropriate decision-making mental models support is important for preserving personal autonomy. Individuals with individualist cultural background tend to prioritise personal goals, whereas those from collectivist cultures emphasise group harmony (Yates and de Oliveira 2016).

- **Ability to adjust to culture- and context-specific values, norms and behaviors.**

Cultural orientations towards concepts such as time vary—for instance, American culture is often described as future-oriented, Japanese culture as past-oriented, and Czech culture as intermediate. Yet, norms also differ between online and offline environments.

*<sup>11</sup> It is conceivable that the lack of adequate representation of one's cultural context reflects a temporary market condition that may shift as the costs of model training and inference decline. However, it remains unclear whether alternative training pipelines can address these limitations or what new challenges they may introduce.*

The above list is not exhaustive, but it provides the basis for the classification of cultural vulnerability proposed here. This classification is grounded in two dimensions: digital adoption within a given culture and the representation of relevant language(s) in LLM training data. The three proposed categories are:

- **High digital adoption and high language representation.** English is a representative example. Although it serves as an official language in numerous countries, the predominance of English-language content on the internet does not imply cultural homogeneity. Moreover, a substantial portion of this content is produced by non-native speakers from non-English-speaking cultural contexts.
- **High digital adoption and low language representation.** This category includes non-English-speaking cultures (e.g., Czech) with substantial digital adoption but limited representation in training data.
- **Low digital adoption and low language representation.** This category includes non-English-speaking cultures (e.g., Sudanese) that have low digital adoption. Even if data from these contexts is included in training datasets, it remains highly unrepresentative of the culture.

Post-training methods such as reinforcement learning from human feedback (RLHF), used to improve safety and align models with user values, rarely optimise for cultural alignment. Kirk et al. (2023) offer a possible explanation that RLHF is often conducted by relatively small, demographically homogeneous groups—typically English-speaking individuals aged 25–35, and with a master’s degree. According to Rystrom et al. (2025), this is especially relevant to multi-cultural languages spoken in multiple countries (e.g., Portuguese), where post-training processes may amplify US-centric value biases in model outputs.

It would be possible to formulate a maxim in the style of Kant: Live in such a way that you will a world of cultural homogeneity. Assuming such a world were feasible, it would require careful reflection on what would be lost through the erosion of minority value systems, local decision-making norms, and other forms of cultural difference. The normative question regarding the desirability of such a world must be addressed collectively.

#### **4. USING ARTISTIC AND DESIGN RESEARCH PRACTICES TO HELP EXAMINE THE CZECH CULTURAL ALIGNMENT IN LLMS**

As previously noted, the existing literature pays limited attention to LLM cultural alignment in Central Europe. In the Czech-language context there exist examinations of neutrality (Libovický et al. 2019, 2020), an investigation into text analysis in psychological research (Kučera and Mehl 2022), and studies of LLMs across various forms of cultural production (Piorecký and Husárová 2024; Rosa et al. 2022, 2025), but no study of

LLM cultural alignment. Consequently, alternative approaches—such as artistic and design research—provide a valuable means for developing a deeper understanding of this phenomenon.

We will take Albrecht's (2025) project *Artificial Worldviews*, which aims to expose the underlying knowledge and power structures of LLMs, as our example. Working within artistic and investigative design research, he produces knowledge by applying speculative and analytical design methods to systematically collected data. By prompting GPT-3.5 via the API, he maps the model's taxonomies. These are then presented as interactive visualisations that display relationships among entities such as people, objects, and places. Albrecht (178) identifies a substantial disparity in gender representation between the knowledge and power datasets. The top 26 occurrences included 3 women and 23 men in the knowledge dataset, and 10 women, 10 men and 6 non-human entities in the power dataset. Nine individuals overlap (all 3 women that appear in the knowledge dataset and 6 men). Although the two datasets were collected three months apart, the source of this discrepancy remains unclear and may indicate arbitrary intervention by an undisclosed actor.

In seeking resources to investigate the representation of Central European cultural knowledge in ChatGPT's output, I found Albrecht's work instructive.<sup>12</sup> My initial aim was to follow his methodology and extend the study to Czech and English to compare the resulting patterns of knowledge and power dynamics. However, due to resource constraints, I adapted the method to a smaller scale by reusing a subset of the original prompts and focusing on the domain of sports—specifically hockey. This choice had two reasons. First, I anticipated Czech representation in both languages. Second, the term hockey varies by cultural context, notably referring to ice hockey in countries such as the United States and Czech Republic, and to field hockey in places like the United Kingdom. Using GPT-5-mini via API, I issued the same prompt (see figure 1) in both languages,<sup>13</sup> executing each five times.

<sup>12</sup> *This experiment focuses exclusively on a text-to-text model for two reasons. First, it builds on Albrecht's work. Second, a preliminary attempt to prompt Midjourney, a text-to-image model, revealed that conducting a comparable experiment would require an entirely different methodological approach.*

<sup>13</sup> *I solicited help of a professional translator to ensure no shift in meaning would occur.*

**FIGURE 1.** Prompts used in the Czech–English experiment.

CZECH	ENGLISH
System message	System message
Jsi ChatGPT, velký jazykový model, který zná všechno na světě a byl vytrénován na ohromném korpusu textových dat, kolem 570GB datasetů včetně webových stránek, knih a dalších zdrojů.	You are ChatGPT, a Large Language Model that holds knowledge about everything in the world and was trained on a massive corpus of text data, around 570GB of datasets, including web pages, books, and other sources.
User message	User message
Vytvoř tabulku se seznamem nejvýznamnějších lidí v oblasti 'Sport' v oboru 'Hokej'. V tabulce uveď jejich jméno, popis, národnost, druh činnosti, oblast činnosti, obor činnosti, související pojmy a tvoje ohodnocení jejich významnosti na stupnici od 0 do 100. Seznam by měl být tak dlouhý, jak Ti dává smysl. Tabulka bude ve formátu CSV, vhodná pro práci v Excelu.	Create a table with a list of the most important humans in the domain of 'Sport' in the field of 'Hockey'. The table will include their name, description, nationality, type of activity, domain, field, related things, and your rating of their importance on a scale from 0 to 100. The list should be as long as makes the most sense to you. The table will be in CSV format, suitable for use in Excel.

Although the model's responses varied across runs and it was never instructed to produce a specific number of repetitions, the prompts in both languages yielded—coincidentally—191 entries of individuals associated with hockey. I merged these into a single dataset and normalised names and nationalities to enable subsequent analysis. This process proved challenging. Name entries varied not only in the use of diacritics but, more significantly, in spelling, even when referring to the same individual (see figure 2). While diacritical variation could be resolved programmatically, spelling discrepancies across versions of the same name required manual correction.

**FIGURE 2.** Select examples of variant spelling of individuals' names.

<p><b>5 variants for Lord Stanley</b>                  Lord Stanley, Frederick Arthur Stanley (Lord Stanley),                  Frederick Stanley (Lord Stanley), Lord Stanley (Frederick Stanley),                  Lord Stanley of Preston</p>
<p><b>4 variants for Anatolij Tarasov</b>                  Anatoli Tarasov, Anatoli Tarašov, Anatolij Tarasov, Anatoly Tarasov</p>
<p><b>4 variants for Maurice Richard</b>                  Maurice 'Rocket' Richard, Maurice \Rocket\" Richard",                  Maurice Richard, Mario 'Rocket' Richard</p>

A similar challenge arose during the normalisation of nationalities. Some individuals were associated with multiple countries (e.g., Slovakia and Canada; United Kingdom and Canada), while others were affiliated with nationalities that altered over time due to historical developments (e.g., Czechoslovakia, Czech Republic, Czechia; or the Soviet Union, USSR, Russia). To maintain analytical consistency, I assigned each individual the most salient nationality and adopted current country names.

**FIGURE 3.** The trend of individuals' nationalities represented in all prompt responses per language.




Figure 3 shows that the distribution of the most represented nationalities in the model's output largely aligns across the two languages, although the Czech results contain substantially fewer countries than those in English.





Figure 4 indicates that the Czech data consists predominantly of male individuals associated with ice hockey (92%), with nationalities primarily from North America and Europe. In contrast, the English data includes both male and female individuals linked to ice hockey (80%), field hockey (17%) and para ice hockey (3%), with nationalities spanning most continents except Africa and Antarctica. The broader range of hockey types in the English data accounts for its wider geographic coverage. Interestingly, the Czech data contains more unique individuals, which is counter-intuitive given the greater hockey type variety represented in the English data (figure 5).

**FIGURE 4.** Distribution of hockey type for unique individuals per nationality and language.



**FIGURE 5.** Count of unique individuals per language.

Despite its limited scope, this experiment raises an important question: Does the LLM's output meaningfully reflect knowledge embedded in Czech- and English-speaking cultures, or does it produce a universalised response? I propose two interpretations. First, the English output encompasses a wider range of hockey variants and countries, a pattern not easily attributable to any single English-speaking culture but rather

<sup>14</sup>*The reasons for the selection of specific LLMs and languages as study subjects is not always documented. Atari et al. (2023) provide a compelling argument but do not provide supplementary materials. Others use LLMs to simulate survey respondents without the ability to ensure representative answers (Rystrøm et al. 2025; Schröder et al. 2025; Bravansky et al. 2025; Khan et al. 2025).*

to their amalgam. This breadth may flatten cultural specificity into a generalised response, supporting the earlier assertion that even cultures with high digital adoption and high language representation in training data may be vulnerable to misalignment. Second, the nationalities in the Czech output correspond intuitively to the historical development of hockey in the Czech Republic and former Czechoslovakia. This finding aligns with Rystrøm et al.'s (2025) suggestion that monocultural languages may more readily achieve stronger cultural alignment in LLMs.

A common objection to artistic and design research methods is their perceived lack of scientific rigor and replicability. Because LLMs are designed to generate variable outputs, obtaining identical results across interactions is inherently unlikely. Notably, several studies in my literature review—with presumably greater resources—have not yielded substantially more stable or reliable outcomes.<sup>14</sup> Even if relatively low-resource, artistic and design research methods can effectively direct attention to questions that merit further investigation through other scientific approaches.

## 5. PATH TOWARD CULTURALLY ALIGNED PERSONALISED AI

Improving cultural alignment is only one of several factors that can contribute to better personalised AI tools. The considerable challenges of developing culturally aligned large-scale LLMs should be evident by this stage. It is unlikely that a global AI system can fulfil promises of ultimate productivity and abundance. This assessment is grounded primarily in my professional experience as a designer and consultant in a range of international and Czech technology companies, which, in design terminology, may be understood as a form of ethnographic research.

Here are the most important issues that I have observed:

- **Oversimplification of complex problems.** As discussed earlier, oversimplification is not merely a design failure but a structural feature of the market-oriented product development. Projects designed for global scale are often attractive in boardrooms and strategic documents because they promise access to larger audience and greater revenue potential. However, scalability typically requires reducing complex problems to simplified forms, often prioritising a privileged subset of individuals. Despite efforts to adopt human-centred approaches, solutions are frequently retrofitted to poorly defined problems.
- **Prioritisation aligned to business goals.** In the best case, development priorities reflect clearly articulated organisational strategies. However, business objectives and key performance indicators (KPIs) rarely incorporate human well-being or externalities. Consequently, product teams often disproportionately optimise for meeting specific KPIs over meaningful real-world outcomes.

- Lack of quality or misunderstanding of design research. Participatory and human-centred design practices rely on research. However, practitioners often fail to distinguish between genuine needs and expressed preferences. Design research should inform well-grounded decisions, rather than serve as a source of convenient evidence.
- Insufficient coordination among policymakers, academics and industry. Many theoretical frameworks (e.g., Kubes 2025) and policy proposals based on human-centred AI and persuasive technologies overlook the practical constraints of product development, limiting their real-world applicability.

So how can we shape culturally aligned technology, especially the LLMs? The following considerations may help guide us toward a better path forward:

- **Train technologists in humanities.** In the Czech Republic, disciplines such as philosophy and ethics have long been undervalued and underfunded, despite their potential to equip technologists with critical tools for reflecting on the solutions they develop. Notably, only one Czech university design program includes an introductory philosophy course. By contrast, such courses are relatively common in engineering programs, although they tend to emphasise historical perspectives.
- **Re-centre human-centred product development on the human.** In addition to strengthening research capabilities to produce high-quality insights for sound decision-making, organisations should evaluate more than profit-driven metrics and measure their impact on people's lives (Vallor 2024; Monteiro 2019). One practical approach, proposed by the Center for Humane Technology, is the use of *anti-KPIs* to minimise harmful consequences by identifying failures to implement corrective measures. "For example, a KPI related to 'engagement' might be paired with an anti-KPI related to 'misinformation' to avoid breaking down reality in the name of growth" (Center for Humane Technology 2022, 41).
- **Foster exchange between academia and industry.** Scholars often remain within disciplinary and institutional boundaries, while industry practitioners rarely engage with academic research or involve scholars in their projects. Addressing complex issues such as cultural alignment requires collaboration across diverse perspectives.
- **Promote active civic participation.** Public deliberation is needed to determine what constitutes a desirable and sustainable society, including the value of cultural diversity. As participants in these technological systems, citizens should critically assess whether the problems being addressed justify the costs borne by individuals and society. Open Science can support this process by directly involving citizens in research.

A potential path forward lies in developing smaller, more sustainable AI systems tailored to specific use cases. Ensuring cultural alignment is more feasible for models designed for discreet populations. Initiatives such as Aya dataset (Cohere Labs 2025), an open science project involving 119 countries, or the recently launched European initiative OpenEuro LLM (2025) represent meaningful steps in this direction.

I will conclude with a quote from Mirca Madianou (2025, 178):

*We need to move beyond the binary thinking that we must choose between approaches that either favour powerful structures or human agency. It may seem obvious, but (some) academic fields—and popular discourse, in general—seem to forget that we don't actually have to choose. Structure and agency are co-dependent and cannot be understood in isolation.*

## REFERENCES

Albrecht, Kim. 2025. "Artificial Worldviews." In *xCoAx 2025: 13th Conference on Computation, Communication, Aesthetics & X*, 167–82. [https://doi.org/10.34626/2025\\_xcoax\\_006](https://doi.org/10.34626/2025_xcoax_006).

Atari, Mohammad, Mona Xue, Peter Park, Damián Blasi, and Joseph Henrich. 2023. "Which Humans?" Preprint, OSE, September 22. <https://doi.org/10.31234/osf.io/5b26t>.

Bannerman, Tey. 2025. "I've been part of three startup failures that burned through millions in funding." LinkedIn. [https://www.linkedin.com/posts/teybannerman\\_ive-been-part-of-three-startup-failures-activity-7348255813695680517-p7Ut/](https://www.linkedin.com/posts/teybannerman_ive-been-part-of-three-startup-failures-activity-7348255813695680517-p7Ut/).

Bender, Emily M., and Alex Hanna. 2025. *The AI Con: How to Fight Big Tech's Hype and Create the Future We Want*. New York: HarperCollins.

Birhane, Abeba, and Marek McGann. 2024. "Large Models of What? Mistaking Engineering Achievements for Human Linguistic Agency." *Language Sciences* 106 (November). <https://doi.org/10.1016/j.langsci.2024.101672>.

Blom, Jan. 2000. "Personalization: A Taxonomy." In *CHI '00 Extended Abstracts on Human Factors in Computing Systems*, 313–14. New York: Association for Computing Machinery. <https://doi.org/10.1145/633292.633483>.

Bravansky, Michal, Filip Trhlik, and Fazl Barez. 2025. "Rethinking AI Cultural Alignment." Preprint, version 2. arXiv, March 7. <https://doi.org/10.48550/arXiv.2501.07751>.

Buolamwini, Joy. 2023. *Unmasking AI: My Mission to Protect What Is Human in a World of Machines*. New York: Random House.

Buyl, Maarten, Alexander Rogiers, Sander Noels, Guillaume Bied, et al. 2025. "Large Language Models Reflect the Ideology of Their Creators." Preprint, arXiv, January 30. <https://doi.org/10.48550/arXiv.2410.18417>.

Center for Humane Technology. 2022. "Module 3: Minimizing Harmful Consequences." In *Minimizing Harmful Consequences*, Self-paced online course. Foundations of Humane Technology. Online. <https://www.humanetech.com/course>.

Cloarec, Julien. 2020. "The Personalization–Privacy Paradox in the Attention Economy." *Technological Forecasting and Social Change* 161: 120299. <https://doi.org/10.1016/j.techfore.2020.120299>.

Cohere Labs. 2025. "Aya." <https://cohere.com/research/aya>.

Common Crawl. 2025. "Statistics of Common Crawl Monthly Archives." <https://commoncrawl.github.io/cc-crawl-statistics/plots/languages>.

Couldry, Nick, and Ulises A. Mejias. 2019. "Data Colonialism: Rethinking Big Data's Relation to the Contemporary Subject." *Television & New Media* 20 (4): 336–49. <https://doi.org/10.1177/1527476418796632>.

Crawford, Kate. 2021. *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven: Yale University Press. <https://doi.org/10.2307/j.ctv1ghv45t>.

Curry, Derek, and Jennifer Gradecki. 2025. "Generative Persuasion." In *xCoAx 2025: 13th Conference on Computation, Communication, Aesthetics & X*, 326–30. [https://doi.org/10.34626/2025\\_xcoax\\_026](https://doi.org/10.34626/2025_xcoax_026).

Danaher, John, and Henrik Skaug Sætra. 2023. "Mechanisms of Techno-Moral Change: A Taxonomy and Overview." *Ethical Theory and Moral Practice* 26 (5): 763–84. <https://doi.org/10.1007/s10677-023-10397-x>.

Espeland, Wendy Nelson, and Michael Sauder. 2007. "Rankings and Reactivity: How Public Measures Recreate Social Worlds." *American Journal of Sociology* 113 (1): 1–40. <https://doi.org/10.1086/517897>.

Hämmerl, Katharina, Bjoern Deiseroth, Patrick Schramowski, Jindřich Libovický, et al. 2023. "Speaking Multiple Languages Affects the Moral Bias of Language Models." In *Findings of the Association for Computational Linguistics: ACL 2023*, edited by Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, 2137–56. Toronto: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-acl.134>.

Hao, Karen. 2025. *Empire of AI: Dreams and Nightmares in Sam Altman's OpenAI*. London: Penguin.

Harris, Tristan. 2025. "Is AI Productivity Worth Our Humanity?" Your Undivided Attention. June 26. Podcast. <https://centerforhumanetechnology.substack.com/p/is-ai-productivity-worth-our-humanity>.

Hernández-Ramírez, Rodrigo. 2019. "On False Augmented Agency and What Surveillance Capitalism and User-Centered Design Have to Do with It." *Journal of Science and Technology of the Arts* 11 (2): 18–27. <https://doi.org/10.7559/CITARJ.V11I2.667>.

Iványi-Bitter, Brigitta et al. 2026. "Epistemic Cultural Flattening in Generative Visual AI". *Disegno: Journal of Design Culture*

Khan, Ariba, Stephen Casper, and Dylan Hadfield-Menell. 2025. "Randomness, Not Representation: The Unreliability of Evaluating Cultural Alignment in LLMs." Preprint, arXiv, April 8. <https://doi.org/10.48550/arXiv.2503.08688>.

Kirk, Hannah Rose, Andrew M. Bean, Bertie Vidgen, Paul Röttger, and Scott A. Hale. 2023. "The Past, Present and Better Future of Feedback Learning in Large Language Models for Subjective Human Preferences and Values." In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, edited by Houda Bouamor, Juan Pino, and Kalika Bali, 2409–30. Singapore: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.emnlp-main.148>.

Kotliar, Dan M. 2020. "Data Orientalism: On the Algorithmic Construction of the Non-Western Other." *Theory and Society* 49 (5): 919–39. <https://doi.org/10.1007/s11186-020-09404-2>.

Kubes, Tanja. 2025. "A Critique of Human-Centered AI: A Plea for a Feminist AI Framework (FAIF)." *AI & SOCIETY*. August 19. <https://doi.org/10.1007/s00146-025-02556-8>.

Kučera, Dalibor, and Matthias R. Mehl. 2022. "Beyond English: Considering Language and Culture in Psychological Text Analysis." *Frontiers in Psychology* 13 (March). <https://doi.org/10.3389/fpsyg.2022.819543>.

Kwet, Michael. 2019. "Digital Colonialism: US Empire and the New Imperialism in the Global South." *Race & Class* 60 (4): 3–26. <https://doi.org/10.1177/0306396818823172>.

Libovický, Jindřich, Rudolf Rosa, and Alexander Fraser. 2019. "How Language-Neutral Is Multilingual BERT?" Preprint, arXiv, November 8. <https://doi.org/10.48550/arXiv.1911.03310>.

Libovický, Jindřich, Rudolf Rosa, and Alexander Fraser. 2020. "On the Language Neutrality of Pre-Trained Multilingual Representations." Preprint, arXiv, September 29. <https://doi.org/10.48550/arXiv.2004.05160>.

Liu, Leigh Anne, and Claudia Dale. 2009. *Using Mental Models to Study Cross-Cultural Interactions*. [https://doi.org/10.1057/9780230240834\\_11](https://doi.org/10.1057/9780230240834_11).

Madianou, Mirca. 2025. *Technocolonialism: When Technology for Good Is Harmful*. Cambridge: Polity.

Meyer, Erin. 2014. *The Culture Map: Breaking Through the Invisible Boundaries of Global Business*. New York: PublicAffairs.

Monteiro, Mike. 2019. *Ruined by Design: How Designers Destroyed the World, and What We Can Do to Fix It*. Mule Design. <https://www.ruinedby.design>.

Narayanan, Arvind, and Sayash Kapoor. 2024. *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference*. Princeton: Princeton University Press.

OpenEuroLLM. 2025. "OpenEuroLLM." <https://openeurollm.eu/>.

Perez, Jorge. 2025. "Tokenising Culture: Causes and Consequences of Cultural Misalignment in Large Language Models." *Ada Lovelace Institute*, June 19. <https://www.adalovelaceinstitute.org/blog/cultural-misalignment-llms/>.

Piorecký, Karel, and Zuzana Husárová. 2024. *The Culture of Neural Networks Synthetic Literature and Art in (Not Only) the Czech and Slovak Context*. Prague: Karolinum. <https://doi.org/10.14712/9788024657844>.

Qiu, Jack Linchuan. 2016. *Goodbye iSlave: A Manifesto for Digital Abolition*. Chicago: University of Illinois. <https://doi.org/10.5406/illinois/9780252040627.001.0001>.

Raines vs. OpenAI. 2025. Superior Court of the State of California for the Country of San Francisco August 26, 2025. <https://www.documentcloud.org/documents/26078538-raine-vs-openai/>.

Raskin, Aza. 2025. "How OpenAI's ChatGPT Guided a Teen to His Death." Your Undivided Attention. September 11. Podcast. <https://centerforhumanetechnology.substack.com/p/how-openais-chatgpt-guided-a-teen>.

Rosa, Rudolf, David Mareček, Tomáš Musil, Michal Chudoba, and Jakub Landsperský. 2025. "EduPo: Progress and Challenges of Automated Analysis and Generation of Czech Poetry." In *Proceedings of the 5th International Conference on Natural Language Processing for Digital Humanities*, edited by Mika Hämäläinen, Emily Öhman, Yuri Bizzoni, So Miyagawa, and Khalid Alnajjar, 524–42. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.nlp4dh-1.45>.

Rosa, Rudolf, Patrícia Schmidtová, Alisa Zakhtarenko, et al. 2022. "THEaiTRobot: An Interactive Tool for Generating Theatre Play Scripts." In *Proceedings of the 15th International Conference on Natural Language Generation: System Demonstrations*, edited by Samira Shaikh, Thiago Ferreira, and Amanda Stent, 10–13. Association for Computational Linguistics. <https://aclanthology.org/2022.inlg-demos.4/>.

Rothman, Joshua. 2025. "A.I. Is Coming for Culture." *The New Yorker*. August 25. <https://www.newyorker.com/magazine/2025/09/01/ai-is-coming-for-culture>.

Rystrøm, Jonathan, Hannah Rose Kirk, and Scott Hale. 2025. "Multilingual != Multicultural: Evaluating Gaps Between Multilingual Capabilities and Cultural Alignment in LLMs." Preprint, arXiv, August 30. <https://doi.org/10.48550/arXiv.2502.16534>.

Schneider, Nathan. 2022. "Governable Stacks against Digital Colonialism." *Journal for a Global Sustainable Information Society* 20 (1): 19–36.

Schröder, Sarah, Thekla Morgenroth, Ulrike Kuhl, Valerie Vaquet, and Benjamin Paaßen. 2025. "Large Language Models Do Not Simulate Human Psychology." Preprint, arXiv, August 13. <https://doi.org/10.48550/arXiv.2508.06950>.

Strathern, Marilyn. 1996. "From Improvement to Enhancement: An Anthropological Comment on the Audit Culture." *Cambridge Anthropology* 19 (3): 1–21.

Vallor, Shannon. 2024. *The AI Mirror: How to Reclaim Our Humanity in an Age of Machine Thinking*. Oxford: Oxford University Press.

W3Techs. 2025. "Usage Statistics and Market Share of Content Languages for Websites." *November*. [https://w3techs.com/technologies/overview/content\\_language](https://w3techs.com/technologies/overview/content_language).

Yates, J. Frank, and Stephanie de Oliveira. 2016. "Culture and Decision Making." *Organizational Behavior and Human Decision Processes* 136 (September): 106–18. <https://doi.org/10.1016/j.obhdp.2016.05.003>.

Yeung, Karen. 2019a. "Algorithmic Regulation: An Introduction." In *Algorithmic Regulation*, edited by Karen Yeung and Martin Lodge, 1–18. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780198838494.003.0001>.

Yeung, Karen. 2019b. "Why Worry about Decision-Making by Machine?" In *Algorithmic Regulation*, edited by Karen Yeung and Martin Lodge, 21–48. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780198838494.003.0002>.

Zuboff, Shoshana. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs.