Anaëlle Beignon\* a.beignon@unistra.fr Université de Strasbourg Strasbourg, France

Missing number, treated as zero.

Thomas Thibault\* contact@thomas-thibault.fr Designers Éthiques Paris, France Nolwenn Maudet\* nmaudet@unistra.fr Université de Strasbourg Strasbourg, France



## Figure 1: Overleaf offering us to fix this paper with AI.

### Abstract

Generative AI is being massively deployed in digital services, at a scale that will result in significant environmental harm. We document how tech companies are transforming established user interfaces to impose AI use and show how and to what extent these strategies fit within established deceptive pattern categories. We identify two main design strategies that are implemented to impose AI use in both personal and professional contexts: imposing AI features in interfaces at the expense of existing non-AI features and promoting narratives about AI that make it harder to resist using it. We discuss opportunities for regulating the imposed adoption of AI features, which would inevitably lead to negative environmental effects.

## Keywords

AI, Deceptive Patterns, Design, Environmental Effects, Regulation

## 1 Introduction

"Computing is evolving again. We spoke last year about this important shift in computing from a mobile-first to an AI-first approach. ... In an AI-first world, we are rethinking all our products and applying machine learning and AI to solve user problems." Sundar Pichai keynote in 2017<sup>1</sup>

We are currently witnessing a paradigm shift in the software industry, towards an AI-first mindset. Not only do many new companies develop AI-based products, but many existing tech companies are integrating AI-based features into their existing products.

This massive turn towards generative AI has worrying social and environmental effects that researchers are just starting to estimate. Studying the environmental impact of an AI service based on Stable Diffusion, Berthelot et al. explain that "with 360 tons of carbon equivalent emission, an impact on metal scarcity equivalent to the production of 5659 smartphones, and an energy footprint of 2.48 Gigawatt hours, it is clear that the impact of Gen-AI should be a matter of concern and not only for its carbon footprint." [7]. In addition

<sup>1</sup>https://www.youtube.com/live/Y2VF8tmLFHw?t=411s

to carbon emissions and mineral use, the massive development of AI requires large quantities of water, conflicting with human use and putting a strain on access to this key resource [24]. Worse, AI is on a path to undermine all optimization efforts and instead worsen the environmental impact of ICT. During the first semester of 2023, North America set a record with a 25% increase in the construction of new data centers<sup>2</sup> and a 2024 report by the International Energy Agency estimated that electricity usage by data centers *"is set to more than double to around 945 TWh by 2030"*[1].

Following the cornucopian paradigm described by Preist et al. [36], tech companies are currently anticipating a growth in demand for their gen-AI products. Accordingly, they have started reshaping and building new infrastructure (data-centers, energy transport and production, etc.) to support this future expected demand, in the process worsening existing environmental issues. The concern for winning the "AI war" has become a national one: for example, the US plans to extend the lifespan of coal power plants to support the increasing energy demand due to AI [15]. AI use has thus been confidently anticipated.

But will users actually adopt AI-based features? This is not yet guaranteed. Generative AI is beset by various controversies that could limit people's willingness to adopt such features or pay for them [5]. The fallibility of gen-AI has long been criticized, and is even debated in mainstream media outlets: a BBC study has for instance found that AI assistants give inaccurate responses more than half of the time <sup>3</sup>. Moreover, many activist-led campaigns have sought to raise awareness about the social and environmental impacts of AI.<sup>4</sup> More broadly speaking, trust issues regarding AI drive "algorithm aversion" [12] or "human favoritism" [43] and lower its appeal to potential users [12].

To push for adoption, tech companies have been investing in large-scale marketing efforts, backed by extensive media coverage in which AI products and features are often presented as revolutionary. More surprisingly, in an arguably unprecedented development at this scale, companies have also been leveraging UX and UI design strategies to promote the adoption of AI-based features. In early 2024, multiple AI buttons began appearing into our daily digital

<sup>\*</sup>All authors contributed equally to this research.

LIMITS '25, 11th Workshop on Computing Within Limits, June 26-27, 2025, Online 2025.

 $<sup>^2</sup> https://www.business$ insider.com/ai-data-centers-land-grab-google-meta-openai-amazon-2023-12

 $<sup>\</sup>label{eq:asymptotic} ^3 https://www.bbc.co.uk/aboutthebbc/documents/bbc-research-into-ai-assistants.pdf \ ^4 see for example https://savethe.ai/, https://www.bbc.com/news/articles/cwyd3r62kp50$ 

tools, in a particularly conspicuous manner compared to other new features.

Our initial goal with this work was to document and analyse the design strategies employed by tech companies in order to bring about the announced *AI revolution* [30]. While new digital services were designed as components of this "revolution", we chose to focus on the strategies that aim at favoring user engagement with AIbased features on digital platforms that were initially not conceived as AI-first products.

Having started to collect extensive anecdotal evidence on the reluctance of users feeling pressured to use AI-based features in their daily tools, we also want to understand to what extent these design strategies fall within existing deceptive design patterns categories. In doing so, we could build on Tocze et al.'s work on unsustainable patterns [39] by describing the ways in which design patterns can favor or force the adoption of digital uses that are environmentally unsustainable, for the benefit of companies. Our hope for this work is that it will spark debate on unsustainable use adoption and on the opportunities and challenges involved in the regulation of the forced adoption of unsustainable digital products and services.

## 2 Background

While Artificial Intelligence (AI) has always been a subject of debate, recent developments have stoked controversies surrounding biases, misinformation and privacy issues [5]. Our work builds upon research on the environmental effects of AI and explores the challenges involved in exposing the role of design in intensifying demand and related indirect effects.

## 2.1 Environmental effects of generative AI

Research have been denouncing the environmental costs of largescale AI development for years [32, 37, 38]. The impact extends beyond climate, affecting 6 out of 9 planetary boundaries[18]; additionally the environmental benefits of *AI for green* solutions are now called into question [25, 32]. Yet, generative AI services remain based on energy-intensive architectures [26]. Also, AI systems and products are produced in a context that favors constant scaling up to match increasing computational possibilities. Even in AI research, a "*bigger-is-better*" narrative perpetuates the assumption that the bigger the AI system, the better its performance and value[40].

A recent study led between September 2023 and August 2024 on the electricity consumption of data centers found they account for over than 4% of the US's electricity consumption, with 56% of these originating from fossil fuels[20]. The authors also project that the *AI race* will play a role in the doubling of the worldwide energy consumption of data centers, which would outgrow the power used in Canada (the sixth highest energy-consuming country in the world) by 2026. As they put it, *"the AI race significantly increases computational demands and, consequently, energy consumption of data centers"*[20]. While little data exists on the use phase of large models, Varoquaux et al. [40] have estimated that the energy use due to inference (i.e. the application phase of AI) would outgrow the training phase amount of energy use only in a few weeks in the case of ChatGPT. Greenhouse gas emissions due to inference are slated to worsen the already alarming impact of the biggest tech companies: in 2023, Google's emissions increased by 48% compared to 2019 and Microsoft's increased by 29% compared to 2020[20].

Luccioni et al. [27] point to the risks of rebound effects if AI development keeps on following an exponential growth trend. They also note that while the power usage effectiveness (PUE) of AI training data centers is improving, their global energy use is on the rise. This rebound effect is tightly connected with increased AI use due to an expanding market, the impact of which cannot yet be fully assessed. This relates to the cornucopian paradigm [36], which describes how demand is stimulated by the design of digital services, which in turn drive the expansion of the digital infrastructure, e.g. manufacturing more GPU chips and building new hyperscale data-centers. In this paper, we seek to document and analyse the strategies deployed by companies to stimulate AI adoption and ultimately lead to rebound effects.

## 2.2 Deceptive Patterns

As gen-AI is a controversial technology, the question of user agency regarding its adoption is important. Since 2010, dark, deceptive and manipulative design patterns have been developed, and have ended up "limit[ing] user autonomy and decision making" [19].

Ibrahim et al. [21] have explored both the individual and social implications of the design of AI interface features. They discuss the **"cascading impacts"** of AI use when deceptive design choices lead to addictive behaviors or encourage the spread of misinformation. Our study builds on this work by focusing specifically on the design strategies used to push for AI adoption, to connect interface feature design with environmental concerns.

While we argue that certain design choices are harmful, we do not argue that they are deliberately harmful. In this sense, we align with Di Geronimo et al.'s [16] approach of deceptive patterns which analyses when an "interface seem[s] to benefit the app rather than the use". We focus on interface choices that have benefited generative AI use, understood here as an environmentally problematic technology, and identify design strategies that have not been conducive to users' autonomy, freedom of choice and understanding. We adopt a collective welfare lens [31] to study deceptive design patterns, meaning that we focus less on individual consumer rights and more on socio-technical transformations liable to deteriorate living conditions on earth.

A previous LIMIT publication [39] coined the term "unsustainable design patterns" which describe how the arrangement of different components (e.g. "artifact", "intent", "participant", "context") results in unsustainable shortcomings. Our study focuses on the "artifact" as a means to highlight potential unsustainable design patterns. Our work expands in particular on "planned rebound effect" pattern, whereby platform providers would be "pretending to ignore what could be a consequence of launching a new service". Our analysis uncovers how this pretending strategy is embedded in interface choices in the case of AI-based features in 2024. We link this analysis to previous work using Gray et al's [19] ontology of patterns, which aims for a "shared language" between legislative and academic communities. We reflect and speculate on how design strategies imposing AI use could be denounced and eventually regulated.

LIMITS '25, 11th Workshop on Computing Within Limits, June 26-27, 2025, Online

## 3 Method

In early 2024, we started noticing that many AI icons were appearing in the software and applications we use. From September 2024 onwards, we started, as a team of 3 HCI researchers, to systematically collect screenshots and descriptions of each new AI-related UI or UX change we encountered. We also collected and documented cases that we came across on social media.

To diversify our corpus, we first asked for input from colleagues, family and friends through word of mouth. We also launched open calls on Mastodon via the Limites Numériques account, which has over 1,000 followers<sup>5</sup>. Over the course of a 4 month period, we collected **90** screenshots and descriptions of recent AI related changes in interfaces in **53** different software and app interfaces.

We only collected examples explicitly labeled or presented as "AI". We also only focused on software and application that previously operated without generative AI-based features. We should note that our corpus is not representative of the state of global software usage and reflects our bias towards the applications we use on a daily basis, which were mainly developed in the US and to a lesser extent in Europe. However, in this study we focus on identifying and analysing the strategies used to favor AI and not on such representativeness issues, which is why we do not provide counts or percentages when describing our findings.

Our corpus includes many of the biggest software companies, including Google, Apple, Meta, Microsoft and Adobe. We found examples in both software geared towards professionals such as Slack or Adobe Photoshop, but also numerous examples of leisure, entertainment or social media applications such as Spotify, Microsoft Paint, Instagram or Snapchat. To put our corpus in perspective, we complemented our analysis with written sources such as press releases and ads targeting a general audience but also more technical communications (companies' blog posts...) intended for designers and developers.

We constructed our method of analysis in two steps. The three authors first collectively analysed the corpus using a clustering technique. We grouped examples that were using the same strategies together and coded these strategies with tags such as "takes up space", "magic", "unremovable", etc.). As strategies started to consolidate, we went back to the corpus and tagged all the corpus items with these strategies to make sure that we were able to reliably identify which strategies were being used. From this analysis, we observed recurring strategies in terms of wording, information architecture and graphic design choices, as well as interaction patterns.

Based on this initial analysis, we wanted to verify whether and to what extent these recurring patterns were consistent with the ontology of existing deceptive patterns. We chose to use Gray's comprehensive ontology[19] to map the specificities of the patterns we encountered.

## 4 Results

Based on our analysis, we evidenced two types of strategies that rely on a combination of deceptive patterns. We first discuss the visual and interaction design strategies that affect the expected affordances (4.1), information architecture and user flows in order

to favor AI adoption, generally at the expense of existing digital practices on the application or platform. We then turn to a second more insidious strategy that relies on the manipulation of narratives around technology use through interface design choices (4.2). Relving on the analysis of interface elements, we used Gray et al's ontology of deceptive patterns [19] to situate identified strategies in relation to known deceptive patterns. Our analysis highlights the use of two specific meso-level deceptive patterns 2: "Manipulating visual architecture" and "Emotional and Sensory Manipulation". In our case, however, it is the intersection and combination of manipulative tactics across many widely used platforms that really gives us insights on how AI use is imposed to users. These strategies also all align with the "planned rebound effect" pattern from Toczé et al. [39], as the wide scale adoption of AI technology will increase the ICT's sector footprint, as well as with "Environmental blindness" and "concealed impacts", as the strategies employed are making AI's environmental impact less visible.

# 4.1 Imposing AI features at the expense of existing non-AI features

As Janlert and Stolterman have shown [22], visual interfaces are finite space, i.e. a bottleneck, which means designers have to make choices about what features to display and how much space can be devoted to each of them. In our analysis we found that, in comparison with non-AI features, AI-based features tend to occupy a place of honor and to call for users' attention in many different ways. AI-based features are positioned and presented in a way that maximizes their use, regardless of whether this benefits the user experience.

4.1.1 **Visual prominence of AI features**. AI features tend to take up a lot of space and are often given the most valuable space in interfaces, toolbars and menus. Linkedin's interface, for example, features a messaging popover and a banner advertising an AI-based feature that occupy more than half of the space. In Notion, AI features take up nearly a third of the toolbar (Figure 4). In both Notion and DeepL, whenever users select words, a floating toolbar appears in which the buttons closest to the cursor are AI features. This indicates that they are treated as the most important features.

Interfaces often feature one salient call to action and these are also being used to encourage the use of AI-based features. In Google Keep, a floating action button, which is the component that helps people take primary actions, displays an AI-powered feature (Figure 3). Similarly, Google's Gemini is embedded by default on all Galaxy S25 devices and can be activated by a long press on the power button.

AI features stand out visually as well as spatially. In numerous cases (Figure 5), AI-based features are highlighted using a distinctive color, a colorful gradient, or even using an animated icon (Notion) in contrast with all the other features, which are usually displayed using static black and white icons. We also found that AI features are given multiple access paths and that users are constantly reminded about these features. In Acrobat Reader, for example, AI features are presented and suggested to users on multiple occasions (Figure 6). The same "AI assistant" feature is presented in both left and right toolbars, as a button in the menu bar, as a tooltip in the interface, as a GUI prompt and also when users select text. Similarly,

<sup>&</sup>lt;sup>5</sup>https://mastodon.design/@limitesnumeriques/113431042717320191

Design strategies for Al adoption					
Imposing AI features at the expense of existing non-AI features	Establishing misleading technology narratives				
Description of the design strategies					
The features are positioned and presented in a way that maximizes their use, regardless the benefit for user experience. 	The design conveys a <b>positively connoted</b> <b>meaning</b> that conflicts with the understanding of the technical functioning of the features and the related effects.				
Benefit to plat	tform providers				
<ul> <li>Stimulating user adoption to justify a price upgrade.</li> <li>Exhibiting AI features for credibility regarding innovation standards</li> </ul>	<ul> <li>Concealing technical weaknesses of gen-Al.</li> <li>Positive framing covering up the degradation of UX</li> <li>Explaining the iterative process to arrive at the expected results</li> </ul>				
Deceptive design p (Gray et	patterns categories al, 2024 )				
Manipulating visual choice architecture; Visual Prominence; Nagging; False hierarchy; Visual Prominence; Bad default; Interface interference;	Emotional or sensory manipulation; Interface interference; Social Engineering; Positive or Negative Framing; Cuteness				

## Figure 2: Summary table of our two strategies and the related deceptive patterns

Obstruction

Whatsapp has updated its applications to integrate even more AIbased features (Figure 8). In addition to the addition of a "Meta AI" chat feature, the search bar originally dedicated to retrieving contacts and conversations now also redirects to Meta's AI assistant.

This visual prominence of these examples of AI features contrasts with the discretion that is exercised when it comes to how user data is being used to train AI models on many applications. If the option to disable such personal data use exists, its existence is not notified to users and the dedicated parameter is very often hidden among others (Fig: 7).

These design strategies are a good match for Gray et al.'s [19] 'Visual Prominence' pattern, which places "an element relevant to user goals in visual competition with a more distracting and prominent element". This omnipresence is even reinforced with strategies consistent with the pattern "nagging" [19], where AI features are



## Figure 3: Floating button for an list generation feature on Google Keep



Figure 4: AI-related features are the first on Notion's toolbar takes up a the third of the space



Figure 5: Icons of AI-based features, enhanced through graphic design.

presented repeatedly and interrupt users in their actions. For example, when opening Adobe Photoshop, users are first welcomed

#### Beignon, Thibault and Maudet

LIMITS '25, 11th Workshop on Computing Within Limits, June 26-27, 2025, Online



Figure 6: Adobe acrobat reader interface, where AI features are redundant.



Figure 7: The path for deactivating personal data training on X involves 5 steps



Figure 8: WhatsApp interface: the search bar has deviated from its initial purpose and can now be used for prompt inputs.

with a popup encouraging them to "explore the power of generative AI". After closing this initial popup, a tooltip opens in the interface to suggest the same feature. Similarly, when opening Microsoft Skype, users were often greeted by its AI-chatbot Copilot, which invited users to use it as soon as the interface opened, regardless of whether they wanted to use that feature or not. These strategies may induce users who were originally uninterested in using AI to click on the features out of curiosity or by mistake. This pervasive advertising of AI also leverages the perception that AI use is an unavoidable technological turn.

4.1.2 AI features are privileged on the interface. We noticed that in some cases, the strategies go beyond making AI features visually prominent. We found examples of AI features that are presented in the same fashion as other features or other interface elements through their visual organization and form. AI assistant features are implemented on messaging apps in the form of a conversation, meaning that assistants use the same UI elements as regular discussions with contacts (Figure 9). ). On Snapchat, on the other hand, the AI assistant, unlike all the other conversations, always remains on top and cannot be removed, regardless of whether it is actually used or not. While other conversations are gradually dropped down if they are never updated, AI assistants are not subject to the same treatment and limit the screen space dedicated to active conversations. The hierarchy rules that apply to similar UI elements do not apply to AI features, in a way that advantages the visibility and the interaction with the feature. Similarly, as noted earlier, AI features are sometimes highlighted through the presence of by multiple buttons on the screen, where the other features generally only have one (Figure 6).

These examples match Gray et al.'s [19] "False Hierarchy" pattern where the strategy is to "give one or more options visual or interactive prominence over others, particularly where items should be in parallel rather than hierarchical". This manipulation of visual choice architecture influences the choice of features to interact with: AI-based features are discretely favored among others using UI/UX design.

4.1.3 Interfering with non Al uses. We found that the ways AI features are pushed on interface users also interfere with the use of non-AI features and interrupt users' usual workflows. Many tools advertise their AI features using banners in the interface, including Slack and Google Docs. The tooltips and banners promoting AI in interfaces generally require at least a click from users to dismiss them. Moreover, the preeminence given to AI features inevitably leads to people triggering AI features, sometimes accidentally. For example, in Notion, it is very easy to launch the AI assistant inadvertently, since pressing the extremely frequently used space bar will open it. In comparison, the keyboard shortcut for any other command requires users to press '/', which is less conducive to this type of error. A similar hijacking of user experience on the platform happens when actions routinely performed before the AI feature update take more time and become more fastidious. On Qwant, the results of a web search engine have been downgraded, now appearing below a huge box of AI-generated responses, requiring users to scroll longer.

These interferences with the use of non-AI features present significant similarities with Gray et al.'s deceptive pattern *Adding Steps*. Even if it has so far mostly been used to describe privacy related issues, *Adding Steps* to perform a given online task is a deceptive pattern that is arguably very relevant in contexts in which a company has an interest in influencing users' adoption of features. In the context of AI, this type of obstruction strategy disrupts users' habits on the platforms. Users who do not want to use AI will have to tolerate this friction on their existing digital habits, or to look for a way to disable -if possible- the feature that is bothering them.

<b>Q</b>	Chat	0 C+		G	Messages	Q	2
Tout Non lu 🕦	Groupes	Répondre	En∨	A	N. S. MARKEN AMERICAN	need -	1 min
My Al 😫 Capture d'écra	an - 76sem		0	+	Gemini Boniour, je sujs Gemini d	tans Goo	14:06
<u>a</u>	9min · 46	5 🆤	Q	R			ven.
👷 Kalendor.	- 4 🎔		0		ene Succession de la composition	and the	
	207 🎔		0	F	an bisana wa		ven.
🧶 🐜 👘	413 🖤		Ø	8	orte Print Carlotte Carl Scher Carlotte Malatte Carlotte	es.	jeu.
S Landing of	3h		Ø	0	in an Air An China ann a <mark>199</mark> 1	ar da a	jeu.
			0	8	Sector States	· • • • •	jeu.
🧿 🐂 2j			\$	8	apida apiday 🖓 Démarrer	une discussio	on
in the second se					06 13 56 62 79		atar.

Figure 9: AI assistants on Snapchat and Messages, displayed as contacts, regardless of whether users interact with the assistant.



Figure 10: Qwant search results.

4.1.4 **Imposing AI use by default**. As we have seen, AI features are extremely easy to trigger, but they are also sometimes enabled by default or deployed before users are allowed to choose to deactivate them.

For example, the sport tracking application Strava introduced AI comments on activities without offering users the option to turn them off. YouTube also started to automatically dub videos. Users can turn it off in each individual video but to my knowledge there is no option to disable it entirely. In Notion, an "AI Summary" property is featured by default when creating a new database, whereas many other properties are not (Figure 11). The AI property requires explicit action from users looking to delete it. Unlike all the other properties, which can be deleted using the "delete property" menu option, the AI summary property can only be dismissed using the "ignore property" menu option. Even when applications provide ways to disable AI-based features, they involve friction. This may first be a result of the wording, as refusing to use AI based feature can rarely be a firm no. Instead, applications often offer to turn off the feature temporarily, using words such as "ignore for the Beignon, Thibault and Maudet

moment" or "maybe later" (Fig: 12). What this tells users is that they will certainly come to use it in the long run.

Al summary IA	🐅 Essayer sur cette page	
➡ Ajouter une propriété	5	
🗄 Al summary 🛛 🗛	🐅 Essayer sur cette page	4

## Figure 11: Notion database implementing a AI summary field by default

The default activation of AI features, making it much harder for users to deactivate or altogether remove from the interface, matches Gray's pattern category "*Bad Defaults*" [19], where default settings are not in the users' best interest; in this case, these default settings also go against sustainable digital practices, which is arguably against the long-term best interests of users (and non-users).



Figure 12: On Google Play, turning off an AI-based feature is not a permanent decision.

# 4.2 Establishing misleading technology narratives

Having documented a first strategy to push AI-based features that consists in making them stand out on the interface and/or interfere with non-AI user experience, in the following we expand on two technology narratives that are distinct in form but have similar effects, related respectively to the covering up of technical weaknesses and the controversies around artificial intelligence.

4.2.1 **Magification of AI technicality**. First, AI tends to be presented as a form of magic. Lupetti et al. [29] for example, have shown how the magic metaphor can be used as a way to control the social meanings embedded in technology use:

"Designing and communicating AI things as supernatural-enchanted-products, in fact, shapes the social perceptions of these systems, taking them out of the realm of mere technical tools to be regarded as socially capable agents [...] and/or socially valuable applications"

In user interfaces as well, AI is represented as magic to favour its adoption. The most common graphic symbol for representing an AI-based feature is the spark icon. Generally associated with something 'special', exciting, new, but also with innovation and wonder, the icon contributes to a representation of an intrinsically 'good' functionality. Unlike other icons that literally or metaphorically illustrate features, the icon used for AI sets it apart from other interface functionalities. Likewise, mauve-centric gradients, commonly associated with magic and the impalpable, are predominantly used for AI.



Figure 13: Qwant's home page in April 2024 (on the left) and November 2024 (on the right). On the right screen, a baseline introduces the new AI features: "Enjoy instantaneous answers with Qwant's artificial intelligence" (Translated from French)

Qwant offers by default what it calls an 'instant' response to searches thanks to a LLM feature. Ironically, this 'instant' response actually appears long after the web results. The words and metaphors used here make the machine's calculations invisible, concealing the fact that it is much slower than a traditional search engine and much more resource-intensive . We don't know how something magical should behave: is it normal for it to take time? Have I done something wrong? Is its consumption of resources justified? The invisibility of the machine enabled by the metaphor of magic puts people in a different relationship to the tool, in a way that diverts attention from its weaknesses - here, the increased response delay. It is no longer possible to pay attention to the machine. While metaphors are useful to the wide-scale appropriation of complex and new technologies, previous work has shown that digital metaphors tend to have an impact on the technological culture, and in turn normalize digital behaviors associated with a limitless, non-impacted, "immaterial" digital infrastructure [9, 10].

By revising their graphic identity using the visual semantics of magic (Figure 13), companies position AI as an all-purpose tool, without ever saying what it does not do. They benefit from setting such vague expectations among users, as it is more difficult to criticize the results of an action that had no precise goal in the first place, in the same way that surprise is the expected outcome of a magic trick. The promised versatility of AI tools, embodied in magical evocations, also results in the invisibilization of the materiality of the many operations involved in AI, spreading the idea that AI necessarily makes things easier, better and faster.

Interestingly, both magic and deceptive patterns can be closely connected to trickery. The language of magic conjures a sense of enchantment, i.e., *"the experience of being caught up and carried away"* [6], where knowing what happens behind the curtain is not only unnecessary, but also detrimental to the enjoyment of the full *magical* experience. Such metaphors are practical because they make it possible to sidestep questions of efficiency and environmental or political effects, framing an opaque process of data generation into a magical event.

4.2.2 **The dedicated assistant**. Another recurring form taken by generative AI in our corpus is the assistant.

The assistants offer to help people 'learn new ways, plan events, write thank you notes, and more' (Gemini on Google), 'chat, create and find anything' (Notion), or 'suggest unique ideas' (Skype). AI assistants are embodied using metaphors that evoke human characteristics, for example given names like 'Aria' (Opera) or 'Leo' (Brave), and are also assessed through the same UI elements as human collaborators or contacts (e.g. Figure 9). The tradition of attributing human characteristics to machines and robots has already been discussed in 1986 by Caporael [11], who unpacked the cognitive biaises behind the assumption that a machine with human attributes will perform better. Deploying particular metaphors of human roles or human characteristics leverages "previously known mental schemas for social relationships"[34], meaning that it shapes users' expectations of technology and thus the way they will approach it. In our case, presenting AI as having human characteristics or skills contributes to the discourse presenting AI as desirable and high-performing. This emphasis on versatile human skills also helps to justify the introduction of AI features into our digital services: it can do everything, so it can be everywhere.

Generative AIs' names often evoke a role in the professional environment, such as "Copilot" (Microsoft), "Agentforce" (Slack), "Cocreator" (Paint), "Companion" (Microsoft, Zoom), "analyst" (Google Chat) and, of course, "Assistant" (Protonmail, Brave, Webex). These names evoke collaboration among peers, sometimes represented by a personified emoji (Notion). They also echo the fantasy of alleviating demanding tasks through technology, a narrative already pushed in Apple's 1987 keynote [3]: thanks to the AI assistant, uninteresting or difficult tasks can be delegated, allowing users to focus on more important duties. The sentences that present the integration of AI assistants generally give an impression of availability and dedication (ex: "Whenever you need me" (Notion), "Need help?" (Brave).

AI assistants are presented as polyvalent subordinates, available for any task that can be formulated through a textual prompt. AIs are positioned under us in the hierarchy in a context of fear of labor degradation for skilled workers. Its polite and "discreet"[6] demeanor also reassures users about the safety of their personal data. Such *"undue trust"* in a robot is a deceptive pattern referred by Gray et al. as *Cuteness*. It conveys the appealing figure of an agent that combines the properties of a tool and of a teammate, leaning towards the *"agent-as-a-teammate paradigm"*[34], where *"we want tireless agents that help us while treating us like teammates, even if we don't treat them like teammates"*.

This second strategy is distinct as it is part of global narratives that have been disseminated beyond platforms (e.g. in commercials, on promotional websites) for decades [34]. Among other effects, these metaphors contribute to framing a competitive climate where AI use equals professional performance [35]. In an environment where AI is presented as unbiased, more efficient and unavoidable, managers and employees might be more prone to adopt useless AI-based features out of a fear of appearing outdated or incompetent. This impact at the social scale links to *Social engineering* deceptive patterns [19], where design choices influence perceived norms around technology use. One important shared trait of the magic and assistant metaphors is that they disseminate unrealistic expectations of technology to favor its blanket adoption in multiple sectors.

Both of the metaphors align well with Gray et al.'s design pattern Emotional or sensory manipulation [19], where design choices are used to "evoke an emotion or manipulate the senses in order to persuade the user into a particular action". Natale has shown how analogies used in the scientific literature contribute to shaping the "AI myth" [33]; likewise, design choices for AI features contribute to the construction of technological myths. In the case of AI, these narratives are strongly connected to ongoing debates on the weaknesses and concerns around artificial intelligence. For example, both metaphors are intended to mitigate "algorithm aversion" [12], i.e. the fact that users value more human abilities more than AI's. Indeed, magic blurs and positively frames the outcomes of AI-mediated operations, while the assistant metaphor suggests a quality equivalent to human-produced content. Conversely and more surprisingly, these metaphors also put users in a position to expect inaccurate results from AI features. The magic metaphor explains approximations in AI responses and encourages users to make other tries by adding some tweaks until they find the right trick. Comparably, the assistant has to be briefed *correctly* to function, meaning the responsibility of AI interaction's outcome falls on users, who need to adjust their prompts to steer the assistant in the right direction. In both ways, the positive framing of these narratives is combined with an indication that the AI feature might not immediately work as we would like it to.

## 5 Discussion

While the design strategies used to favour AI adoption do fit within the existing deceptive pattern categories described by Gray et al.[19], Brignul's initial overall definition of deceptive patterns ("tricks used in websites and apps that make you do things that you didn't mean to, like buying or signing up for something"<sup>6</sup>) falls short in our case. Here the patterns are not deceptive per se: users know that they are using AI features and no immediate damage occurs, like an unintended purchase. The issue here is that use is provoked by leveraging powe dynamics that are beyond any user's agency. In the case of AI's deceptive patterns, the tricks are also performed on another temporality and on another scale, reminiscent of Toczé et al.'s design pattern "planned rebound effects"[39]. Chordia et al. [14] have coined the term "deceptive infrastructure" as a way to present the interaction of socio-economical factors and power dynamics with interface features with a negative impact on the user. The deceptive patterns identified in this study similarly interact with broader economic and environmental contexts, which results in long-term systemic impacts rather than in instantaneous individual damage. Because the end goal of these patterns is not deception but enforced adoption, design strategies aimed at imposing the domination of unsustainable digital features on the market need to be better characterized, so that in turn we know on which grounds they could be regulated.

We have relied on the concept of deceptive patterns to analyze our findings because it has already been used to bridge design and regulation actors [19]. Regulating AI deployment on the basis of its environmental impact poses a challenge, as do all environmental regulations, which is why different opportunities for regulating the imposed adoption of AI are suggested here. Deceptive patterns could be used as a way to support regulations on a variety of dimensions that are key for reducing the environmental footprint of the ICT sector more generally: connecting sustainability interests with fair market competition; favoring individual sustainable practices; and promoting environmental transparency.

### 5.1 On growth, market capture and lock-ins

"Attempts to present AI as desirable, inevitable, and as a more stable concept than it actually is follows well-worn historical patterns: one of the most important ways for a technology to gain market share and buy-in is to present it as an inevitable and necessary part of future infrastructure, and in turn to encourage the adaptation or building of new, anticipatory infrastructures around it." Widder and Hicks [41]

These radical and synchronous interface and interactive experience modifications are repercussions of an "AI race for technological advantage" [13]: ": AI is developed with the intention to reach technological superiority and gain market shares. Gen-AI services are particularly capital-intensive, which gives the edge to actors that already dominate the digital market. This is tied to an economic approach of rapid product scaling and market capture. If demand is weak or non-existent, it must be stimulated. In this context, the importance of user adoption crystallizes around AI features, which makes the highlighting of AI features in interfaces all the more important. Because AI adoption has been anticipated with massive investments and growth announcements <sup>7</sup>,<sup>8</sup>, the use of deceptive patterns is not surprising, as user engagement on AI features is highly strategic. Users are expected to pay, either directly through subscriptions or indirectly through data, the price of the investments in AI systems, and this could remain unnoticed as everything is made to sweep away users' doubts, disrupt existing digital habits and discourage alternative digital practices. In this situation, companies might have recourse to deceptive patterns to ensure that users will stay on the service despite a price increase.

In such competitive circumstances, more regulations on fair competition in the digital market could also converge with environmental interests. Monopolistic digital platforms are already known for deliberately increasing users' dependency through various means

<sup>&</sup>lt;sup>6</sup>https://www.deceptive.design/

<sup>&</sup>lt;sup>7</sup>https://www.sequoiacap.com/article/ais-600b-question/

<sup>&</sup>lt;sup>8</sup>https://www.ft.com/content/634b7ec5-10c3-44d3-ae49-2a5b9ad566fa

LIMITS '25, 11th Workshop on Computing Within Limits, June 26-27, 2025, Online

(e.g. lack of interoperability [17], path dependency [4], cognitive dependency [8]). This makes it hard for users to leave digital platforms when, for example, they cannot export their data in a format based on an open standard, or when they cannot collaborate or communicate with other people who do not agree to switch platforms. The situation with AI adoption is even more worrying, as AI features are used as an excuse for increasing subscription fees in a saturated market with weak growth perspectives. A good example of this trend is Microsoft 365's January 2025 announcement of a 30% fee increase on family subscriptions and 43% on personal subscriptions<sup>9</sup>. As companies continue to invest huge sums in AI deployment<sup>10</sup>, we can reasonably anticipate further price hikes. If users knew in advance that their subscription price could be doubled, would they stay on the service and continue creating files, adding contacts, planning meetings on the service?

In this regard, there is an interesting connection between environmental interests and installed digital market regulations. Legal dispositions such as the Digital Markets Act <sup>11</sup> could integrate deceptive patterns to ensure that platforms in the most advantageous market positions are not using deceitful strategies to limit user exodus and taking advantage of lock-ins to dramatically increase subscription prices.

Toczé et al.'s *"lock-ins"* [39] are used in a capitalistic logic that goes against sustainable practices, e.g. by undermining the consumers power to threaten to leave the platform. While environmental sustainability might not be the main motivation behind further regulations of market competition in the digital sector, a better regulation of the enforcement of user adoption could give more sustainable alternative platforms a chance.

# 5.2 On discouraging moderation in the use of digital services

Deceptive design patterns can be drawn upon to demonstrate that design choices discourage the adoption or continuation of more sustainable digital uses. Because most products on the digital market are service-based (SAAS), consumers have little power to refuse incremental evolution of the services: cloud-hosted digital services allow companies to transform interfaces overnight and instantly affect all their users. Users who choose to use a platform for certain features, interface quality and/or prices cannot be ensured that all of these characteristics will remain unchanged the next day. In particular, the first strategy we presented describes how interfaces' affordances and information architecture have evolved and now demand more efforts to pursue digital activities without the use of generative AI: the interfaces' choice architecture [31] pushes users towards the least sustainable features of the interface, thus degrading the user experience for users who seek to avoid increasing their individual footprint.

For example, the authors of this paper are at the time of writing collaborating on Overleaf and we cautiously dismissed the language suggestions that started to pop up a few months ago. How can we trust that we will still be able to disable such suggestions in the

 $^{11}$ https://digital-markets-act.ec.europa.eu/index<sub>e</sub>n

future? If we switch to different software for sustainability reasons, how can we be sure that it will not follow the same path? So far, users have very little control over their software on this matter. In this forced AI transition, tech companies have largely taken advantage of this grey area regarding consumer choice.

We argue that users who chose a software for its features (or because it doesn't have certain features) should be better protected in order to promote the adoption and continued practice of *moderate digital uses* [42], a term that encompasses all digital practices that consume less data. As deceptive patterns are useful to demonstrate how users' choices are influenced for corporate purposes, we believe applying the concept to sustainability issues could also be fruitful, by helping to establish whether design choices intentionally or unintentionally support sustainable engagement with the digital service. While placing the responsibility of feature adoption on the user won't solve the situation in itself, designers and tech companies could be held responsible for ensuring at each update that users could still use the service without increasing their environmental footprint, with no impact on user experience.

## 5.3 On greenwashing and lack of transparency

The EU's AI Act, which came into effect in August 2024, requires companies to report on the environmental impact of their activities. Sustainability and environmental responsibility reports are actually not new, and are already published by large tech companies every year <sup>12</sup>. While the reliability of self-reported data is questionable and significant limitations exist in the reporting of environmental effects (e.g. disclosing energy consumption or water consumption is not required from the AI Act) [2], this is a step towards more transparency. Similar arguments have been put forward in favour of transparent energy ratings for AI chatbots as a way to encourage minimizing the environmental costs of these services [28].

To prevent greenwashing, such as when companies like Meta <sup>13</sup> or Microsoft <sup>14</sup> purport to strive for net-zero emissions and water consumption while actually employing design strategies that maximize such emissions, companies could be held accountable for the way their platform designs steer demand on features toward higher environmental impacts. Similarly, energy rating labels do not prevent users from intensifying their AI use. In this sense, we agree with the argument articulated by Kender et al. [23], who in the case of social media found that "while functional design (e.g. algorithmic content suggestion, data collection) is increasingly being regulated, aesthetic design is barely considered in discussions about the ef[f]ects". What they call "aesthetic design power" could be harnessed to address situations in which design is employed to offset or contradict ethical or legal corporate commitments.

## 6 Future work

This work aimed at documenting a special moment of the longer transformations effected through interface and interaction design in relation to AI use. The applications and software we studied are changing so rapidly that part of our corpus has already been

<sup>&</sup>lt;sup>9</sup>https://www.microsoft.com/en-us/microsoft-365/blog/2025/01/16/copilot-is-nowincluded-in-microsoft-365-personal-and-family/

<sup>10</sup> https://www.ft.com/content/634b7ec5-10c3-44d3-ae49-2a5b9ad566fa

<sup>&</sup>lt;sup>12</sup>See for example: https://citizen.snap.com/?lang=en-US

<sup>13</sup> https://sustainability.atmeta.com/

<sup>&</sup>lt;sup>14</sup>https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/msc/ documents/presentations/CSR/Microsoft-2024-Environmental-Sustainability-Report.pdf#page=10

obsolete as we submit are submitting this article. Further work could explore how AI-based features evolve and are updated, and how strategies and deceptive patterns evolve or develop in relation to generative AI diffusion. We also purposely didn't focus on companies or designers' intents in producing deceptive patterns. The design of AI-based features in these forms could be driven by different motives, including an effort to make AI technologies accessible for a new public. However, it is likely that the context of surveillance capitalism in which these AI-based features are developed highly influences their design, as this economic model relies on mass data collection, not only to offer custom features to users, but also to sell better-profiled data to advertisers. Further studies could uncover the different motivations that explain why AI-based features are designed the way they are. Finally, we see great potential in further exploring deceptive patterns as a means to bring about debates and regulations on the effects of design choices. Digital design needs to be better regulated to mitigate ICT's environmental effects, along with other effects, including social and psychological ones.

Aside from impacts at the regulatory level, further research aiming at better characterizing deceptive design patterns regarding sustainability could prove useful to designers looking to avoid perpetuating the growth paradigm of the digital sector.

## 7 Conclusion

This work has explored the role of interface and interaction design in the strategic diffusion of AI-based features on prominent digital platforms. We have identified two main strategies that reflect various deceptive design patterns: imposing AI features at the expense of existing non-AI features and promoting misleading technology narratives. In different ways, these strategies steer users towards the adoption of AI-based features by leveraging companies' dominant position with users and in the digital market. The paper discusses the opportunities offered by the "deceptive patterns" concept for advancing the environmental regulation of the design of digital services.

### 8 Acknowledgments

We thank the reviewers for their appreciation and constructive comments. We are also grateful for Jean-Yves Bart's careful proofreading and language corrections. This work was supported by a French government grant managed by the Agence Nationale de la Recherche as part of the France 2030 program, reference ANR-22-EXEN-0006 (PEPR eNSEMBLE / PC5).

### References

- [1] 2025. Energy and AI Analysis. https://www.iea.org/reports/energy-and-ai
- [2] Nicolas Alder, Kai Ebert, Ralf Herbrich, and Philipp Hacker. 2024. AI, Climate, and Transparency: Operationalizing and Improving the AI Act. https://doi.org/ 10.48550/arXiv.2409.07471 arXiv:2409.07471 [cs].
- [3] Apple Inc. 1987. EDUCOM 1987 John Sculley Keynote. http: //archive.org/details/Apple-Computer-Educom-1987-John-Sculley-Keynote-Address-On-The-Knowledge-Navigator
- [4] W. Brian Arthur. 1989. Competing Technologies, Increasing Returns, and Lock-In by Historical Events. <u>The Economic Journal</u> 99, 394 (1989), 116–131. https: //doi.org/10.2307/2234208 Publisher: [Royal Economic Society, Wiley].
- [5] Maria Teresa Baldassarre, Danilo Caivano, Berenice Fernandez Nieto, Domenico Gigante, and Azzurra Ragone. 2023. The Social Impact of Generative AI: An Analysis on ChatGPT. In Proceedings of the 2023 ACM Conference on Information

Technology for Social Good (GoodIT '23). Association for Computing Machinery, New York, NY, USA, 363–373. https://doi.org/10.1145/3582515.3609555

- [6] Jane Bennett. 2016. <u>The Enchantment of Modern Life: Attachments, Crossings,</u> and Ethics. Princeton University Press. <u>https://doi.org/10.1515/9781400884537</u>
- [7] Adrien Berthelot, Eddy Caron, Mathilde Jay, and Laurent Lefèvre. 2024. Estimating the environmental impact of Generative-AI services using an LCA-based methodology. <u>Procedia CIRP</u> 122 (Jan. 2024), 707–712. https://doi.org/10.1016/j. procir.2024.01.098
- [8] Gunter Bombaerts, Tom Hannes, Martin Adam, Alessandra Aloisi, Joel Anderson, Lawrence Berger, Stefano Davide Bettera, Enrico Campo, Laura Candiotto, Silvia Caprioglio Panizza, Yves Citton, Diego DâAngelo, Matthew Dennis, Nathalie Depraz, Peter Doran, Wolfgang Drechsler, Bill Duane, William Edelglass, Iris Eisenberger, Beverley Foulks McGuire, Antony Fredriksson, Karamjit S. Gill, Peter D. Hershock, Soraj Hongladarom, Beth Jacobs, Gábor Karsai, Thomas Lennerfors, Jeanne Lim, Chien-Te Lin, Mark Losoncz, David Loy, Lavinia Marin, Bence Péter Marosán, Chiara Mascarello, David McMahan, Jin Y. Park, Nina Petek, Anna Puzio, Katrien Schaubroek, Jens Schlieter, Brian Schroeder, Shobhit Shakya, Juewei Shi, Elizaveta Solomonova, Francesco Tormen, Jitendra Uttam, Marieke Van Vugt, Sebastjan Vörös, Maren Wehrle, Galit Wellner, Jason M. Wirth, Olaf Witkowski, Apiradee Wongkitrungrueng, Dale S. Wright, and Yutong Zheng. 2024. From an attention economy to an ecology of attending. A manifesto. https://doi.org/10.48550/arXiv.2410.17421 arXiv:2410.17421 [cs].
- [9] Alan Borning, Batya Friedman, and Deric Gruen. 2018. What pushes back from considering materiality in IT?. In <u>Proceedings of the 2018 Workshop on</u> <u>Computing within Limits (LIMITS '18)</u>. Association for Computing Machinery, New York, NY, USA, 1–6. https://doi.org/10.1145/3232617.3232627
- [10] Alan Borning, Batya Friedman, and Nick Logler. 2020. The 'invisible' materiality of information technology. <u>Commun. ACM</u> 63, 6 (May 2020), 57–64. https: //doi.org/10.1145/3360647
- [11] L. R. Caporael. 1986. Anthropomorphism and mechanomorphism: Two faces of the human machine. <u>Computers in Human Behavior</u> 2, 3 (Jan. 1986), 215–234. https://doi.org/10.1016/0747-5632(86)90004-X
- [12] Noah Castelo, Maarten W. Bos, and Donald R. Lehmann. 2019. Task-Dependent Algorithm Aversion. Journal of Marketing Research 56, 5 (2019), 809–825. https: //www.jstor.org/stable/26967271 Publisher: [Sage Publications, Inc., American Marketing Association].
- [13] Stephen Cave and Seán S. ÓhÉigeartaigh. 2018. An AI Race for Strategic Advantage: Rhetoric and Risks. In Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society (AIES '18). Association for Computing Machinery, New York, NY, USA, 36–40. https://doi.org/10.1145/3278721.3278780
- [14] Ishita Chordia, Lena-Phuong Tran, Tala June Tayebi, Emily Parrish, Sheena Erete, Jason Yip, and Alexis Hiniker. 2023. Deceptive Design Patterns in Safety Technologies: A Case Study of the Citizen App. In Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23). Association for Computing Machinery, New York, NY, USA, 1–18. https://doi.org/10.1145/ 3544548.3581258
- [15] Amanda Chu. 2024. US slows plans to retire coal-fired plants as power demand from AI surges. Financial Times (May 2024).
- [16] Linda Di Geronimo, Larissa Braz, Enrico Fregnan, Fabio Palomba, and Alberto Bacchelli. 2020. UI Dark Patterns and Where to Find Them. <u>Proceedings of the</u> 2020 CHI Conference on Human Factors in Computing Systems (2020). https: //doi.org/10.1145/3313831.3376600
- [17] Cory Doctorow. 2021. Competitive compatibility: let's fix the internet, not the tech giants. <u>Commun. ACM</u> 64, 10 (Oct. 2021), 26–29. https://doi.org/10.1145/3446789
- [18] Sophia Falk, Aimee van Wynsberghe, and Lisa Biber-Freudenberger. 2024. The attribution problem of a seemingly intangible industry. <u>Environmental Challenges</u> 16 (Aug. 2024), 101003. https://doi.org/10.1016/j.envc.2024.101003
- [19] Colin M. Gray, Cristiana Teixeira Santos, Nataliia Bielova, and Thomas Mildner. 2024. An Ontology of Dark Patterns Knowledge: Foundations, Definitions, and a Pathway for Shared Knowledge-Building. In <u>Proceedings of the 2024 CHI</u> <u>Conference on Human Factors in Computing Systems (CHI '24)</u>. Association for Computing Machinery, New York, NY, USA, 1–22. https://doi.org/10.1145/ 3613904.3642436
- [20] Gianluca Guidi, Francesca Dominici, Jonathan Gilmour, Kevin Butler, Eric Bell, Scott Delaney, and Falco J. Bargagli-Stoffi. 2024. Environmental Burden of United States Data Centers in the Artificial Intelligence Era. https://doi.org/10.48550/ arXiv.2411.09786 arXiv:2411.09786 [cs].
- [21] Lujain Ibrahim, Luc Rocher, and Ana Valdivia. 2024. Characterizing and modeling harms from interactions with design patterns in AI interfaces. http://arxiv.org/ abs/2404.11370 arXiv:2404.11370 [cs].
- [22] Lars-Erik Janlert and Erik Stolterman. 2017. <u>Things That Keep Us Busy: The</u> Elements of Interaction. The MIT Press.
- [23] Kay Kender and Christopher Frauenberger. 2022. The Shape of Social Media: Towards Addressing (Aesthetic) Design Power. In <u>Proceedings of the 2022 ACM</u> <u>Designing Interactive Systems Conference</u> (DIS '22). Association for Computing Machinery, New York, NY, USA, 365–376. https://doi.org/10.1145/3532106. 3533470

LIMITS '25, 11th Workshop on Computing Within Limits, June 26-27, 2025, Online

- [24] Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren. 2025. Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models. https://doi.org/10.48550/arXiv.2304.03271 arXiv:2304.03271 [cs].
- [25] Anne-Laure Ligozat, Julien Lefevre, Aurélie Bugeau, and Jacques Combaz. 2022. Unraveling the Hidden Environmental Impacts of AI Solutions for Environment Life Cycle Assessment of AI Solutions. <u>Sustainability</u> 14, 9 (April 2022), 5172. https://doi.org/10.3390/su14095172
- [26] Alexandra Sasha Luccioni, Yacine Jernite, and Emma Strubell. 2024. Power Hungry Processing: Watts Driving the Cost of AI Deployment?. In <u>The 2024</u> <u>ACM Conference on Fairness, Accountability, and Transparency.</u> 85–99. https: //doi.org/10.1145/3630106.3658542 arXiv:2311.16863 [cs].
- [27] Alexandra Sasha Luccioni, Emma Strubell, and Kate Crawford. 2025. From Efficiency Gains to Rebound Effects: The Problem of Jevons' Paradox in AI's Polarized Environmental Debate. https://doi.org/10.48550/arXiv.2501.16548 arXiv:2501.16548 [cs] version: 1.
- [28] Sasha Luccioni, Boris Gamazaychikov, Sara Hooker, Régis Pierrard, Emma Strubell, Yacine Jernite, and Carole-Jean Wu. 2024. Light bulbs have energy ratings—so why can't AI chatbots? <u>Nature</u> 632, 8026 (2024), 736–738. https: //www.nature.com/articles/d41586-024-02680-3 Publisher: Nature Publishing Group UK London.
- [29] Maria Luce Lupetti and Dave Murray-Rust. 2024. (Un)making AI Magic: A Design Taxonomy. In Proceedings of the 2024 CHI Conference on Human Factors in <u>Computing Systems (CHI '24)</u>. Association for Computing Machinery, New York, NY, USA, 1–21. https://doi.org/10.1145/3613904.3641954
- [30] Spyros Makridakis. 2017. The forthcoming Artificial Intelligence (AI) revolution: Its impact on society and firms. <u>Futures</u> 90 (June 2017), 46–60. https://doi.org/ 10.1016/j.futures.2017.03.006
- [31] Arunesh Mathur, Mihir Kshirsagar, and Jonathan Mayer. 2021. What Makes a Dark Pattern... Dark? Design Attributes, Normative Considerations, and Measurement Methods. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21). Association for Computing Machinery, New York, NY, USA, 1–18. https://doi.org/10.1145/3411764.3445610
- [32] Clément Morand, Anne-Laure Ligozat, and Aurélie Névéol. 2024. How Green Can AI Be? A Study of Trends in Machine Learning Environmental Impacts. https://doi.org/10.48550/arXiv.2412.17376 arXiv:2412.17376 [cs].
- [33] Simone Natale and Andrea Ballatore. 2020. Imagining the thinking machine: Technological myths and the rise of artificial intelligence. <u>Convergence</u> 26, 1 (Feb. 2020), 3-18. https://doi.org/10.1177/1354856517715164 Publisher: SAGE Publications Ltd.

- [34] Amanda K. Newendorp, Mohammadamin Sanaei, Arthur J Perron, Hila Sabouni, Nikoo Javadpour, Maddie Sells, Katherine Nelson, Michael Dorneich, and Stephen B. Gilbert. 2024. Apple's Knowledge Navigator: Why Doesn't that Conversational Agent Exist Yet?. In <u>Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)</u>. Association for Computing <u>Machinery, New York, NY, USA</u>, 1–14. https://doi.org/10.1145/3613904.3642739
- [35] Matteo Pasquinelli, Cristina Alaimo, and Alessandro Gandini. 2024. AI at Work: Automation, Distributed Cognition, and Cultural Embeddedness. <u>Tecnoscienza</u> – Italian Journal of Science & Comp. Technology Studies 15, 1 (July 2024), 99– 131. https://doi.org/10.6092/ISSN.2038-3460/20010 Artwork Size: 5MB Publisher: Tecnoscienza – Italian Journal of Science & Technology Studies.
- [36] Chris Preist, Daniel Schien, and Eli Blevis. 2016. Understanding and Mitigating the Effects of Device and Cloud Service Design Decisions on the Environmental Footprint of Digital Infrastructure. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems. ACM, San Jose California USA, 1324– 1337. https://doi.org/10.1145/2858036.2858378
- [37] Alexandra Sasha Luccioni, Sylvain Viguier, and Anne-Laure Ligozat. 2022. Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model. https://doi.org/10.48550/arXiv.2211.02001 ADS Bibcode: 2022arXiv221102001S.
- [38] Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2020. Energy and Policy Considerations for Modern Deep Learning Research. <u>Proceedings of the AAAI Conference on Artificial Intelligence</u> 34, 09 (April 2020), 13693–13696. <u>https://doi.org/10.1609/aaai.v34i09.7123 Number: 09.</u>
- [39] Klervie Toczé, Maël Madon, Muriel Garcia, and Patricia Lago. 2022. The Dark Side of Cloud and Edge Computing: An Exploratory Study. In <u>Computing within</u> Limits. LIMITS. https://doi.org/10.21428/bf6fb269.9422c084
- [40] Gaël Varoquaux, Alexandra Sasha Luccioni, and Meredith Whittaker. 2025. Hype, Sustainability, and the Price of the Bigger-is-Better Paradigm in AI. https: //doi.org/10.48550/arXiv.2409.14160 arXiv:2409.14160 [cs].
- [41] David Gray Widder and Mar Hicks. 2024. Watching the Generative AI Hype Bubble Deflate. https://doi.org/10.48550/arXiv.2408.08778 arXiv:2408.08778 [cs].
- [42] Kelly Widdicks, Christian Remy, Oliver Bates, Adrian Friday, and Mike Hazas. 2022. Escaping unsustainable digital interactions: Toward "more meaningful" and "moderate" online experiences. <u>International Journal of Human-Computer</u> <u>Studies</u> 165 (Sept. 2022), 102853. https://doi.org/10.1016/j.ijhcs.2022.102853
- [43] Yunhao Zhang and Renée Gosline. 2023. Human favoritism, not AI aversion: People's perceptions (and bias) toward generative AI, human experts, and human-GAI collaboration in persuasive content generation. Judgment and Decision Making 18 (2023), e41. https://doi.org/10.1017/jdm.2023.37