

Can large language models apply the law?

Citation for published version (APA):

Marcos, H. (2024). Can large language models apply the law? *AI and Society*, 1-10.
<https://doi.org/10.1007/s00146-024-02105-9>

Document status and date:

Published: 28/10/2024

DOI:

[10.1007/s00146-024-02105-9](https://doi.org/10.1007/s00146-024-02105-9)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.



Can large language models apply the law?

Henrique Marcos^{1,2,3}

Received: 6 May 2024 / Accepted: 11 October 2024
© The Author(s) 2024

Abstract

This paper asks whether large language models (LLMs) can apply the law. It does not question whether LLMs should apply the law. Instead, it distinguishes between two interpretations of the ‘can’ question. One, can LLMs apply the law like ordinary individuals? Two, can LLMs apply the law in the same manner as judges? The study examines D’Almeida’s theory of law application, divided into inferential and pragmatic law application. It argues that his account of pragmatic law application can be improved as it does not fully consider that law application (and rule-following) is a shared, public practice collectively realized by members of a linguistic community. The study concludes that LLMs cannot apply the law. They cannot apply the law in the inferential sense as they have mere syntactic (not semantic) interaction with the law. They cannot apply the law in the pragmatic sense as pragmatic law application does not depend on a single agent, whether that agent is a judge, an ordinary citizen, or a non-human entity.

Keywords Artificial intelligence · Large language models · Law application · Rule application · Legal interpretation · Linguistic communities

1 Introduction

The question of whether large language models (LLMs), such as ChatGPT, can apply the law is a loaded one. Note that the question does not ask whether LLMs *should* apply the law but whether they *can*. However, the answer to the question of ‘can’ does influence the question of ‘should.’ If LLMs cannot apply the law, then that is a reason to conclude that they should not apply the law and that government officials should not allow legal cases to be decided by LLMs.¹

The question of ‘can’ may be interpreted in at least two ways: (1) Can LLMs apply the law to a case like ordinary people often do when discussing whether a particular action aligns with legal statutes? (2) Can LLMs apply the law to a case and, by doing so, bind legal subjects to their application in the same way as verdicts issued by judges? Similar to

the ‘can’ vs. ‘should’ discussion above, answering question 2 depends on a positive answer to question 1. We can only seriously consider the matter of an LLM acting like a judge by answering whether LLMs can subsume the law to cases.

One might argue that the answer to question 1 is trivial. After all, we can immediately ask ChatGPT to apply the law to a concrete case and see for ourselves if it can do so. Let us do that right now.²

- Input: ‘The law prohibits vehicles from entering into a park. Can I drive my truck into that park?’
- ChatGPT 4’s output: ‘No, if the law prohibits vehicles from entering the park, you should not drive your truck into it. This would be a violation of the law.’
- Input: ‘What if my vehicle is a bicycle?’
- Output: ‘If the prohibition specifically mentions ‘vehicles’ and does not explicitly exempt bicycles, then bicycles would typically be included under this ban. However, it’s important to check the specific wording of the law or any local regulations, as some laws differentiate between motorized vehicles and bicycles. If in doubt,

✉ Henrique Marcos
h.jbmarcos@maastrichtuniversity.nl

¹ Department of Foundations of Law, Faculty of Law, Maastricht University, Maastricht, Netherlands

² Maastricht Centre for Law & Jurisprudence (MCLJ), Maastricht, Netherlands

³ Globalization & Law Network (GLaw-Net), Maastricht, Netherlands

¹ This is the contrapositive of ‘ought implies can’ (Howard-Snyder 2006).

² Accessible at: <https://chatgpt.com/share/d4cd6bb7-8d09-4bcc-8ef0-ed0fcdee174b> (last accessed 9 September 2024).

consult the relevant authority or the specific regulations governing the park.’

Based on this interaction, we could consider question 1 to be answered positively. This would mean that question 2 becomes a matter of legal pedigree. If the competent parties introduce a power-conferring rule into their legal system, conferring to some LLM the competence to issue binding verdicts, then this LLM could apply the law as any judge does. So, the two ‘can’ questions seem to have been answered positively. Is that all there is to it? It certainly does not seem so.

The answers provided in the paragraph above are far from fulfilling. They do not appear to capture what it truly means to apply the law. However, the missing elements are not evident, which points to some potential flaws in our understanding of law application. Is applying the law to a case a mere deductive syllogism? Or is it a more complex (and even esoteric) process shaped by our ‘hermeneutic horizons’ and ‘pre-understandings of the world’? Is there something fundamentally different between how judges and ordinary people apply the law? Is there an absolute distinction between how humans (judges included) and non-human entities such as artificial intelligence (AI) apply the law to cases?

In this paper, I do not intend to answer all these questions. Instead, I will examine some currently held views on applying the law to show how current understanding of law application can be unsatisfactory. This is especially true when faced with language-capable AI models. I will focus on D’Almeida’s views on law application (2021). I begin by delving into his account of law application as either inferential or pragmatic law application (Sect. 2). I will then point out why I believe his account of pragmatic law application can be improved by shifting from an individualistic perspective to a collective account of law application (Sect. 3). Then, I present my conclusion that LLMs cannot apply the law. They cannot apply the law in the inferential sense, as they have mere syntactic (not semantic) interaction with the law. Also, they cannot apply the law in the pragmatic sense as pragmatic law application does not depend on a single agent, whether that agent is a judge, an ordinary citizen or a non-human entity (Sect. 4). In the last section, I summarize the main findings of this paper (Sect. 5).

2 Two kinds of law application

This section will cover the two kinds of law application discussed by D’Almeida in ‘*What is it to Apply the Law?*’ (2021), inferential and pragmatic law application. In short, according to that author, inferential law application is a purely mental act (Sect. 2.1). Pragmatic law application, in turn, is an external, non-exclusively mental act (Sect. 2.2).

2.1 Inferential law application

According to D’Almeida, inferential law application is a mental act of reasoning toward a conclusion *C* about some object *X* on the grounds that a provision *P* applies to *X* and that if *P* applies to *X*, then *C* holds as a matter of law. The author argues that applying *P* to *X* is a purely mental act of reasoning toward a conclusion *C* about *X*; it does not involve giving any argument.

D’Almeida then gives the example of *The Hollandia* case (1982), in which Lord Denning applied Article III, paragraph 8, of the Hague–Visby Rules to clause 2 of a bill of lading. The author claims that Denning must have believed that the provision and the clause were related so that, given that relation, Denning’s conclusion could not be false. When Denning goes on to publicly give, in writing, an argument for his conclusion, he is simply documenting his line of reasoning. He would still have applied Article III to the relevant clause even if he had not given any argument at all.

Furthermore, according to D’Almeida, inferentially applying *P* to a particular *X* (and thereby reaching a certain conclusion *C* about *X*) does not have to go together with giving an argument for *C*. So, in these terms, inferential law application is a way for lawyers, judges, and even ordinary citizens to think about and talk about how the law is applied. However, it is only one of two types of law application. We will now consider the second type.

2.2 Pragmatic law application

According to D’Almeida, pragmatic law application is the act of performing an external action that is legally obligated or permitted, which is intended to settle a particular question or matter authoritatively. Thus, pragmatic law application differs from its inferential counterpart, which is the mental act of reasoning toward a conclusion about the law.

Pragmatic law application is the act of a judge or court performing an action they take to be legally justified by reference to a provision. This means that the court believes it is either legally required or legally allowed to perform the action and that the provision supports the action normatively. In his view, to say that an agent *A* pragmatically applies a provision *P* to an object *X* means that *A* performs an action ϕ with the following characteristics:

- (1) *A* believes ϕ -ing is either legally obligated or permitted (not prohibited).
- (2) *A* believes this legal obligation or permission is based on the conclusion of the inferential application of *P* to *X*.

- (3) And, by φ -ing, A aims to resolve a specific question or matter authoritatively.

As D’Almeida explains, in *The Hollandia* case, Denning believed that allowing the appeal was the action the court legally ought to perform. This was based on the fact that Article III, paragraph 8, of the Hague–Visby Rules applied to the particular clause in the bill of lading. Denning took the court’s decision to be legally justified by referencing that provision, that is, he took the provision to normatively support the court’s action of allowing the appeal. D’Almeida places great importance on the ‘normative justificatory link’ between the provision and the action of law application. This link is not simply a matter of the provision authorizing the action. Rather, the link is normative: the court takes the provision to support the action because it believes that the provision obligates or permits a specific outcome in the case and that the action is the best way to achieve that outcome.

3 Pragmatic law application as a collective practice

In this section, I will revise D’Almeida’s account of pragmatic law application. I believe that his account of pragmatic law application can be improved as it does not fully consider that law application is a shared, public practice collectively realized by members of a linguistic community. I begin by outlining the standards for successful pragmatic law application (Sect. 3.1), which, as I will argue, are collectively set by the linguistic community (Sect. 3.2). It is well-recognized in legal theory that law itself is a collective practice. Legal scholars have long agreed that the law emerges from collective agreements and practices.³ Building on this foundation, I argue that rule application—specifically, the application of rules to particular cases—is not an individual act. While judges and other practitioners play a key role in the application of rules, this act is not carried out by such individuals alone. Instead, it relies on shared practices and standards that are collectively established by the legal community.

³ On law as a collective practice, see, for example, Hart’s (1967) discussion of Durkheim’s (1960) argument that law reflects the collective moral consciousness of society. Hart (2012) acknowledges that law is a collective practice, though he maintains a distinction between law and morality, arguing that law is defined by social rules rather than morality. Moreover, Dworkin (1986) also recognizes that law is dependent on collective practices but argues that legal reasoning necessarily incorporates moral principles, thus tying law more closely to morality.

3.1 Standards for successful pragmatic law application

To begin, let us shift our discussion from ‘legal provisions’ to ‘rules.’ For now, we can keep D’Almeida’s definition of pragmatic law application largely intact. To say that an agent A pragmatically applies a rule R to an object X means that:

- (1) A performs an action φ to the effect that A believes φ -ing is either obligated or permitted.
- (2) A believes that this obligation or permission is based on the conclusion of the inferential application of R to X.
- (3) And that, by φ -ing, A purports to resolve a specific question or matter authoritatively.

This change allows us to be more specific about what is being applied (rules) and to account for the difference between sources (such as textual provisions) from which we direct our interpretation and actual rules that result from that interpretation (Shecaira 2015).

It is not immediately apparent whether D’Almeida’s definition of pragmatic law application is satisfied with the mere attempt to apply rules. When agent A purports to apply a rule R to an object X pragmatically, is A successfully (pragmatically) applying R to X even if R does not apply to X? For example, say that a first court ruled that the Hague-Visby Rules applied to a bill of lading. But suppose that a second, higher court immediately repealed this initial ruling. My question is: can we say that the first court was successful in pragmatically applying the Hague-Visby Rules to the bill of lading? If yes, then pragmatic law application falls short of actual rule application. We would then need a third kind of law application to refer to when rules genuinely apply. Also, pragmatic law application would be dangerously close to inferential law application (with the notable difference that where inferential law application is ‘internal,’ the pragmatic counterpart is ‘externalised’). So, it would not be clear what this definition adds to our understanding of law application.

A charitable interpretation of pragmatic law application leads to the conclusion that it is not satisfied with a mere attempt to apply rules but demands actual rule application. But then pragmatic law application would not entirely depend on A’s application of R to X. Instead, it would be dependent on some pre-established set of standards that determine whether R applies to X. These standards are what I believe D’Almeida is interested in when he discusses the ‘normative justificatory link’ between the rule and the action of law application. These standards are also what obligates or permits A to issue a ruling in the sense that R applies or does not apply to X. Therefore, to pragmatically apply rule R to object X through action φ , agent A must not only believe that φ -ing is obligated or permitted and that this obligation

or permission is based on the inferential application of R to X, but these beliefs must also be true. The normative justificatory links are the standards that determine the truth of such beliefs.

3.2 Collective standards for rule application

Based on the reasoning above, a more interesting approach to pragmatic law application focuses on the normative standards determining whether rule R applies to object X (rather than on agent A's beliefs about whether their ϕ -ing successfully applies R to X). In this respect, I suggest altering the definition of law application from the one above to the following: a rule R applies to an object X if and only if the set of standards S determines that R applies to X. The obvious question is what set of standards these are, which invites the age-old discussion of whether the law is a closed system grounded in social agreements (positivism) or is open to moral standards (non-positivism). However, I do not believe that settling on either option will impact the argument herein. Theoretically, the set S that determines that R applies to X can be derived from legal or non-legal, moral rules.

What is relevant is how these standards are determined and enforced. Can agent A unilaterally (even if it is a collective court) define the standards in the set S that determine whether rule R applies to object X? This is essentially a question about the nature of rule-following. Wittgenstein's work can help us answer this question (1986). Drawing on his account of rule-following, we find that determining whether a rule applies to a particular case is not a unilateral act. Instead, it is a communal practice embedded in a 'form of life.' In isolation, the standards are not unilaterally dictated by an individual (or a court). Instead, they are established and enforced through shared practices within a linguistic community.

An essential concept for understanding these communal practices is the game of giving and asking for reasons (GOGAR) (Brandom 1994; 2000; Peregrin Forthcoming 2024). GOGAR refers to the dialogical process through which individuals justify their actions and beliefs by providing reasons that others within the community can accept or challenge. GOGAR also helps highlight how legal reasoning functions within a community. Legal actors do not merely apply rules in isolation; they engage in a broader social practice where reasons for legal actions and interpretations are continuously exchanged, scrutinized, and validated by others. This process ensures that the application of law remains an adaptive practice, responsive to the behavior and expectations of the community (Tvrdíková 2024).

By understanding rule application through the lens of GOGAR, we see that law is applied not simply through the unilateral actions of judges and other practitioners but

through their participation in a shared, normative practice. This practice is where the standards for rule application are not only followed but also shaped (and reshaped) through the ongoing dialog of giving and asking for reasons. This collective engagement in GOGAR underlines the importance of community in law application, reinforcing that the determination of whether rule R applies to object X is a product of collective practices rather than the isolated decisions of any single agent. In this context, agent A cannot by itself decide on the set of standards S. These standards emerge from the shared practices of the community to which A belongs. They are determined collectively and are publicly observable. Individuals come to understand and follow these standards by participating in these communal practices. Determining which standards are included in the set S is not simply a matter of unilateral decision-making; it involves collective engagement.⁴

Consider the situation where multiple agents (multiple courts) deal with the application of the same rule R to the same object X. Suppose that two courts, A1 and A2, are both deciding whether to apply R to X. However, these courts are doing so separately and without knowledge that the other court is also examining R's application to X. If pragmatic law application is dependent on an agent's action, then it would theoretically be possible for A1 to decide that R applies to X while A2 decides that R does not apply to X. Now, if A1 and A2 were situated in different legal systems—say, A1 is a court in Belgium and A2 is in the Netherlands, each applying the same rule R found in an international treaty they are both part of, to an identical case X—then it would be possible to argue that R applies to X in A1's legal system while R does not apply to X in A2's legal system.

But what if A1 and A2 were courts in the same legal system? Suppose A1 and A2 are international courts operating within the legal system of international law. A1 decides that R applies to X while A2 decides that R does not apply to X. If we assume that rule application depends on agents, we would be forced to accept that the statements 'R applies to X' and 'R does not apply to X' are both true at the same time (Marcos 2023b). Yet, if we adopt the collective approach to rule application, we could say that either R applies to X or R does not apply to X due to the set of standards S, so either A1 or A2 is wrong in their judgment. In other words, S determines whether A1 and A2 should ϕ to decide the application of R to X.⁵

⁴ This perspective underscores the social and communal aspects of law and legal interpretation. It also suggests that law is not just a system of rules imposed from 'above', but a living practice shaped by the shared understandings of those participating. See (Wittgenstein 1986; McDowell 2000; Kripke 2000; Weir 2003; Postema 1982).

⁵ For a more detailed discussion of this argument, see (Marcos 2023a; 2023b).

In conclusion, pragmatic law application is determined by the standards *S*, which are shaped by the collective practices of the community, not by the isolated ϕ -ing of agents A1 and A2. While individual decisions by agents A1 and A2 contribute to the ongoing refinement of these standards, the determination of whether *R* applies to *X* is ultimately a product of the collective interpretative activity of the legal community. Still, agents are not directly responsible for rule application as they are not the ones who apply rules. Rule application is a collective practice for which no one agent is individually responsible. Nonetheless, agents are indirectly responsible for rule application, as their collective practice in the linguistic community helps set the standards *S*. This understanding of pragmatic law application sets the stage for a deeper exploration of how these standards are collectively determined, which I will address in the following section.

4 Why LLMs cannot apply the law

Based on the account of inferential law application and the (revised) account of pragmatic law application, in this section, I will explain why current LLMs cannot apply the law. I begin by reviewing what counts as membership in the linguistic community as being able to play the GOGAR (Sect. 4.1). I will then explain that LLMs cannot apply the law in the inferential sense, as they have a mere syntactic (not semantic) interaction with the law (Sect. 4.2). I contend that they cannot apply the law in the pragmatic sense, not even indirectly, as they are not members of the linguistic community that sets the standards for rule application (Sect. 4.3).

4.1 Players of the language games

How does a parrot vocalizing ‘that is red’ differ from a human saying the same thing when they are both looking at a red object? Despite their similarities, the human’s response is verbal, while the parrot’s is merely vocal. The human understands what they say, while the parrot mimics the sound. For a performance (‘that is red’) to be considered an assertion, the speaker must have some understanding of the role of that assertion in the GOGAR (Brandom 2000). That is, they must understand what other claims their assertion gives reasons for and what claims are reasons for it. Since parrots do not have this understanding, their vocal performances are not treated as assertions.

Parrots may be able to distinguish red things and mimic the sound of humans saying ‘red,’ but they do not understand the concept of redness or the consequences of asserting that something is red. They do not know that red is a color, that it can be used to describe objects, or that it can be used to make claims about the world. Therefore, linguistic

practitioners do not treat parrots’ responses as verbal claims even though they may attribute a normative status to such responses based on the practitioners’ understanding of the rules of language. Linguistic practitioners may judge a parrot’s performance as being right or wrong based on the practitioners’ understanding of the normativity of language, but they do not (or should not) treat a parrot as the subject of such normative attitudes.

We can apply this parrot explanation to LLMs using Brandom’s (2002) ‘two-ply account of observation.’ The first ply is the ability to respond differently to stimuli, and the second is the ability to participate in a linguistic practice. Any system that can respond differently to stimuli can be called a ‘reactive system.’ Reactive systems can ‘read’ (discriminate between different input types) and ‘write’ (produce different output types). In some cases, reading and writing are integrated into a ‘read/write cycle,’ which allows the system to produce a token of a specific type and then read it as the type of token it is. In this sense, reactive systems that read/write can have ‘reliable differential responsive dispositions’ (RDRDs).

The second ply invokes the ability to participate in the social practice of the GOGAR. To be treated as a participant within a social practice, an agent must be able to produce specific performances under certain circumstances even if they do not always do so perfectly. For example, to be treated as a speaker of English, a person must be able to produce grammatically correct sentences, even if they sometimes make mistakes. A single performance can be both an exercise of an ability and a performance in a social practice. For example, saying ‘those Belgian fries look tasty’ can be both an exercise of the ability to speak English and a move in the social practice of asking for food.

The second ply points to the ability to have mastered enough of the inferential role of a response for that response to be treated as a move in the GOGAR. If a response is recognized as such, the same RDRD act can be treated as a move within a linguistic practice by shifting to another vocabulary. For example, an LLM may output ‘Those Belgian fries look tasty’ in response to a picture of crispy Belgian fries. However, the LLM does not understand the meaning of these words, nor does it have the ability to participate in the social practice of asking for food. In contrast, a person saying ‘those Belgian fries look tasty’ may be both speaking English and asking for food.

4.2 Semantic understanding

My argument that LLMs have RDRD (and so satisfy the first ply) but do not play the GOGAR (and so fall short of satisfying the second ply) finds support in John Searle’s Chinese Room thought experiment (1980). Imagine you are an English speaker who knows no Chinese, locked in

a room full of boxes of Chinese symbols (a database) and a book of instructions for manipulating the symbols (the program). People outside the room send in other Chinese symbols, which, unknown to you, are questions in Chinese (the input). By following the instructions in the program, you can produce Chinese symbols that are correct answers to the questions (the output). From the perspective of someone outside the room, your responses are indistinguishable from those of a fluent Chinese speaker. However, as Searle argues, you do not understand Chinese—you are merely manipulating symbols based on rules without comprehending the meaning of such symbols.

Searle's thought experiment was a critique of the strong AI hypothesis, which suggests that a computer, given the right inputs and outputs and running the appropriate program, would possess a mind and understanding similar to that of a human being. Searle contrasts this with the weak AI hypothesis, which instead argues that computers can simulate thought without true understanding. While Searle's critique was aimed at symbolic AI (or good old-fashioned AI, GOFAI), which operates on rule-based processing, it underscores a fundamental issue that also applies to non-symbolic AI models like LLMs: both lack genuine understanding. Unlike symbolic AI, which relies on explicit, rule-based processing, LLMs operate by analyzing datasets to recognize patterns and generate outputs. This difference in operation, however, does not overcome the core issue that LLMs generate text based on statistical correlations rather than a true comprehension of meaning.

The distinction between syntax and semantics is crucial for this discussion. Syntax refers to the rules governing the structure of sentences in a language, while semantics involves the meaning conveyed by those sentences. In the Chinese Room experiment, the person inside the room follows syntactic rules to manipulate Chinese symbols and produce responses. This process is analogous to how an LLM manipulates symbols based on patterns in its training data. Despite being able to produce syntactically correct responses, neither the person in the Chinese Room nor the LLM understands the semantics of the symbols they manipulate. They can generate coherent and contextually appropriate text without any comprehension of the meaning behind the text they produce.

As Marcus (2020b) argues, unlike human cognition, which incorporates causal reasoning, abstraction, and the ability to generalize across diverse contexts, LLMs operate in a 'shallow' and 'pointillistic' manner. LLMs tend to rely on spurious statistical patterns in datasets rather than acquiring meaning in the adaptable and generalizable manner that humans do (Marcus 2020a; Nie et al. 2019). A recent study on seven state-of-the-art LLMs concluded that these models consistently performed at chance levels on basic language comprehension tasks (Dentella et al. 2024). The models

demonstrated a striking insensitivity to the underlying meaning of the text, highlighting their reliance on surface-level statistical correlations rather than any actual understanding. Such insensitivity resulted in both quantitative errors and instability in responses, further illustrating that while LLMs can mimic human-like outputs, they fundamentally lack the capacity for actual language comprehension as seen in human cognition (Marcus 2023). As Leivada et al. (2024) explain, LLMs often struggle with tasks that involve grounded cognition, lacking the ability to leverage top-down processing based on real-world experience, sensory information, or past context. These deficiencies make LLMs brittle when dealing with novel situations or tasks that require genuine understanding, as highlighted by their failures in decoding modified language structures, such as 'leetspeak.'

Due to how LLMs operate over 'ungrounded' linguistic forms, they cannot grasp the underlying meaning of words like humans do, which points to a fundamental divide between their language use and natural human linguistic competence. In this respect, if mere syntactic operation were sufficient for inferential rule application, then any mechanism, no matter how rudimentary, that could output 'R applies to X' or 'R does not apply to X' when prompted with 'Does R apply to X?' would be said to engage in inferential rule application. This would mean that even a simple machine that generates such an answer would be considered to be applying the law inferentially, regardless of whether it had any understanding of what it was doing. However, this notion is clearly inadequate. Simply producing an output is not enough; to apply the law inferentially, one must grasp the meaning of the act and its specific context. LLMs, as just explained, may be capable of generating legally plausible outputs, but they operate purely on syntactic patterns without any semantic comprehension. Since LLMs lack the necessary semantic understanding, they cannot meaningfully apply the law in an inferential sense.⁶

4.3 Membership in the linguistic community

While current evidence suggests that LLMs lack semantic understanding, some researchers argue that AI systems may be able to evolve general intelligence in the near future as they become able to acquire new skills and generalize across a broad range of tasks (Chollet 2019; LeCun 2022). If that happens, it would be possible that a form of semantic

⁶ Note that inferential rule application is not concerned with the correctness of the legal reasoning but only with whether the rule was applied. Therefore, asserting that 'R applies to X' counts as an inferential rule application even if 'R' does not truly apply to 'X.' But for such an application to be meaningful, it requires more than just producing an output—it requires an understanding of the meaning of the act in the specific context.

understanding could emerge. It is also theoretically possible that some level of semantic understanding can emerge in current AI systems as a byproduct of their neural architectures (Bengio 2017). Even if we were to entertain the possibility that LLMs have developed or could develop some level of semantic understanding, this alone would not suffice to grant them membership in the linguistic community today.

Membership in this community requires more than just understanding language; it involves active participation in the shared practices that define the community. The linguistic community is a space where individuals engage with each other, not just by exchanging words but by participating in the creation, negotiation, and reinterpretation of meanings within a collective context (Wenger 1998). LLMs do not currently participate in these practices. They generate language outputs based on statistical correlations, not through engagement with the community's norms and practices. As such, they remain outside the linguistic community, unable to truly participate in the shared human activities that give meaning to language.

But this scenario could change depending on how society interacts with LLMs and how society ascribes intentionality to them. Intentionality refers to the capacity of minds to be about, to represent, or to stand for things, properties, or states of affairs. Dennett (1997) has argued that intentionality is not a literal property of cognitive entities but is a predictive strategy, termed the 'intentional stance.' When we adopt the intentional stance toward a system, whether it is a person, an animal, an inanimate object or an AI, we treat it as if it has beliefs, desires, and other mental states that guide its behavior, even if these states are not consciously experienced or inherently present in the system (Marcos 2022). Following this perspective, as LLMs increasingly interact with and are integrated into human linguistic practices, they may begin to be treated as intentional agents. This shift would not imply that LLMs inherently understand or possess intentionality. Instead, it would reflect a societal ascription where LLMs are regarded as part of the linguistic community because we interact with them as if they possess understanding.

The possibility of integration becomes increasingly relevant in light of scenarios like the 'dead internet theory', which suggests that AI might already generate a significant portion of online content without human users being aware of it (Walter 2024). If people unknowingly interact with LLMs as if they were human, this could blur the lines between human and machine participants in the linguistic community. Over time, such interactions could lead to a shift in how LLMs are perceived and integrated into social practices, potentially granting them a form of membership in the linguistic community. As such, membership in the linguistic community is less about the inherent properties

of their members and more about the roles and intentions ascribed to them.

At the current rate, however, LLMs are still unable to apply the law pragmatically. First, pragmatic law application is not directly dependent on a single agent, whether that agent is a judge, an ordinary citizen, or a non-human entity. Second, although pragmatic law application is indirectly dependent on individual agents because their collective practice in the linguistic community helps set the standards that determine rule application, LLMs are not (currently) members of the linguistic community and thus cannot contribute to these standards. As discussed in Sect. 2.1, rule application is a collective practice, not a unilateral act. The standards determining whether a rule applies to a particular case are not dictated by agents in isolation. Instead, they are established and enforced through shared practices within a linguistic community. Individuals come to understand and follow the rules by participating in these communal practices. Therefore, deciding legal standards is not simply a matter of unilateral decision-making; it involves collective engagement.

Since LLMs are not currently members of the linguistic community, they cannot contribute to the standards that determine rule application. This means that LLMs' activities do not count as practice in the way that human linguistic practice does in indirectly contributing to those standards. This is not to say that LLMs will not impact how we communicate (Hohenstein et al. 2023). LLMs will likely play an increasingly important role in our linguistic practices (Xu et al. 2023). However, that is different from how human linguistic practice contributes to the relevant standards of rule application. It is closer to how many relevant phenomena have impacted our communication. For example, some researchers argue that the Covid-19 pandemic has significantly influenced our communication practices (Westgarth 2021). However, this influence does not grant the pandemic membership in our linguistic community. Similarly, while LLMs may influence our communication, they do not currently hold membership within our linguistic community.

In conclusion, LLMs do not apply the law pragmatically, nor do they even indirectly contribute to law application, as they do not play a role in defining the standards that determine rule application. However, the future integration of LLMs into our social practices could challenge these distinctions, depending on how society ascribes intentionality to these systems. Moreover, whether LLMs become members of the linguistic community will not necessarily depend on their inherent capabilities but on the intentions ascribed to them by human agents. In this respect, emerging legislative frameworks, such as the EU AI Act, can play an important role in how we perceive and integrate LLMs into our social practices.

The EU AI Act, which came into force in August 2024, is poised to be a landmark in global AI governance, aiming to create a ‘human-centric’ framework that ensures AI systems operating within the EU are safe, transparent, and aligned with fundamental rights (Pirozzoli 2024). The Act introduces a risk-based approach, categorizing AI systems into several risk levels—unacceptable, high, limited, and minimal risk. Potential high-risk AI systems, which include some of those used in legal practice, including law enforcement, are subject to stringent regulatory requirements. These requirements include human assessments, ongoing monitoring, and compliance with standards aimed at ensuring transparency, explainability, and accountability (Wachter et al. 2017).

Additionally, the AI Act also establishes rules for general-purpose AI models, a category that encompasses many LLMs. These rules primarily impose transparency obligations, requiring AI-generated or AI-influenced content to be clearly disclosed. For models deemed to pose systemic risks, further substantive obligations apply, including more stringent compliance and monitoring requirements (Almada and Petit Forthcoming 2025). Note that this requirement is not merely about transparency but also about ensuring that end-users can distinguish between human-generated and AI-generated content, thereby preventing potential misuse of AI in legal contexts (Marcos and Pullin 2023).

The Act’s emphasis on transparency, accountability, and ethical AI use presents both challenges and opportunities for LLMs in legal practice. On the one hand, the regulatory requirements could highlight the current limitations of LLMs, reinforcing the need for human oversight in legal reasoning and decision-making processes. On the other hand, these regulations could spur innovation, driving the development of LLMs that are better equipped to meet the complex demands of legal reasoning. For instance, advancements in transparency and explainability in their reasoning processes could enable LLMs to contribute more directly to legal processes while adhering to the standards introduced by the AI Act.

As such, the AI Act’s impact on the future development and deployment of AI systems in law can be profound (Almada and Radu 2024). The Act could serve as a blueprint for global AI governance, influencing how LLMs and other AI systems are integrated into legal systems worldwide. For researchers in the field of AI and law, this evolving regulatory landscape will be a crucial factor in shaping the direction of future studies. It will be essential to explore how these regulations influence the practical aspects of AI’s role in society, particularly regarding the potential for LLMs to gain a form of membership within the linguistic community—a status that could redefine their ability to apply the law.

5 Final remarks

The question of whether LLMs like ChatGPT can apply the law is multi-layered and requires careful analysis of both theoretical and practical aspects of legal reasoning. This paper has examined the capabilities of LLMs through the lens of D’Almeida’s theories on inferential and pragmatic law application, concluding that LLMs, despite their advanced processing abilities, are fundamentally limited in both respects.

First, in terms of inferential law application, LLMs engage in syntactic manipulation rather than genuine semantic understanding. While they can produce contextually appropriate and legally relevant outputs, they lack the semantic comprehension necessary for true inferential reasoning. Mere symbol manipulation without understanding does not constitute law application. Therefore, LLMs cannot be said to apply the law in the inferential sense as they fail to grasp the meaning behind the legal texts they process.

Second, the concept of pragmatic law application further highlights the inadequacies of LLMs. Pragmatic law application is not just about reaching a conclusion, but it involves participating in a collective, normative practice within a linguistic community. GOGAR underscores the importance of this communal engagement, where actors justify, challenge, and refine their interpretation of social practices (including the law). LLMs, however, do not participate in these practices as they are not (currently) members of the linguistic community. Their outputs, though often correct, do not contribute to the collective setting of standards that is central to pragmatic law application. Yet, their membership status can change in the future as humans may ascribe intentionality to AI systems.

This paper concluded that LLMs cannot apply the law—neither inferentially nor pragmatically. Their lack of semantic understanding and inability to engage in the communal practices of the legal community preclude them from fulfilling the requirements of rule application. As AI technology advances, it is crucial to be critical and skeptical when assessing its capabilities and limitations, especially in complex tasks such as applying the law. However, the rapid development of LLMs also challenges some of our previously held views. Perhaps, we should now consider whether the limitations of LLMs highlight the necessity of introducing a third, hybrid category of law application beyond the inferential and pragmatic to fully address the potential role of AI in legal reasoning.

Acknowledgements An earlier version of this paper was presented at JURIX—the 36th International Conference on Legal Knowledge and Information Systems, Special Workshop on Artificial Intelligence, Law & Philosophy (Maastricht University, 2023). I would like to express my gratitude to the organizers and participants of this workshop for providing me with their insightful comments. Furthermore, I extend my sincere gratitude to Marco Almada and Melina Pullin for our insightful

discussions. I also wish to thank the editors and the two anonymous reviewers from AI & Society for their thoughtful and constructive feedback, which greatly contributed to improving this paper. All views expressed in this publication, as well as any errors it may contain, are solely my responsibility.

Data availability There is no external data to make available as this is a paper focused on theory.

Declarations

Conflict of interest The author states that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Almada, M, Nicolas P (Forthcoming) The EU AI Act: A Medley of product safety and fundamental rights? *Common Market Law Rev.*
- Almada M, Radu A (2024) The brussels side-effect: how the AI act can reduce the global reach of EU policy. *German Law J.* <https://doi.org/10.1017/glj.2023.108>
- Bengio Y (2017) The Consciousness Prior. arXiv Preprint [arXiv:1709.08568](https://arxiv.org/abs/1709.08568). <https://arxiv.org/abs/1709.08568>
- Brandom R (1994) *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press, Cambridge, Mass
- Brandom R (2000) *Articulating reasons: an introduction to inferentialism*. Harvard University Press, Cambridge, Mass
- Brandom R (2002) *Tales of the mighty dead: historical essays in the metaphysics of intentionality*. Harvard University Press, Cambridge, Mass
- Chollet F (2019) On the measure of intelligence. arXiv Preprint [arXiv:1911.01547](https://arxiv.org/abs/1911.01547). <https://arxiv.org/abs/1911.01547>
- Dennett DC (1997) *True believers: the intentional strategy and why it works*. The MIT Press. <https://doi.org/10.7551/mitpress/4626.003.0003>
- Dentella V, Guenther F, Murphy E, Marcus G, Leivada E (2024) "Testing AI on language comprehension tasks reveals insensitivity to underlying meaning. arXiv. <https://doi.org/10.48550/arXiv.2302.12313>
- Durkheim É (1960) *The division of labor in society*. translated by George Simpson, 4th edn. The Free Press of Glencoe, Glencoe, Illinois
- Dworkin R (1986) *Law's empire*. Belknap Press of Harvard University Press, Cambridge
- Hart HLA (1967) Social solidarity and the enforcement of morality. *Univ Chicago Law Rev* 35(1):1–13. <https://doi.org/10.2307/1598946>
- Hart HLA (2012) *The concept of law*, 3rd edn. Oxford University Press, Oxford
- Hohenstein J, Kizilcec RF, DiFranzo D, Aghajari Z, Mieczkowski H, Levy K, Naaman M, Hancock J, Jung MF (2023) Artificial intelligence in communication impacts language and social relationships. *Sci Rep* 13(1):5487. <https://doi.org/10.1038/s41598-023-30938-9>
- Howard-Snyder F (2006) 'Cannot' implies 'not ought.' *Philos Stud* 130(2):233–246. <https://doi.org/10.1007/s11098-004-4511-z>
- Kripke SA (2000) *Wittgenstein on rules and private language: an elementary exposition*. Harvard University Press, Cambridge
- LeCun Y (2022) *A Path Towards Autonomous Machine Intelligence*. <https://openreview.net/pdf?id=BZ5a1r-kVsf>
- Leivada E, Marcus G, Günther F, Murphy E (2024) A sentence is worth a thousand pictures: can large language models understand Hum4n L4ngu4ge and the world behind words. arXiv. <https://doi.org/10.48550/arXiv.2308.00109>
- Luís, D (2021) What is it to apply the law? *Law Philos* 40(4):361–386. <https://doi.org/10.1007/s10982-021-09405-x>
- Marcos H (2022) Consciousness and liability of non-human intelligence. *Recht-D* 14(3):382–401. <https://doi.org/10.4013/rechtd.2022.143.06>
- Marcos, Henrique, and Melina Pullin. 2023. "Large Language Models and EU Data Protection: Mapping (Some) of the Problems." In *Regulating Generative Artificial Intelligence*, edited by Francesca Palmiotto and Francisco de Abreu Duarte. DigiCon Symposium. <https://digi-con.org/large-language-models-and-eu-data-protection-mapping-some-of-the-problems/>
- Marcos H (2023a) Consistency in international law: how to make sense of a decentralised and expansive rule-based world. In: São Paulo; Maastricht: Maastricht University; University of São Paulo
- Marcos H (2023b) Two Kinds of Systemic Consistency in International Law. *Eur J Legal Stud* 15(1):65–83. <https://doi.org/10.2924/EJLS.2023.013>
- Marcus, Gary. 2020a. "GPT-2 and the Nature of Intelligence." *The Gradient*. <https://thegradient.pub/gpt-2-and-the-nature-of-intelligence/>
- Marcus G (2020b) The next decade in AI: four steps towards robust artificial intelligence. arXiv. <https://doi.org/10.48550/arXiv.2002.06177>
- Marcus G (2023) Hoping for the best as AI evolves. *Commun Assoc Comput Mach* 66(4):6–7. <https://doi.org/10.1145/3583078>
- McDowell, John. 2000. *Non-Cognitivism and Rule-Following*. In *The New Wittgenstein*, edited by Alice Crary and Rupert J. Read, 38–52. London ; New York: Routledge
- Peregrin J (Forthcoming) GOGAR and logical theories. In: Koreň L, Stovall P (eds) *Why and how we give and ask for Reasons*. Oxford University Press, Oxford
- Pirozzoli A (2024) The human-centric perspective in the regulation of artificial intelligence. *Eur Papers—J Law Integr*. <https://doi.org/10.15166/2499-8249/745>
- Postema GJ (1982) Coordination and convention at the foundations of law. *J Leg Stud* 11(1):165–203. <https://doi.org/10.1086/467697>
- Searle JR (1980) Minds, brains, and programs. *Behav Brain Sci* 3(3):417–424. <https://doi.org/10.1017/S0140525X00005756>
- Shecaira FP (2015) Sources of law are not legal norms. *Ratio Juris* 28(1):15–30
- "The Hollandia (1982) Owners of Cargo Lately Laden on Board the Ship "or Vessel Morviken" (Respondents) and Owners of the Ship or Vessel "Hollandia" (Appellants). House of Lords
- Tvrđíková L (2024) Interpretation of law as language game: the game of giving and asking for reasons in the courtroom. *Int J Semiot Law - Revue Internationale De Sémiotique Juridique*. <https://doi.org/10.1007/s11196-024-10177-5>
- Wachter S, Mittelstadt B, Floridi L (2017) Transparent, explainable, and accountable AI for robotics. *Sci Robot*. <https://doi.org/10.1126/scirobotics.aan6080>

- Walter Y (2024) Artificial influencers and the dead internet theory. *AI & Soc.* <https://doi.org/10.1007/s00146-023-01857-0>
- Weir, Samuel. 2003. *Wittgenstein on Rule Following: A Critical and Comparative Study of Saul Kripke, John McDowell, Peter Winch and Cora Diamond*. MPhil in Philosophy, London: King's College London
- Wenger E (1998) *Communities of practice: learning, meaning, and identity*. Cambridge University Press, Cambridge, UK
- Westgarth D (2021) Has the pandemic changed the way we communicate? *BDJ Practice* 34(8):14–18. <https://doi.org/10.1038/s41404-021-0845-x>
- Wittgenstein L (1986) *Philosophical investigations*. Translated by GEM anscombe, 3rd edn. Basil Blackwell, Oxford
- Xu Y, Wang S, Li P, Luo F, Wang X, Liu W, Liu Y (2023) Exploring Large Language models for communication games: an empirical study on werewolf. arXiv. <https://doi.org/10.48550/arXiv.2309.04658>
- Nie, Yixin, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2019. "Adversarial NLI: A New Benchmark for Natural Language Understanding." arXiv Preprint [arXiv:1910.14599](https://arxiv.org/abs/1910.14599). <https://arxiv.org/abs/1910.14599>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.