

The Dual Challenge from AI and the Cognitive Sciences for Law and Legal (Reasoning) Practices

ANTONIA WALTERMANN

Maastricht University

ABSTRACT

What do a man claiming that a brain tumour caused his paedophilic behaviour, an AI system called DABUS being recognised as an inventor for the purposes of a patent application, and the modelling of legal reasoning in computational form have in common? They are examples of developments in and increasing insights from the fields of artificial intelligence and the cognitive sciences that challenge extant legal (reasoning) practices in various ways. This paper proposes an analytical framework for the purposes of situating these different challenges vis-à-vis each other and our legal (reasoning) practices. The aim of this framework is to facilitate understanding how they relate to each other, what it is they really challenge, and to critically reflect on them.

KEYWORDS

practical reasoning, legal reasoning, legal concepts, cognitive sciences, artificial intelligence

CITATION

Waltermann A. 2023. *The Dual Challenge from AI and the Cognitive Sciences for Law and Legal (Reasoning) Practices*, in Brigaglia M., Roversi C. (eds.), *Legal Reasoning and Cognitive Science: Topics and Perspectives*, «Diritto & Questioni pubbliche», Special Publication, August 2023.

DIRITTO & QUESTIONI PUBBLICHE / RECOGNISE

Legal Reasoning and Cognitive Science: Topics and Perspectives
Special Publication / August, 2023, pp. 485-504

© 2023, *Diritto e questioni pubbliche*, Palermo.

ISSN 1825-0173



The Dual Challenge from AI and the Cognitive Sciences for Law and Legal (Reasoning) Practices

ANTONIA WALTERMANN

1. Introduction – 2. Theoretical building blocks & analytical framework – 2.1. Of and within a practice – 2.2. Mackor's four questions – 2.3. The building blocks combined – 3. Application – 3.1. Human cognition – 3.1.1. Brain tumours and paedophilia – 3.1.2. Does law undermine human flourishing? – 3.1.3. Hungry judges – 3.2. Artificial cognition – 3.2.1. DABUS, the AI inventor – 3.2.2. Can AI engage in legal reasoning? – 3.2.3. Legal responsibility for AI? – 3.3. First conclusions and some caveats – 4. How do we respond to these challenges? – 4.1. No inconsistency, no challenge – 4.2. Inconsistency, but no challenge – 4.3. Inconsistency & challenge – 4.3.1. Individual – 4.3.2. Systematic – 5. Conclusion

1. Introduction

What do a man claiming that a brain tumour caused his paedophilic behaviour, an AI system called DABUS being recognised as an inventor for the purposes of a patent application, and the modelling of legal reasoning in computational form have in common? They all *prima facie* challenge our legal (reasoning) practices (JONES 2013), albeit in seemingly very different ways. It may be useful, then, to be able to situate these and other challenges vis-à-vis each other and our legal (reasoning) practices to understand how they relate to each other, what it is they really challenge, and to critically reflect on them. That is what this chapter aims to do: the main claim of this chapter is that increasing insights into both human and artificial cognition challenge our understanding of legal reasoning as well as the core concepts we reason with. It proposes a framework that allows us to understand and situate these challenges, and briefly addresses how to respond to them.

To do this, the chapter proceeds as follows: following this introduction, section 2 offers the theoretical building blocks that are necessary for the analytical framework: it outlines the distinctions between human and artificial cognition, object- and meta-level or within and of a practice, and four ways in which extant legal (reasoning) practices can be challenged by insights from outside the law. It then combines these building blocks into an analytical framework within which different kinds of challenges from AI and the cognitive sciences can be situated—which section 3 then does. It considers several examples of challenges to our legal (reasoning) practices coming from the cognitive sciences and situates these within the analytical framework, thereby demonstrating a. the main claim of the paper and b., the way the framework works. Section 4, finally, considers different responses to these challenges, also touching on the question whether they really constitute challenges at all. The chapter concludes in section 5.

The value of this chapter does not lie in a systematic and substantive treatment of a particular field related to legal reasoning and the cognitive sciences. In fact, this chapter provides no such treatment at all. Instead, its value lies—ideally—in offering tools that allow readers to critically reflect on different challenges, to construct systematic, substantive treatments themselves, or to situate such treatments or challenges in a larger context. As such, the aim of this chapter is not primarily to make substantive claims (although it contains some), but to facilitate clear thinking.

* The author thanks Melina Poulin for her editorial and research assistance and the anonymous reviewer for the helpful comments. All remaining mistakes are, naturally, my own.

2. Theoretical building blocks & analytical framework

This chapter proposes a framework that can be used to identify, situate, and understand the different ways in which insights from the cognitive sciences broadly construed challenge law. To build this framework, we need a number of ‘theoretical building blocks’ without which our framework does not make sense. The building blocks that we need are:

- First, some terminological clarifications that set the stage for the analytical framework. This is predominantly the distinction between what I here call human and artificial cognition (section 2.1).
- Second, the distinctions between justifications (or concepts) within particular practices and outside (or of) the same practices (section 2.2).
- Third, four questions or challenges developed by Anne Ruth Mackor (section 2.3).

In section 2.4, we combine these building blocks into our analytical framework.

2.1. Human and artificial cognition

This chapter talks about a dual challenge from insights about human and artificial cognition: but what does this mean?

Cognitive science in general can be defined quite broadly as «the interdisciplinary study of mind and intelligence, embracing philosophy, psychology, artificial intelligence, neuroscience, linguistics, and anthropology» (THAGARD 2020). In this chapter, I distinguish between cognitive science that focuses on human cognition (e.g. the human mind, human intelligence, human reasoning, etc.) and artificial cognition (e.g. implementation of reasoning practices into artificial intelligence). I do this not because I think there are necessary differences between human and artificial cognition, but because the approaches of study and—often—the expertise of those in the field tends to differ (THAGARD 2020, 2). This does not take away that this distinction can be subsumed under the cognitive sciences more generally, or can be further refined.

2.2. Of and within a practice

Different scholars in different contexts have distinguished between different levels of abstraction with regard to different objects. In 1933, for example, the logician Alfred Tarski distinguished between object- and metalanguage for the purposes of establishing criteria for a truth definition: in Tarski’s model, object language is the language under discussion, while metalanguage is the language used to discuss the object language (HODGES 2018). HAMPTON (1997) writes about this distinction between levels of inquiry and analysis, which she holds can be used for other areas of inquiry as well:

«Consider the famous “liar’s paradox,” illustrated with the sentence, “This sentence is false.” The sentence cannot be true when it tells us it is false; but if it is false, then given the assertion it is making, it would seem to be true. Alfred Tarski resolved this paradox by distinguishing two kinds of language, which he called the “object language” and the “metalanguage.” The metalanguage is used to talk about the object language but is not itself part of that language. By understanding the predicates “is true” and “is false” to belong only to the metalanguage, we avoid the paradox» (HAMPTON 1997)¹.

¹ Note, however, that some authors argue—rightly, in my view—that an understanding of legal reasoning practices (our legal logic) can be better construed as not distinguishing between object- and meta-language (cf. HAGE 2005). Nonetheless, I think it is helpful for present purposes to make the distinction, without thereby wanting to make any

In 1955, the philosopher John Rawls distinguished between justifications within and justifications of a practice (RAWLS 1955, 3-32): with regard to the practice of punishment, he distinguished between applications of the practice, that is, the punishment of a particular individual, and the practice as such. In each case, one can ask about the justification: one might, first, ask about the justification within a practice, or, second, about the justification of a practice. What justifies the punishment of an individual person? Likely, one would here refer to the rules of the practice. What justifies the institution of punishment as such? Here, one cannot refer to the rules of the practice, but needs to look outside of the practice for justification.² In this connection, one can talk of justifications either internal and external (to the practice). This is also relevant in the next section.

For present purposes, we are not interested in establishing conditions for truth or in justifying punishment. Nonetheless, the *type* of distinction that both Tarski and Rawls make is useful for us. More precisely, it is useful for us to distinguish between two levels with regard to legal reasoning. These are:

- 1) the level of arguments we make within legal reasoning, including the conceptual building blocks of these arguments, such as LIABILITY, RESPONSIBILITY, GUILT, BLAMEWORTHINESS, TORT, CRIME, and much more.
- 2) the level of arguments we make about legal reasoning, including the building blocks we use to conceptualise these arguments, such as SYLLOGISM, VALIDITY, ANALOGY, or even LEGAL REASONING itself.

If we combine this distinction between the two levels with the distinction between human and artificial cognition of the previous section, we get the following:

	Within reasoning	Of/about reasoning
Human cognition		
Artificial cognition		

TABLE 1

The starting point of this paper is that insights from the cognitive sciences—or insights about human and artificial cognition, to use the distinction from 2.1—challenge both the concepts within our legal reasoning practices and of or about our legal reasoning practices.

2.3. Mackor's four questions

In 2013, Anne Ruth Mackor published a chapter asking what neuroscience can say about legal responsibility. In this chapter, she distinguishes between four questions, two internal to the practice of responsibility and two external to the practice. From least to most radical, these questions are:

- 1) whether the neurosciences have something to say about the proper application of conceptions of responsibility ('application');
- 2) whether findings of the neurosciences can change conceptions of responsibility (that is, the criteria for application of one or more concepts of responsibility) ('conception');

claims about the underlying (formal) logical structures of our reasoning.

² The picture is more complicated than this: different modes of justification are possible, such as conventionalist or coherentist. However, this is not the focus of this chapter, so we will not delve into it here.

- 3) whether neuroscientific research can show that responsibility practices are pointless to the extent that they cannot fulfil the goal they are intended to fulfil ('implication');
- 4) whether neuroscientific research can show that responsibility practices are fundamentally untenable because they rest on false presuppositions ('presuppositions') (MACKOR 2013).

The first of these question ('application') does not attack the practice or its concepts and their criteria of application in any way. It only asks whether our understanding of when the criteria for application are fulfilled is correct. In the context of neuroscience and responsibility, for example, this would be the question whether

«neurosciences can help to show that specific (categories of) persons whom we used to think that fit the criteria of application, in fact do not fit the criteria, or, conversely, that specific (categories of) persons whom we thought not to fit the criteria, turn out to fit the criteria after all» (MACKOR 2013).

The first question thus leaves the criteria for application intact. This is different with the second question ('conception'), which calls the criteria for application of a concept into question. Do the neurosciences, to remain with that example, change the criteria for application of one or more concepts of responsibility (MACKOR 2013)?

The third question ('implication') calls into question the practice as such, by asking whether the practice does not have implications that show it to be pointless, i.e., because it fails to contribute to the realisation of the goals or purposes of the practice. In the case of neuroscience and responsibility, the question is «whether neurosciences can offer evidence to support the claim that holding people responsible is pointless» (MACKOR 2013). The fourth question, similarly calls the practice as such into question by asking whether it rests on false presuppositions that make it fundamentally flawed or untenable (MACKOR 2013).

The first two of these questions are internal ones because they are situated within the responsibility practices in question; the latter two are external ones because they assess the practice from an external perspective. Abstracting from the context of neuroscience and responsibility respectively, we can draw up the following framework:

Internal	External
<p>Application Challenges to the application of concepts within the practice.</p>	<p>Implications Challenges our practice(s) by demonstrating that they cannot achieve their purpose</p>
<p>Conceptions Challenges the conceptions (applicability criteria) of concepts within the practice.</p>	<p>Presuppositions Challenges our practice(s) by demonstrating that they rest on wrong presuppositions and starting points</p>

TABLE 2

The practice to which these questions or challenges apply can in principle be any practice. For present purposes, we are looking at our legal reasoning practice and, following section 2.2, challenges from both human and artificial cognition to both the concepts within the practice of legal reasoning as well as of the practice of legal reasoning. In section 3, we will apply these different questions to cases, which will make them more concrete and serve to illustrate how our analytic framework functions. Before doing so, however, the different building blocks need to be put together.

2.4. The building blocks combined

If we combine Table 1 and Table 2, we get the following picture:

	Within reasoning		Of/about reasoning	
Human cognition	Internal Application	External Implications	Internal Application	External Implications
	Conceptions	Presuppositions	Conceptions	Presuppositions
Artificial cognition	Internal Application	External Implications	Internal Application	External Implications
	Conceptions	Presuppositions	Conceptions	Presuppositions

TABLE 3: COMBINED

Table 3 shows the combination of the dual challenge from human and artificial cognition to and within our legal reasoning practices and Mackor's four questions: insights from either human or artificial cognition might challenge the application, conceptions, presumed implications or presuppositions of our legal reasoning practices or the concepts with which we reason. This gives us sixteen possible combinations and thus, sixteen possible positions along which to classify different (kinds of) challenges. We have four quadrants: the first (blue in the table above) is where insights about human cognition challenge the concepts we use within our reasoning practices; the second (red) indicates challenges from human cognition to the concepts of and about our reasoning practices; the third (yellow) brings us to insights from artificial cognition challenging concepts within our reasoning practices; the fourth (green) indicates challenges from artificial cognition to our reasoning practices as such. Within each quadrant, we can distinguish between internal and external challenges: challenges to the application or conception of the concepts in question are internal; if, however, the challenge is that the implications or presuppositions of the concept are fundamentally flawed, it is external.

This gives us our analytical framework in full. But how does it look in action and how does it help? In the next section, we will apply this framework by considering different examples and situating them on the 'map' above. We will see that attempting to do so will make us engage with the examples critically.

3. Application

In this section, we will consider a number of examples to see where they are situated in the analytical framework developed in the previous section. In doing so, we will see that insights from human and artificial cognition challenge both the concepts and practices that play a role within our legal reasoning and the concepts and practice of legal reasoning as such (the main claim of the paper). We will also gain a better understanding of how the analytical framework functions.

3.1. *Human cognition*

In this section, we will consider examples of challenges to law that arise from increasing insights about human cognition, and situate these examples within our analytical framework. We will consider three examples: the challenge raised by brain tumours causing paedophilia for criminal liability and sentencing, the charge that law undermines human flourishing, and the oft-discussed study suggesting that the time since the last (food) break is more decisive than factors such as risk of recidivism when it comes to whether inmates would get parole.

3.1.1. *Brain tumours and paedophilia*

There have been, at the time of writing, two cases (one in the US and one in Italy) of men who developed uncontrollable paedophilia, which coincided with the growth of a brain tumour. Once the tumour was removed, the paedophilic urges and behaviour also ceased. In the US case, when the tumour regrew, the paedophilia also returned—and disappeared again after the tumour was removed once more.

What, if any, role should the insight that acquired paedophilia can be a consequence of cancer in the orbitofrontal region of the brain (BURNS, SWERDLOW 2003), play in sentencing of individuals? If we ask ourselves, for example, whether they should receive a reduced sentence, we are asking whether the criteria for application of a reduced sentence are fulfilled by a—in this case very specific—group of people who were not previously considered as fulfilling those criteria. That is an example of insights about human cognition challenging the application of concepts we use within our legal reasoning practices, i.e., it can be situated within the first cell ('application') of the first quadrant ('human cognition & within reasoning') of Table 3.

Similarly, consider the following argument in a news item published in response to the US case mentioned above:

«The U.S. Supreme Court has ruled that executing mentally retarded murderers is unconstitutionally cruel because of their diminished ability to reason and control their urges.

Chris Adams, a death penalty specialist for the National Association of Criminal Defense Lawyers, thinks the next logical step would be to include people who have brain tumors.

“Some people simply don't have the frontal lobe capacity to stop what they're doing,” he said» (NBC 2003).

This, too, can be situated in the first cell of the first quadrant of Table 3, i.e., it is an example of insights about human cognition challenging the application of the concepts with which we reason within the law. In essence, the argument is that (some) people with brain tumours satisfy the criteria that would make it unconstitutional to execute them, namely diminished ability to reason and control their urges.

Situating these questions and arguments in this way allows us to understand what precisely is—and is not—challenged.

3.1.2. *Does law undermine human flourishing?*

Alces and Sapolski, in a 2022 article, use the US case described above (calling the man in question by the pseudonym 'Mr Oft') and embed it in a larger and more encompassing argument. For them, it is one example of many neuroscientific insights that demonstrates that we are essentially—if complexly—mechanical entities:

«Our normative systems conceive of law and morality as [...] a product of sufficient choice to attach blame, fault, and concepts of desert. But on what basis do we draw the distinctions between physical and normative malady: Are not both just (generally) distinguishable manifestations of mechanical causes? If human agents are essentially mechanical entities, on what basis could we find a normative difference between, say, tuberculosis and selfishness or insufficient ability to feel compassion for others?» (ALCES & SAPOLSKY 2021-22).

They hold that «extant legal doctrine and practices (civil as well as criminal) actually undermine human thriving: they are not merely a distraction; they are an impediment» (ALCES & SAPOLSKY 2021-22, 1081). This is because our current legal doctrines and practices are «relying on a misconception of what it means to be human» (ALCES & SAPOLSKY 2021-22, 1084).

According to them, it

«is easy to see how Oft's tumor challenges a legal system equating intent with responsibility or volition. But it is not the sort of case that can be generalized easily enough to revolutionize legal thinking. This is because of the uniqueness of its clarity, where massively abnormal behaviour is caused by the singular and massive abnormality of a brain tumor, literally demonstrable at the scene of the crime. What contemporary science shows is that the intent behind our best and worst behaviors, and all those ambiguously in between, is as much the end product of factors outside our control as was Oft's intentional criminality. However, it is far harder to appreciate this than the case of Oft for at least three reasons: (1) unlike the singularity of his tumor, our behaviour mostly arises from a multitude of biological factors that subtly interact, (2) no single factor has remotely the overt sledgehammer causality of a tumor, and (3) many of the factors were set into action long before the behaviour occurred (with some even long before the individual in question was born» (ALCES & SAPOLSKY 2021-22, 1102).

According to Alces and Sapolsky, insights from contemporary science about human cognition and the causes of human behaviour conflict with the model of who we are as human beings that the law presupposes (ALCES & SAPOLSKY 2021-22, 1112). Applying our analytical framework to their argument, we can situate them in the fourth cell ('presupposition') of the first quadrant ('human cognition & within reasoning'): their challenge to the law is that it rests on wrong assumptions.

That is not the only challenge they raise, however: when they hold that the law, in making these wrong assumptions, undermines and actively impedes human thriving, this is another challenge, if one with an implicit premise. Implicitly, their argument suggests—and I think many would agree—that law does or should have as its aim to enable, rather than hinder, human thriving. If so, they raise a second challenge: that by relying on a misconceived view of humankind, law undermines the very purpose for which we have it. This thrust of their challenge can be situated in the second cell ('implication') of the first quadrant of our analytical framework.

3.1.3. *Hungry judges*

The previous two examples have both been situated in the first quadrant of our analytical framework, meaning that they have been examples of challenges from insights about human cognition to the concepts and practices that play a role within our legal reasoning. However, insights about human cognition also challenge our legal reasoning practices and the concepts we use to describe legal reasoning. This section is about one example of this.

A 2011 study by Danziger, Levav and Avnaim-Pesso has received much attention. The authors summarise their study as follows:

«We test the common caricature of realism that justice is “what the judge ate for breakfast” in sequential parole decisions made by experienced judges. We record the judges' two daily food breaks, which result in

segmenting the deliberations of the day into three distinct “decision sessions.” We find that the percentage of favorable rulings drops gradually from $\approx 65\%$ to nearly zero within each decision session and returns abruptly to $\approx 65\%$ after a break. Our findings suggest that judicial rulings can be swayed by extraneous variables that should have no bearing on legal decisions» (DANZIGER et al. 2011, 6889-6892).

Bublitz, in assessing this study, writes that:

«One intuitively shares their sentiment. Rulings have gone astray, they were influenced by something that ought not be there, something that affects or taints or contaminates decisions and renders them incorrect. Capricious rulings resulting from such extralegal factors appear as a paradigmatic instances of what it means to be ‘at the whim of a judge’. Or, in the words of the former US Supreme Court Justice William Douglas: there will be no “justice under law if a negligence rule is applied in the morning but not in the afternoon.” In other words, if justice is indeed what judges ate for breakfast, there is no justice» (BUBLITZ 2020, 6).

If this is correct, it provides a challenge to the core of our legal reasoning practice by demonstrating that core presuppositions of our legal reasoning practice—that judgments are rendered on the basis of legally relevant factors—are not borne out in reality. This can be understood as a challenge from human cognition to our legal reasoning practice, i.e., it can be situated in the fourth cell (‘presupposition’) of the second quadrant (‘human cognition & about reasoning’) of our table.

It bears mentioning, however, that Bublitz ultimately draws different conclusions in his chapter and that alternative explanations for the study have been offered (BUBLITZ 2020, 4),—which demonstrates the importance of critically assessing the challenges raised as well as the importance of replication. Whether and in how far the effect of the study and the conclusions drawn from it will bear out in future studies will remain to be seen, but at the same time, the challenge raised by this study should also not be underestimated.

3.2. *Artificial cognition*

In this section, we will consider examples of challenges to law that arise from increasing insights about artificial cognition, and situate these examples within our analytical framework. We will again consider three examples: the case of DABUS, an artificially intelligent system listed as sole inventor on a patent application; the question whether AI can engage in legal reasoning, particularly where legal reasoning is analogical; and the question whether AI can and should be considered an agent responsible for its acts in the eyes of the law.

3.2.1. *DABUS, the AI inventor*

In 2021, the Federal Court of Australia decided that an artificial intelligence system called DABUS could be considered an inventor for the purposes of an international patent application. Courts in other countries, such as the US, UK, and Germany, had rejected the patent application, as had Australian patent offices, but the Federal Court found that it was compatible with the Australian Patent Act to consider DABUS the inventor—but not owner—of the patent (MATULIONYTE 2021).

The Federal Court held that:

«for the following reasons, in my view an artificial intelligence system can be an inventor for the purposes of the Act. First, an inventor is an agent noun; an agent can be a person or thing that invents. Second, so to hold reflects the reality in terms of many otherwise patentable inventions

where it cannot sensibly be said that a human is the inventor. Third, nothing in the Act dictates the contrary conclusion»³.

The Australian Federal Court here applied the concept of inventor in a way it had not previously been applied, while courts of other countries did not do so: we could say that the existence of DABUS and the patent application have challenged how the concept of INVENTOR is applied, with different courts responding differently to this challenge. Where should this be situated on the analytical framework? Is this a case of application—where the criteria for application are not challenged—or one of conception—where the criteria for application are in question?

The criteria for application in the eyes of the court, it seems to me, was whether DABUS could be said to have invented something «that satisfies all of the requirements of patentability in terms of novelty, inventiveness and utility»⁴. What was decisive for the Federal Court was whether the act of inventing was performed, not whether it was performed by a human or non-human being. In this connection, the argument was that

«As an agent noun (like “computer”, “dishwasher” or “lawnmower”) the agent can be a person or a thing. In this context the primary judge noted that, whereas once the word “inventor”, like “computer”, might originally have been apt to describe persons when only humans could make inventions (or perform computations), now the term may be used to describe machines which can carry out the same function»⁵.

If the sole criterion is carrying out the function and technological progress makes it possible for an AI system to carry out the function, this could make the case of DABUS an example fitting under the first cell (‘application’) of the third quadrant (‘artificial cognition & within reasoning’).

However, much of the debate of the DABUS case is precisely about whether this is the correct view. Let us consider why the patent application was initially rejected: the Deputy Commissioner of Patents who rejected the application before the applicant (Dr Thaler), sought judicial review of the rejection and did so for the reason that the ordinary meaning of “inventor” was taken to be inherently human (CURREY & OWEN 2021). On this view, then, DABUS does not challenge ‘merely’ the application of INVENTOR but also the criteria for application: is it a criterion for the correct application of INVENTOR that the entity in question is a human being?

Table 4 gives the two different views in tabular form:

<u>Option 1</u>	<u>Option 2</u>
Criterion for correct application of INVENTOR	Criteria for correct application of INVENTOR
(1) Something was invented	(1) Something was invented
	(2) The invention was brought about by human action

TABLE 4

From this perspective, the case of DABUS challenges the criteria for application of the concept INVENTOR, not just the application. This view was upheld in appeal by the Full Federal Court

³ *Thaler v Commissioner of Patents*, 2021, Federal Court of Australia, sec 10.

⁴ *Thaler v Commissioner of Patents*, 2021, Federal Court of Australia, sec 7.

⁵ *Commissioner of Patents v Thaler*, 2022, Federal Court of Australia, sec 46.

which overturned the initial decision by the Federal Court (single judge)⁶. This demonstrates that what was at stake was not the application, but the conception of INVENTOR. That puts the DABUS case in the third cell ('conception') of the third quadrant ('artificial cognition & within reasoning').

The fact that on different construals, the DABUS case can be situated at different places in the framework demonstrates, in my view, the usefulness of the analytical framework: it demands critical reflection on what, precisely, is being challenged.

3.2.2. Can AI engage in legal reasoning?

In 2001, Cass Sunstein held that artificial intelligence cannot engage in legal reasoning. His argument rests, in part, on the state of technology at the time: the state of technology in 2001 was not (yet) far enough advanced that artificial intelligence could be said to be engaged in legal reasoning. However, he also held that claims, at the time, that AI could engage in legal reasoning rested on an «inadequate picture of what legal reasoning actually is» (SUNSTEIN 2001, 31). This second part of the argument is what I want to focus on here. Before I do so, two short notes are in order: first, Sunstein made this argument in 2001, so any claims about the state of the art of technology are outdated—but claims about legal reasoning are not, which is what makes this argument interesting for our purposes. Second, Sunstein's argument is US-centric and more obviously applicable to the common than to the civil law.

Sunstein argues that legal reasoning is often analogical. Reasoning by analogy, he posits, requires the following:

«[A]nalogizers in law have to ask which case has relevant similarities to the case at hand. It is more accurate still to say that whether a case has relevant similarities to the case at hand depends on the principle for which the initial case is said, on reflection, to stand. It follows that the crucial step in analogical reasoning consists, not in a finding of "more" similarities, not in establishing "many" distinctions, and not even showing "relevant" similarities and differences, but instead in the identification of a principle that justifies a claim of similarity or difference. Because the identification of that principle is a matter of evaluation, and not of finding or counting something, artificial intelligence is able to engage in analogical reasoning only to the extent that it is capable of making good evaluative judgments» (SUNSTEIN 2001, 5).

This means that «the analogizer attempts to make best constructive sense out of a past decision by generating a principle that best justifies it, and by bringing that principle to bear on the case at hand» (SUNSTEIN 2001, 7) and Sunstein saw no reason to think that the AI systems of the time had the capacities to do so. He explicitly leaves open the possibility for this to change, but at the time of writing, mechanisms of abstraction and analogy are still out of reach (PAVLUS 2021).

For our present purposes, it is interesting to see where Sunstein's argument can be situated in our framework. In my view, the argument can be framed as an argument against considering the concept of legal reasoning (in sense of including analogical reasoning) applicable to artificial intelligence, i.e., as responding to a challenge from AI to the application of the concept. Sunstein's argument is not that the criteria for application should or have changed, which would situate the argument under *Conception* but it is an argument that application to AI is not (yet) appropriate. One could, however, construe it also as an argument for a different conception of legal reasoning, certainly if compared to the conception of legal reasoning Sunstein ascribes to those who claim(ed at the time) that AI can engage in legal reasoning. This would make the argument an example of *Conception* rather than *Application*.

⁶ *Commissioner of Patents v Thaler*, 2022, Federal Court of Australia.

Bart Verheij made precisely that kind of argument in his 2020 presidential address to the seventeenth international conference on artificial intelligence and law when he held that some of the hurdles for the development of implementing legal reasoning in AI support the view that a particular model of legal reasoning—which he calls the subsumption model—is false:

«[a]ccording to the subsumption model of law there is a set of laws, thought of as rules, there are some facts,—and you arrive at the legal answers, the legal consequences by applying the rules to the facts [...]. The case facts are subsumed under the rules, providing the legal solution to the case. It is often associated with Montesquieu’s phrase of the judge as a ‘bouche de la loi’, the mouth of the law, according to which a judge is just the one who makes the law speak. All hurdles just mentioned show that this perspective cannot be true» (VERHEIJ 2020, 188).

Which are the hurdles that show this? Verheij lists the following: that legal reasoning is rule-guided, rather than rule governed; that legal terms have an open texture; that legal questions can have more than one answer, but still demand timely and reasonable responses; and that the answers to legal questions can change over time (VERHEIJ 2020, 187 f.).

Two years later, Francesconi hints at a similar challenge to our conception of legal reasoning when he writes that (far advanced) AI

«actually opens up the possibility that a machine, on the basis of deductive rules, facts and categories, can reach the levels of complexity of human legal reasoning, until replacing it. But this perspective is not without question marks. For example, does the human judge argue only by deductive categories? Moreover, which role have the emotions [sic] in taking decisions? Will a digital judge, emotionally neutral, be fairer than a human judge?» (FRANCESCONI, 2022, 157).

While challenges to models (that is, conceptions) of legal reasoning as subsumption or of legal formalism are not new (LEITER 2005), attempting to model legal reasoning in computational form offer support of such challenges or new challenges of their own. ASHLEY (2002) offers a plausible explanation of how and why this is the case when he writes that:

«the virtue of applying AI to research in legal or practical ethical reasoning “is that the nature of the subject forces additional explicitness and clarification, because ultimately its products must be encoded and run on a computer. Thus, hidden and unclear assumptions can often be exposed in such a context”» (ASHLEY 2002, 165, quoting SCHAFFNER 1990).

None of the above challenge the practice of legal reasoning as such, but raise questions about the application and conceptions of concepts of/within the practice.

3.2.3. *Legal responsibility for AI?*

With developments in the field of artificial intelligence, the question whether AI can and should be held morally and legally responsible has gained traction in academic and public debates⁷. I will here use only one argument within that debate as an example, and then only briefly: in 2017, Jaap Hage

⁷ To give only a number of examples: Open Letter to the European Commission Artificial Intelligence and Robotics (<http://www.robotics-openletter.eu/>); Report of COMEST on Robotics Ethics, 2017 (<https://unesco.blob.core.windows.net/pdf/UploadCKEditor/REPORT%20OF%20COMEST%20ON%20ROBOTICS%20ETHICS%2014.09.17.pdf>); ANDERSON & ANDERSON 2011; BROŽEK & JAKUBIEC 2017; BRYSON et al. 2017; CHOPRA & WHITE 2011; COECKELBERGH 2019; DAHIYAT 2021; FLORIDI et al. 2018; FLORIDI & SANDERS 2004; GUNKEL 2012; HIMMA 2009; SOLUM 1992.

has argued that it is possible and may be sensible to hold autonomous agents responsible, that is, legally liable. He summarises (part of) the argument as follows:

«it is argued [...] that agency, responsibility and liability are not found in a mind-independent reality, but rather are attributed to elements of a social practice that will be called the ‘practice of agency’. This practice may be based on the way human beings experience themselves and their fellow humans, but does not necessarily have a firm foundation in the ‘real’, mind-independent world. This practice might have been different from what it actually is and might attribute agency, responsibility and liability to autonomous systems just as easily as it actually attributes these characteristics to human beings. It is argued that a major reason to treat humans and autonomous systems differently in this respect—that humans act intentionally and on the basis of a free will—has lost much of its credibility in the light of modern science» (HAGE 2017, 256).

This argument is interesting for us to consider in light of the analytical framework of this chapter because Hage combines increased insights into artificial cognition and technological developments with insights into human cognition to build the argument. This shows the limitations of the framework of this chapter: not every challenge can be separated into neat analytical categories—rather, some challenges combine different categories. Nonetheless, we can ask what Hage ultimately challenges: his argument holds that our legal responsibility practice (‘within reasoning’ in our framework) rests on flawed presuppositions, namely «that human beings act intentionally and on the basis of a free will» (HAGE 2017). This makes Hage’s challenge comparable to that of Alces and Sapolsky (discussed in section 3.1.2), although Alces and Sapolsky are not concerned with artificial intelligence in their own argument. Nonetheless, they are sharing a view of humankind, and what it means for our responsibility practice.

3.3. *First conclusions and some caveats*

In the sections above, we have considered six examples of how insights from the cognitive sciences—about human or artificial cognition—challenge the law, both when it comes to our concepts and practices of legal reasoning and the concepts and practices within our legal reasoning. Before we consider how to respond to these challenges in the next section, it is worth taking a moment to draw some first conclusions and make some additional remarks. I think the six examples above demonstrate a. that the cognitive sciences challenge the law in different ways (the main substantive claim of this paper) and b. that these challenges can be fruitfully situated in the analytical framework outlined in this chapter.

It bears mentioning, however, that insights from the cognitive sciences (regarding both human and artificial cognition) are not alone or the first to challenge law in these ways. Over the centuries, if not millenia, many challenges that fit within the above theoretical framework have been put forth already⁸. What is new are not the challenges, but the (epistemic) support for these challenges and the means for teasing/implementing legal reasoning techniques in artificial cognition. Equally, the analytical framework of this chapter is not the only possible way in which one can conceptualise and understand these challenges (cf. HAGE 2021).

It also bears mentioning that the insights from the cognitive sciences are, by and large, descriptive or, to put it in a different way, belonging to the realm of theoretical reason. We can distinguish between theoretical and practical reason as follows:

⁸ MORSE (2015) holds, for example, that «[i]n principle, [...] neuroscience adds nothing new, even if neuroscience is a better, more persuasive science than some of its predecessors». And arguments about, for example, free will and determinism have existed long before the rise of neuroscience, cf. DILMAN 1999.

«The aim of theoretical reason is the description and explanation and sometimes also the prediction of events. The aim of practical reason is not to describe and to explain and thereby to understand states of affairs or events in the world, but to answer the question, “*What should I do?*”» (MACKOR 2013)⁹.

From the mere fact that something is the case, e.g. that a tumour causes paedophilic behaviour, we cannot derive any information about what should be done. This is also called *Hume’s guillotine* (HUME 1978): it is impossible to derive what ought to be done *only* from what is. In order to derive ought-judgments, we need a premise containing an ought¹⁰. Imagine, for example, a father telling his child that she ought to eat vegetables. Asked “why”, the father responds, “Because vegetables are healthy”. But this leaves silent the premise that the child ought to eat healthy things. Why ought she do so? Again, an argument is required—and again, this argument will require an ought-premise.

Hume’s guillotine has led some to conclude that challenges such as the ones outlined above are not a real threat to our legal practices, as these practices fall in the realm of practical, not theoretical reason (MACKOR 2013). This brings us to the last section of this chapter, which addresses possible responses to the challenges identified above.

4. *How do we respond to these challenges?*

So far, we have considered a number of examples of challenges from the cognitive sciences to our legal (reasoning) practices. We have also briefly touched on the idea that maybe, these challenges are not *really* challenges to our practices.

In this section, I want to address three different possible approaches to the question whether insights about human or artificial cognition really challenge our legal (reasoning) practices and if so, how we can respond. These approaches hold, respectively, that there is no challenge or inconsistency (section 4.1), that there is inconsistency, but no challenge (section 4.2), or that there is a challenge (section 4.3), in which case it needs to be addressed, which can be done on an individual (4.3.1) or a systematic basis (4.3.2).

4.1. *No inconsistency, no challenge*

One approach is to insist that there can be no inconsistency between our legal (reasoning) practices and insights about human or artificial cognition and therefore also no challenge, because our legal (reasoning) practices and insights about human or artificial cognition talk about different things. This might be because one belongs to the realm of practical and the other to the realm of theoretical reason, as we have seen above. It might also be because law can define its own concepts and as such, the cognitive sciences talk about different things than the law. For example, law determines what kinds of entities can count as INVENTOR in the eyes of the law, irrespective of whether we would regard these entities as capable of inventing things in non-legal terms. If artificial cognition develops to such a degree that we generally think that AI systems can invent things, this does not say anything about whether they can invent things in the eyes of the law. Similarly, from the ‘mere’ fact that artificial cognition

⁹ Note that variations of «What should I do?» such as what others should do and what one should have done fall under practical reason as well.

¹⁰ For a more nuanced explanation of Hume’s guillotine and what taking Hume’s guillotine for granted says about our ontological and epistemological beliefs (that is, beliefs about what exists and how we can have knowledge of it respectively), I refer readers to Jaap Hage’s chapter in this volume on the nature of law and constructivist facts.

has developed to such a degree or that we generally think AI systems can invent things one cannot derive that the law ought to recognise AI systems as inventors.

This view can be illustrated using the metaphor of maps¹¹. A map of Paris and a map of Rome cannot be inconsistent with one another and one cannot challenge the other, because these maps are not showing the same area at all. Analogically, if we take the cognitive sciences and our legal (reasoning) practices to be about different concepts or to belong to different realms altogether, we can say that there can be no inconsistencies between them and that the cognitive sciences cannot challenge our legal (reasoning) practices.

It is not obvious to me, however, that the cognitive sciences and our legal (reasoning) practices and the concepts used in each are maps showing different areas altogether¹². Moreover, even if we maintain a clear distinction between practical and theoretical reason, insights belonging to the realm of theoretical reason can be embedded into practical reason arguments.

4.2. *Inconsistency, but no challenge*

A different approach would be to say that the two are maps that show the same area, but in different ways: one map shows rivers and bodies of water, the other is a roadmap. If we could put one map on top of the other or combine the two into one integrated map without problems, there is no challenge. They might portray the same landscape in different ways, but the map showing a river in one spot is not challenged by another map showing a road running parallel to the river. The picture changes when one map puts a road in the middle of a large lake. This is problematic, at least if the map does not also show that there is a bridge across the lake. In this case, then, there seems to be a challenge—unless we do not assume that the two maps of the same area should be consistent with one another. Say that one map is historical and the other is current. In this case, we can notice the inconsistency between the maps, but the historical map does not challenge the accuracy of the current map if the lake has simply dried out.

What does it look like to translate this analogy back to insights about cognition and our legal (reasoning) practices? It may be the case that our legal (reasoning) practices rely on flawed presuppositions about human cognition, or that in doing so, they sometimes fail to realise their own goals. If we can nonetheless argue that in most cases, our extant legal (reasoning) practices are better than alternatives that are more consistent with insights about human cognition, we have an argument why our legal (reasoning) practices ought not be consistent with insights about human cognition. If we could argue, for example, that it is more just to maintain present practices despite flawed presuppositions, rather than to adapt them to more accurate presuppositions, or that maintaining present practices maximises happiness, we have such an argument. This is the kind of argument that Strawson (1962) makes when he holds that determinism cannot displace the reactive moral attitudes that are deeply ingrained in us and form the basis of our responsibility practices (STRAWSON 1962). In a similar vein, proponents of retributivism as a basis for punishment may hold that to reject retributivism in favour of consequentialist justifications of punishment is undesirable because it presupposes a view of human beings that does not treat them as agents in their own right or respect their dignity (e.g., MURPHY 2017).

¹¹ Jaap Hage and I have used this metaphor already elsewhere: HAGE & WALTERMANN 2021. It is initially taken from HAACK 2008.

¹² For a much more extended version of argument as applied to legal responsibility practices, see HAGE & WALTERMANN 2021.

4.3. *Inconsistency & challenge*

If it cannot be demonstrated that the cognitive sciences and our legal (reasoning) practices are a. about entirely different things or b. ought not be compatible with one another, we are left with the option to recognise that there is an inconsistency that needs to be addressed, that is, that there is a challenge from the cognitive sciences to law. To return to the map metaphor, in this scenario we are looking at two concurrent maps of the same area that show an inconsistency. We now need to ask ourselves how we can address the inconsistency. A likely way to do so is to change one of the maps to make the two maps consistent with one another so that they could, in theory, be combined into one integrated map without any issue.

What could this look like?

4.3.1. *Individual*

We can address challenges from the cognitive sciences on an individual basis. If the cognitive sciences show that under certain circumstances, brain tumours in the orbitofrontal region of the brain cause paedophilic behaviour, we can construct arguments about what this means for our practices of criminal law and criminal sentencing. If a study demonstrates that extra-legal factors such as time since the last meal break play a role in judicial decision-making, we can construct arguments about whether, why, and how to address this insight. For each challenge, we can construct arguments about what ought to change to address the challenge, that is, in order to bring our practice in line with the relevant insight about human or artificial cognition.

By definition, addressing a challenge on an individual basis will serve only to address that particular challenge. Particularly where the challenges from AI and the cognitive sciences challenge implications or presuppositions of law, however, a more integrated approach seems required.

4.3.2. *Systematic*

A more encompassing approach would be to develop a coherent theory of what our legal (reasoning) practices are, what they ought to be, and how insights from the cognitive sciences fit into this picture, and to use this theory as the backdrop of arguments about how to respond to new insights about human or artificial cognition that challenge our legal (reasoning) practices. This would be very demanding: such a theory would include not only state of the art knowledge about human and artificial cognition, but also encompass the aims of our legal (reasoning) practices, the values we pursue with them, an understanding of the nature of law and how it relates to the cognitive sciences, and more. Both the theory and our extant practices would then be revised on a continuous basis against the standard of coherence¹³.

5. *Conclusion*

In this chapter, we have considered a number of examples of ways in which insights about human or artificial cognition (seem to) challenge our legal (reasoning) practices. We have developed an analytical framework within which to situate these challenges in order to facilitate clear and critical thinking about them. This framework is summarily represented in the table below:

¹³ A lot more can and has been said about coherence. By way of example, HAGE 2013; LEHRER 1992; ZIPURSKY 1997.

	Within reasoning		Of/about reasoning	
Human cognition	Internal	External	Internal	External
	Application	Implications	Application	Implications
	Conceptions	Presuppositions	Conceptions	Presuppositions
Artificial cognition	Internal	External	Internal	External
	Application	Implications	Application	Implications
	Conceptions	Presuppositions	Conceptions	Presuppositions

We have also considered whether these seeming challenges really do challenge our legal (reasoning) practices. In this connection, we have considered a number of different possibilities, visualised in Figure 1:




FIGURE 1

Increasing insights from the cognitive sciences about cognition do and will challenge our legal (reasoning) practices, both when it comes to the concepts and practices we use within our reasoning practices and when it comes to the concepts and the practice of legal reasoning itself. While this chapter has addressed this in terms of challenges, these increasing insights also offer chances: in particular, they offer chances for a more encompassing, more nuanced, better—in the sense of more grounded in and coherent with best scientific theories about cognition—understanding of our legal (reasoning) practices. It is my hope that the analytical framework of this chapter will contribute to realising this better understanding.

References

- ALCES P.A., SAPOLSKY R.M. 2021-22. *Nohwere*, in «William & Mary Law Review», 63, 4, 1079 ff.
- ANDERSON S.L., ANDERSON M. 2011. *Machine Ethics*, Cambridge University Press.
- ASHLEY K.D. 2002. *An AI Model of Case-based Legal Argument from a Jurisprudential Viewpoint*, in «Artificial Intelligence and Law», 10, 163 ff.
- BECK S. 2015. *The Problem of Ascribing Legal Responsibility in the Case of Robotics*, in «AI & Society», 31, 473 ff.
- BROŹEK B., JAKUBIEC M. 2017. *On the Legal Responsibility of Autonomous Machines*, in «Artificial Intelligence and Law», 25, 293 ff.
- BRYSON J.J., DIAMANTIS M.E., GRANT T.D. 2017. *Of, for, and by the People: The Legal Lacuna of Synthetic Persons*, in «Artificial Intelligence and Law», 25, 273 ff.
- BURNS J.M., SWERDLOW R.H. 2003. *Right Orbitofrontal Tumour with Pedophilia Symptom and Constructional Apraxia Sign*, in «Archives of Neurology», 60, 3, 437 ff. Available on: <https://pubmed.ncbi.nlm.nih.gov/12633158/>.
- CHOPRA S., WHITE L. 2011. *A Legal Theory for Autonomous Artificial Agents*, University of Michigan Press.
- COECKELBERGH M. 2019. *Artificial Intelligence: Some Ethical Issues and Regulatory Challenges*, in «Technology and Regulation», 2019. Available on: <https://techreg.org/article/view/10999>.
- CURREY R., OWEN J. 2021. *In the Courts: Australian Court Finds AI Systems Can Be “Inventors.”*, in «Wipo Magazine», 3. Available on: https://www.wipo.int/wipo_magazine/en/2021/03/article_0006.html.
- DAHIYAT E.A.R. 2021. *Law and Software Agents: Are They “Agents” by the Way?*, in «Artificial Intelligence and Law», 29, 59 ff.
- DANZIGER S., LEVAV J., AVNAIM-PESSO L. 2011. *Extraneous Factors in Judicial Decisions*, in «Proceedings of the National Academy of Sciences», 108, 6889 ff.
- DILMAN I. 1999. *Free Will: An Historical and Philosophical Introduction*, Routledge.
- FLORIDI L., COWLS J., BELTRAMETTI M., CHATILA R., CHAZERAND P., DIGNUM V., LUETGE C., MADELIN R., PAGALLO U., ROSSI F., SCHAFER B., VALCKE P., VAYENA E. 2018. *AI 4People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations*, in «Minds and Machines», 28, 689 ff.
- FLORIDI L., SANDERS J. W. 2004. *On the Morality of Artificial Agents*, in «Minds and Machines», 14, 349 ff.
- FRANCESCONI E. 2022. *The Winter, the Summer and the Summer Dream of Artificial Intelligence in Law: Presidential Address to the 18th International Conference on Artificial Intelligence and Law*, in «Artificial Intelligence and Law», 30, 147 ff.
- GUNKEL D.J. 2012. *The Machine Question: Critical Perspectives on AI, Robots, and Ethics*, The MIT Press.
- HAACK S. 2008. *Putting Philosophy to Work: Inquiry and Its Place In Culture—Essays on Science, Religion, Law, Literature, and Life*, Prometheus Books.
- HAGE J. 2005. *Studies in Legal Logic*, Springer.
- HAGE J. 2013. *Three Kinds of Coherentism*, in ARASZKIEWICZ M., SAVELKA J. (eds.), *Coherence: Insights from Philosophy, Jurisprudence and Artificial Intelligence*, Springer.

- HAGE J. 2017. *Theoretical Foundations for the Responsibility of Autonomous Agents*, in «Artificial Intelligence and Law», 25, 255 ff.
- HAGE J. 2021. *Are the Cognitive Sciences Relevant for Law?* In HAGE J., BROZEK B., VINCENT N. (eds.), *Law and Mind: A Survey of Law and the Cognitive Sciences*, Cambridge University Press, 17 ff.
- HAGE J., WALTERMANN A. 2021. *Responsibility, Liability, and Retribution*, in HAGE J., BROZEK B., VINCENT N. (eds.), *Law and Mind: A Survey of Law and the Cognitive Sciences*, Cambridge University Press, 255 ff.
- HAMPTON J. 1997. *Political Philosophy*, Westview Press.
- HIMMA K.E. 2009. *Artificial Agency, Consciousness, and the Criteria for Moral Agency: What Properties Must an Artificial Agent Have to Be a Moral Agent?*, in «Ethics and Information Technology», 11, 19 ff.
- HODGES W. *Tarski's Truth Definitions*, in ZALTA E.N. (ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition), <https://plato.stanford.edu/archives/fall2018/entries/tarski-truth/>.
- HUME D. 1978. *A Treatise of Human Nature*, Selby-Bigge edition, Oxford University Press. (Originally published in 1739-40)
- JONES O.D. 2013. *Seven Ways Neuroscience Aids Law*, in «Vanderbilt Public Law Research Paper», 13-28. Available on: <https://ssrn.com/abstract=2280500>.
- LEHRER K. 1992. *Coherentism*, in DANCY J., SOSA E. (eds.), *A Companion to Epistemology*, Blackwell.
- LEITER B. 2005. *American Legal Realism*, in GOLDING M.P., EDMUNDSON W.A. (eds.), *The Blackwell Guide to the Philosophy of Law and Legal Theory*, Blackwell, 50 ff.
- MACKOR A.R. 2013. *What Can Neuroscience Say about Responsibility? Taking the Distinction between Theoretical and Practical Reason Seriously*, in VINCENT N. (ed.) *Neuroscience and Legal Responsibility*, Oxford University Press, 53 ff.
- MATULIONYTE R. 2021. *Australian Court Says that AI Can Be an Inventor: What Does It Mean for Authors?*, in «Kluwer Copyright Blog», 29 September 2021. Available on: <http://copyrightblog.kluweriplaw.com/2021/09/29/australian-court-says-that-ai-can-be-an-inventor-what-does-it-mean-for-authors/>.
- MURPHY J. 2017. *Retribution*, in LUNA E. (ed.), *Reforming Criminal Justice: Punishment, Incarceration, and Release. Volume 4*, Arizona State University.
- NBC NEWS 2003. *Pedophile Lost Urge after Surgery*, 1 November 2003 (Source: The Associated Press). Available on: <https://www.nbcnews.com/health/health-news/pedophile-lost-urge-after-surgery-flna1c9478663>.
- PAVLUS J. 2021. *The Computer Scientist Training AI to Think with Analogies*, in «Artificial Intelligence Quanta Magazine», 6 August 2021. Available on: <https://www.scientificamerican.com/article/the-computer-scientist-training-ai-to-think-with-analogies/>.
- RAWLS J. 1955. *Two Concepts of Rules*, in «Philosophical Review», 64, 1, 3 ff.
- SCHAFFNER, K. F. 1990. *Case-Based Reasoning in Law and Ethics*, Presentation at the Foundations of Bioethics Conference, Hastings Center.
- SOLUM L. B. 1992. *Legal Personhood for Artificial Intelligences*, in «North Carolina Law Review», 70, 4, 1231 ff.
- STRAWSON P.F. 1962. *Freedom and Resentment*, in «Proceedings of the British Academy», 48, 187 ff.
- SUNSTEIN C.R. 2001. *Of Artificial Intelligence and Legal Reasoning*, in «University of Chicago Law School», 8, 29 ff.

THAGARD P. *Cognitive Science*, in ZALTA E.N. (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2020 Edition), <https://plato.stanford.edu/archives/win2020/entries/cognitive-science/>.

VERHEIJ B. 2020. *Artificial Intelligence as Law: Presidential Address to the Seventeenth International Conference on Artificial Intelligence and Law*, in «Artificial Intelligence and Law», 28, 2, 181 ff.

ZIPURSKY B.C. 1997. *Legal Coherentism*, in «Southern Methodist University Law Review», 50, 5, 1679 ff.