

What Can Computers Do Now? Dreyfus Revisited for the Third Wave of Artificial Intelligence

Ben Schuering¹, Thomas Schmid^{2,3,4}

¹ Leibniz Universität Hannover

² Martin Luther University Halle-Wittenberg

³ Lancaster University in Leipzig

⁴ Leipzig University

Abstract

In recent years, artificial intelligence (AI) has seen significant advances that have in fact exceeded even optimistic prognoses. Using data-driven AI, namely deep learning techniques, it has been demonstrated that computers may now be equipped with abilities of remarkable scope and quality, such as solving image and text processing tasks at human level. Large language models, in particular, have sparked debates regarding opportunities and challenges of this rapidly developing area. Will remaining fundamental challenges of data-driven AI, such as factual or logical mistakes, be overcome for good if complemented and hybridized with symbolic AI techniques, such as knowledge representation and reasoning? Will systems of artificial general intelligence (AGI) emerge from this, possessing common sense and in fact completing the decades-old quest for AI that motivated the raise of the field in the 1950s? In the light of these questions, we review the likewise, decades-old philosophical debate about capabilities and limitations of computers from a hybrid AI point of view. Here, we discuss how hybrid AI is coming closer to disproving Hubert Dreyfus' famous statements regarding what computers can not do. At the same time, we shed light on a lesser discussed challenge for hybrid AI: the possibility that its developers might be its biggest limiters.

Introduction

Since its very early days, the field of artificial intelligence (AI) has been fascinated by the idea of creating computers with human-like cognitive abilities. Great expectations came with successes of early search and planning strategies in the 1950s, due to which some researchers anticipated machines with human-like capabilities within the next decade already. While this turned out to be too optimistic, new hope grew with more advanced logic AI called expert systems introduced in the 1970s, and even more so with the raise of neural networks in the 1980s. Both paradigms, however, come with particular challenges. Data-driven AI, such as neural networks, is traditionally not particularly good at working with logic concepts. Logic AI, on the other hand, can usually not be automatically adapted to dynamic changes. Therefore, AI researchers have been striving to complement the strengths of these first two waves of AI in a third wave or paradigm, typically referred to as hybridization or hybrid AI.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Neither pure logic AI nor hybrid AI have so far reached human-like abilities and common sense. Data-driven AI, however, has recently seen major advancements due to the concept of large language models (LLMs), at some tasks exceeding all previously known abilities. LLMs, such as GPT-4.0 or BERT, have been successfully used to write reports and poems, to summarize or extend texts (Min et al. 2023). Multimodal LLMs, such as Stable Diffusion, are even able to create images from text input (Zhang et al. 2023). At the same time, the rather unwanted phenomenon of hallucination has gained significant interest (Tonmoy et al. 2024). It illustrates not only the limitations of LLMs and their statistical nature (Lake et al. 2017), but has also sparked interest in overcoming limitations by combining machine learning with knowledge engineering techniques (Colon-Hernandez et al. 2021). With deep learning and LLMs having already reached capabilities at an unprecedented quality, such further improvements raise the question of how close the next generation of AI will actually come to a so-called artificial general intelligence (AGI) with human-like abilities and common sense. Not few in the field assume that the third wave of AI will push the limits of what computers can do significantly, and may even overcome the longest standing barriers separating man and machine.

What Can Computers Not Do?

As motivating the vision of machines with human-like abilities and common sense may have been for early AI researchers, it turned out to be not only ambitious, but in fact its hardest challenge ever. Moreover, this ideal provided an easy target for critics. Philosopher Hubert L. Dreyfus, in particular, became famous for his fundamental criticism on AI formulated in several publications (Dreyfus 1965, 1972; Dreyfus and Dreyfus 1986; Dreyfus 1992). Condensed in the title of his seminal book *What Computers can't do* (Dreyfus 1972), Dreyfus has been arguing that AI will fundamentally never be able to deliver common sense, moral and ethical reasoning, contextual awareness, and emotions.

By the time Dreyfus initially published these thoughts, AI was dominated by mathematical and logic approaches that were built to deal with symbols and representations. Against this background, he emphasized that human cognitive capacities rely strongly on unconscious processes and extend beyond explicit articulation (Dreyfus 1965, 1972). In the

light of further developments in this direction in the area of neural networks, Dreyfus renewed his criticism two decades later with his follow-up book *What Computers Still Can't Do* – then dissecting the predictive nature of data-driven AI, which relies on emulation rather than genuine understanding (Dreyfus 1992).

A core target of Dreyfus' critique on AI was the goal of creating machines with common sense (Dreyfus and Dreyfus 1986). Shaped by diverse factors, such as upbringing and social interactions, individuals develop nuanced ethical and moral reasoning approaches, often grounded in shared understandings – such as the unwritten rule of refraining from checking one's phone during a job interview – commonly recognized as common sense. Dreyfus argued that these intricate reasoning approaches, deeply embedded in human experience, pose significant challenges for AI systems attempting to replicate them successfully (Dreyfus 1972; Descartes 1996).

Not surprising, Dreyfus's propositions provoked prompt and reactions from the AI community in the following years. Papert, for example, reacted already in 1968 with a harsh rebuttal, not only basically denying most arguments brought forward by Dreyfus but even questioning the legitimacy of doing so (Papert 1968). In the year of publication, computer scientist Bruce G. Buchanan reviewed *What Computers Can't Do* critically (Buchanan 1972), encouraging readers to adopt a contemporary view of humanity and the world, distinct from traditional scientific viewpoints. Was the overall approach misaligned? As Crossman elaborates in "The Kiss and the Promise", Dreyfus's work falls short in shedding light on the meaning or understanding of AI, instead focusing on the societal misuse of computers, particularly the potential replacement of human emotional interactions with machine-human interchange (Crossman 1985).

Critics argue that the mistakes lie not in what computers can not do, but rather in what they can do and how they achieve it (Collins 1996). Further, debates surrounding Dreyfus's work delve into the dichotomy of "Body and World." Scholars like Hubert Haugeland argue that intelligent bodies are fundamentally situated, with relevance contingent on the essentially human situation. Intelligence would then reside bodily in the world, implicating not only information processing but also neurobiology and anthropology (Haugeland and Dreyfus 1996). Moreover, Dreyfus focuses excessively on the detailed architecture and physical form of computers (Collins 1996), neglecting real-time interactions and social dimensions shaping AI capabilities.

Keeping up with Dreyfus' tradition, Timothy Koschmann explored the symbolic grounding issue further and extended the critique to the inadequacy of specifying all exception clauses and the conjectural nature of arguments (Koschmann 1996). He emphasized the necessity for AI to employ multiple strategies and approaches, contributing to the broader non-formalist, anti-representationalist debate within the situated cognition controversy (Koschmann 1996). This is largely the case for all of Dreyfus' work, as he was one of the main followers of this philosophy, and is further supported by other prevalent reviewers and philosophers (Buchanan 1972; Koschmann 1987; Papert 1968).

Is the Hybrid Whole More Than The Sum Of Its AI Parts?

Over the course of the last three decades, a variety of approaches to create such hybrid AI systems have emerged, such as genetic programming of neural networks, logic spiking neural networks, hybrid expert systems or knowledge-based neural networks. A popular classical scheme discriminates between three types of hybrid AI: unified, transformational, and modular hybrid systems (McGarry, Wermter, and MacIntyre 1999). While modular design of hybrid systems is still relevant today (Schmid 2023), the other two paradigms have been differentiated in more recent schemes, following the idea of design patterns popularized by modern software engineering (van Bekkum et al. 2021; Witschel et al. 2021).

Usually, the justification for any hybrid AI approaches is the assumption that by combining two opposite approaches, each individual drawback will balance out each other. Or in other words: that the individual strengths will complement each other. In fact, however, this is not guaranteed. What, for example, if the weaknesses of opposite approaches complement each other - instead of the strengths? We will discuss this in the following for common sense and contextual awareness, ethical reasoning, and moral reasoning.

- **Common Sense and Contextual Awareness** A field that has also received great advancements based on hybrid AI, is the world of personal assistants such as Siri or Alexa, and virtual chat assistants. Often, they use a combination of different AI paradigms, such as Natural Language Processing (NLP), Machine learning, text-to-speech(TTS), knowledge graphs and for the virtual chat assistants the Chatbot Framework. Concepts, such as a knowledge graph, allow the system to have an underlying graph modelling the connection and relationship between different entities. This allows the system to process an input and provide the relevant output. Overall, this in function with ML algorithms allows the system to function properly. NLP is the basis for the computer to understand what the user is inputting.

A possible interpretation of some of these systems is that they indeed showcase the ability to have common sense and contextual awareness (Varde et al. 2015). Questions, such as "What is the traffic looking like?", showcase the ability to use contextual awareness. Given that Siri, for example, would respond based on live data of surrounding traffic and the location. This creates a seemingly accurate response, including cases such as weather requests. As this adapts based on various factors that surround the user, contextual awareness may be determined in this case, at least to a certain degree (Signorelli 2018). Common sense lacks an official definition, making it extremely difficult to argue if it has been completed (Varde et al. 2015; Shanahan et al. 2020). However, examples from virtual assistants that learn and better during their usage, allow a certain amount of common sense to be included in the responses given, which creates the argument on the possibility of systems being able to have common sense. For example, if a user asks the virtual personal assistant to book a flight, the system uses its

common sense capabilities to understand the user's intention and provide relevant information and options, such as the available flights, schedules, and prices. This common-sense decision-making process enables the virtual personal assistant to provide a more human-like and intuitive experience for the user, which can help to build trust and satisfaction in the technology.

- **Ethical Reasoning.** For illustration, we will in here consider aspects from the field of autonomous driving, which is today heavily dependent on various AI technologies. From a technical standpoint, autonomous vehicles are relevant because they have to scale to big levels on a relatively simple premise. Hybrid AI systems in use on autonomous vehicles rely on computer vision, control systems, machine learning, and deep learning (Kisačanin 2017). Both logic AI and data-driven AI have been criticized by Dreyfus for not being able to make ethical decisions. Can hybrid AI systems overcome this? Regarding control systems and machine learning algorithms of vehicles, this can not really be confirmed (Fridman et al. 2017). Control systems, as the name suggests, are responsible for the vehicle's controls. This means they determine the vehicles' path, speed, and trajectory based on the sensor's inputs and other algorithms at work.

Highlighting, the fact that if the vehicle had to choose between taking a human's life and hitting a pole, it would in most cases hit the pole. Therefore, seemingly demonstrating ethical reasoning (Gerdes and Thornton 2015). Dreyfus, however, would contradict. Due to the nature of his statements and the definition of ethical reasoning, hybrid AI is not able to make an ethical decision in a specific case: Even though it seems to do so, it can only be made if it is hardcoded (Bonnemains, Saurel, and Tessier 2018). The control systems in this case are logically deciphering what to do, not reasoning by themselves. Overall, this would mean that although it can mirror the effect of ethical reasoning, it is not engaging in it fully.

- **Moral Reasoning.** Both logic AI and data-driven AI have been criticized by Dreyfus for not being able to achieve moral reasoning. This concept itself is hard to define, as many different interpretations exist of what it encompasses. Yet, hybrid AI has seen advancements and successes regarding this topic. In medical recommendation and diagnosis systems (Kulikowski 1980), for example, a machine learning component may be trained on data related to medical ethics, such as best practices for informed consent and respect for patient autonomy. The corresponding expert system component provides a human-like understanding of moral considerations and values and can make informed moral decisions. In a situation where the system is faced with a decision about whether to recommend a risky medical procedure, the system would then imply its moral reasoning capabilities to make a decision that aligns with moral principles and values, such as the patient's autonomy and informed consent (Montani 2008). Similar to the argument about ethical reasoning, however, moral reasoning remains incomplete. It trains on different data, and creates the moral

reasoning through the human input (Kulikowski 1980). It can be argued that the system is not actually reasoning morally, instead remaining logical. In theory, it still only follows if...then, as it does not understand the moral aspect. Instead, it uses its logical reasoning based on the inputs provided by the medical expert to determine what the best solution could be, further proving an incomplete showcase of moral reasoning (Montani 2008).

What Are The Major Bottlenecks for the Third Wave of AI?

While the assumption of complementary strengths makes it tempting to assume unlimited opportunities for hybrid AI, one may still well consider several relevant bottlenecks even where complementation works generally out well. We consider in particular these potential major bottlenecks:

- **Data.** One of the major challenges in data-driven AI is the factor of training data, in particular the labeling of such (Roh, Heo, and Whang 2021). Data selection and quality has direct influence on the quality of the resulting model, and with the increase in available data in today's world, one would expect the possibilities of a model's capabilities to increase. However, this is true to a certain extent as many practitioners will be able to tell. Different developers and stakeholders have diverse criteria of selection, quality, and success. Similar issue arise around the labelling of data, which even if (or potentially due to) being carried out by humans may sometimes be faulty, biased, or insufficient in amount (Roh, Heo, and Whang 2021). These critique points had been raised by Dreyfus' later work (Dreyfus 1992).

Is hybrid AI free from these issue of data-driven AI? Although the functionality of hybrid AI systems do not solely rely on training data like data-driven AI, due to their hybrid nature they remain based at least in part on training data inputted by the human. The hypothesis that the process of system design does not lead to human-like abilities seems to remain true as the problem can stem from even before the model is created, hence, even if the models become better the issues remain and Dreyfus remains correct (Dreyfus 1992).

- **Predictions.** An interesting approach to the argument of what hybrid AI can or cannot do is looking at other description used for the overarching AI systems. One of these alternative names is "Prediction machines", used for example in economics or businesses contexts (Agrawal, Gans, and Goldfarb 2018), but also supported by findings of psychological research (Schrimpf et al. 2021). This perspective highlights some further reasons why it is so difficult to argue against statements made by Dreyfus: The argument is the fact that machines do not actually "know" what they are doing, but rather predict what they should be doing, based on the input given. This is supported by the statement made by Dreyfus that it's incredibly hard to separate the knowledge from the knower (Dreyfus 1992). Famously, humans often know more than they can put into words, let alone put into contextual statements based on if...then.

With this idea, it becomes apparent why the work of Dreyfus with its astonishing polarizing attack on AI as such remains studied. Today's hybrid AI systems are close to demonstrating all the features shown by Dreyfus, but do so in an almost imposter-type way (Bonnemains, Saurel, and Tessier 2018). Ethical and moral reasoning, for example, can be displayed in these systems, even to great success, as shown in the world of autonomous vehicles and medical systems; however, they are not actually reasoning in this way, instead logically reasoning based on rules provided by the expert. Interestingly, some psychologists argue that humans follow the same pattern, as in some cases it is unclear if it is logical reasoning taught as ethical or moral reasoning, compared to actually reasoning morally or ethically (Kulikowski 1980). Instead, the difference seems to be that we understand and can create new and adaptive interpretations of these rules, especially to previously unknown challenges. Together with the fact that these systems rely heavily on predictions, it makes sense as to why they are incomplete (Agrawal, Gans, and Goldfarb 2018).

- **Humans.** Let us take on yet another perspective and assume for a moment that we as a human developer would aim to create a hybrid AI system with one of humankind's most important abilities: creativity. Would we in this respect self limit our capability to produce creative, intelligent machines? In line with the statement made above and supported by Dreyfus: Yes, because we can not actually put into words how to achieve this goal, and hence cannot create the systems needed.

Contrarily, considering that data make the systems function, the realization has to be made, that the issues stem from a much deeper problem. Therefore, it might even be accurate to say, that models and systems are not the sole problem, nor the techniques currently in use. Rather, human developers preparing and employing the data remain one of the primary blocking factors of advancements. While overall data availability is increasing, the problems faced increase in the same manner, such as the difficulty of properly managing all of this data, or labelling it correctly. This influences the training of the models and also its results (Roh, Heo, and Whang 2021). Hence, returning to the statement that the problem of Dreyfus's statements is not entirely related to the technology itself, but largely the human definitions of certain aspects. Further, Dreyfus statements remain so open-ended, and human definitions so precise, yet unexplainable (Dreyfus 1992).

Even if this can change, in hybrid AI development there are always experts, designing the models and creating them. These limitations can be considered vital, due to the size of AI systems, as the "black box" within these becomes increasingly complicated (Yu and Ali 2019). In many systems, not only in LLMs but also hybrid systems, it is actually hard to tell comprehensively on which grounds individual decisions are made. To some extent, this may be reduced by approaches such as rule extraction from neural networks (Jacobsson 2005).

Conclusions

Over the last decades, Dreyfus' ideas and views have been shaping both scientific and public debate about goals and challenges of AI. Although not particularly welcomed by everyone in the AI community, Dreyfus has made significant intellectual contributions to the field. The variety of responses to the existing problems analyzed by Dreyfus have been prompting scholars, researchers, and developers to continually reassess and refine their understanding of the capabilities of AI (Schmid et al. 2021).

While both data-driven and hybrid AI systems have seen significant progress in many areas, they still have limitations and are not capable of fully replicating human intelligence and experience. Major bottlenecks for future systems will be data, the predictive nature of such systems, but also the humans developing these systems. Hybrid AI may be able to make decisions based on pre-defined principles and values, perform natural language processing and recognize patterns in data, but not to replicate human emotions, empathy, creativity, and intuition.

To this end, the advancements in hybrid AI have not completely disproven Dreyfus' statements on the limitations of AI. Remember, he argued that AI systems would never be able to fully replicate human intelligence and experience, due to their lack of understanding. While hybrid AI systems have come closer to disproving the statements regarding commonsense or ethical reasoning since the 1960s, they still lack the conclusion to do – or are only creating a superficial disapproval doing this superficially.

So are Dreyfus' work and arguments still relevant in times of the third wave of AI? They are. This is not solely due to the philosophical nature of his statements, and its interpretability range being so large. There is little indication, for example, that actual moral and ethical reasoning can be carried out in a reliable and communicated in a human-understandable fashion by today's AI systems. Certified safety criteria, for example, will therefore have to be decided by law (Yu and Ali 2019).

References

- Agrawal, A.; Gans, J.; and Goldfarb, A. 2018. *Prediction machines: the simple economics of artificial intelligence*. Harvard Business Press.
- Bonnemains, V.; Saurel, C.; and Tessier, C. 2018. Embedded ethics: some technical and ethical challenges. *Ethics and Information Technology*, 20: 41–58.
- Buchanan, B. G. 1972. *Review of Hubert Dreyfus' What Computers Can't Do: A Critique of Artificial Reason*. (Harper & Row, New York, 1972).
- Collins, H. M. 1996. Embedded or embodied? A review of Hubert Dreyfus' what computers still can't do. *Artificial Intelligence*, 80(1): 99–118.
- Colon-Hernandez, P.; Havasi, C.; Alonso, J.; Huggins, M.; and Breazeal, C. 2021. Combining pre-trained language models and structured knowledge. *arXiv preprint arXiv:2101.12294*.

- Crossman, E. K. 1985. The kiss and the promise: A review of Hubert L. Dreyfus' What Computers Can't Do: The Limits of Artificial Intelligence. *Journal of the Experimental Analysis of Behavior*, 44(2): 271.
- Descartes, R. 1996. *Meditations on First Philosophy with Selections from the Objections and Replies/trans.* and edited by John Cottingham.
- Dreyfus, H.; and Dreyfus, S. E. 1986. *Mind over machine*. Simon and Schuster.
- Dreyfus, H. L. 1965. Alchemy and artificial intelligence.
- Dreyfus, H. L. 1972. What computers can't do: The limits of artificial intelligence.
- Dreyfus, H. L. 1992. *What computers still can't do: A critique of artificial reason*. MIT press.
- Fridman, L.; Brown, D. E.; Glazer, M.; Angell, W.; Dodd, S.; Jenik, B.; Terwilliger, J.; Kindelsberger, J.; Ding, L.; Seaman, S.; et al. 2017. MIT autonomous vehicle technology study: Large-scale deep learning based analysis of driver behavior and interaction with automation. *arXiv preprint arXiv:1711.06976*, 1.
- Gerdes, J. C.; and Thornton, S. M. 2015. *Implementable Ethics for Autonomous Vehicles*, 87–102. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-662-45854-9.
- Haugeland, J.; and Dreyfus, H. L. 1996. Body and world: a review of What Computers Still Can't Do: A Critique of Artificial Reason. *Artificial intelligence*, 80(1): 119–128.
- Jacobsson, H. 2005. Rule extraction from recurrent neural networks: Ataxonomy and review. *Neural Computation*, 17(6): 1223–1263.
- Kisačanin, B. 2017. Deep Learning for Autonomous Vehicles. In *2017 IEEE 47th International Symposium on Multiple-Valued Logic (ISMVL)*, 142–142.
- Koschmann, T. 1996. Of Hubert Dreyfus and dead horses: some thoughts on Dreyfus' What Computers Still Can't Do: (MIT Press, Cambridge, MA, 1992); liii+ 354 pages, 13.95.
- Koschmann, T. D. 1987. *Mind over machine: The power of human intuition and expertise in the era of the computer: Hubert L. Dreyfus and Stuart E. Dreyfus* (Basil Blackwell, Oxford, 1986); 223 pages, £ 15.00.
- Kulikowski, C. A. 1980. Artificial intelligence methods and systems for medical consultation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-2(5): 464–476.
- Lake, B. M.; Ullman, T. D.; Tenenbaum, J. B.; and Gershman, S. J. 2017. Building machines that learn and think like people. *Behavioral and brain sciences*, 40: e253.
- McGarry, K.; Wermter, S.; and MacIntyre, J. 1999. Hybrid neural systems: from simple coupling to fully integrated neural networks. *Neural Computing Surveys*, 2(1): 62–93.
- Min, B.; Ross, H.; Sulem, E.; Veyseh, A. P. B.; Nguyen, T. H.; Sainz, O.; Agirre, E.; Heintz, I.; and Roth, D. 2023. Recent advances in natural language processing via large pre-trained language models: A survey. *ACM Computing Surveys*, 56(2): 1–40.
- Montani, S. 2008. Exploring new roles for case-based reasoning in heterogeneous AI systems for medical decision support. *Applied Intelligence*, 28: 275–285.
- Papert, S. A. 1968. The artificial intelligence of Hubert L. Dreyfus: A budget of fallacies.
- Roh, Y.; Heo, G.; and Whang, S. E. 2021. A Survey on Data Collection for Machine Learning: A Big Data - AI Integration Perspective. *IEEE Transactions on Knowledge and Data Engineering*, 33(4): 1328–1347.
- Schmid, T. 2023. A Systematic and Efficient Approach to the Design of Modular Hybrid AI Systems. In Martin, A.; Hinkelmann, K.; Fill, H.-G.; Gerber, A.; Lenat, D.; Stolle, R.; and van Harmelen, F., eds., *Proceedings of the AAAI 2023 Spring Symposium on Challenges Requiring the Combination of Machine Learning and Knowledge Engineering (AAAI-MAKE 2023)*.
- Schmid, T.; Hildesheim, W.; Holoyad, T.; and Schumacher, K. 2021. The AI Methods, Capabilities and Criticality Grid. *KI - Künstliche Intelligenz*, 35(0): 425–440.
- Schrimpf, M.; Blank, I. A.; Tuckute, G.; Kauf, C.; Hosseini, E. A.; Kanwisher, N.; Tenenbaum, J. B.; and Fedorenko, E. 2021. The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences*, 118(45): e2105646118.
- Shanahan, M.; Crosby, M.; Beyret, B.; and Cheke, L. 2020. Artificial Intelligence and the Common Sense of Animals. *Trends in Cognitive Sciences*, 24(11): 862–872.
- Signorelli, C. M. 2018. Can Computers Become Conscious and Overcome Humans? *Frontiers in Robotics and AI*, 5: 121.
- Tonmoy, S.; Zaman, S.; Jain, V.; Rani, A.; Rawte, V.; Chadha, A.; and Das, A. 2024. A Comprehensive Survey of Hallucination Mitigation Techniques in Large Language Models. *arXiv preprint arXiv:2401.01313*.
- van Bekkum, M.; de Boer, M.; van Harmelen, F.; Meyer-Vitali, A.; and Teije, A. t. 2021. Modular design patterns for hybrid learning and reasoning systems: a taxonomy, patterns and use cases. *Applied Intelligence*, 51(9): 6528–6546.
- Varde, A.; Tandon, N.; Chowdhury, S. N.; and Weikum, G. 2015. Common sense knowledge in domain-specific knowledge bases. *Tech Rep, Max-Planck-Institut für Informatik (Saarbrücken, Germany)*.
- Witschel, H. F.; Pande, C.; Martin, A.; Laurenzi, E.; and Hinkelmann, K. 2021. Visualization of Patterns for Hybrid Learning and Reasoning with Human Involvement. In Dornberger, R., ed., *New Trends in Business Information Systems and Technology: Digital Innovation and Digital Business Transformation*, 193–204. Cham: Springer.
- Yu, R.; and Ali, G. S. 2019. What's Inside the Black Box? AI Challenges for Lawyers and Researchers. *Legal Information Management*, 19(1): 2–13.
- Zhang, C.; Zhang, C.; Zhang, M.; and Kweon, I. S. 2023. Text-to-image diffusion model in generative ai: A survey. *arXiv preprint arXiv:2303.07909*.