

Neural Networks in Legal Theory

Vadim Verenich

Independent researcher
Tartu, Estonia

e-mail: vadimverenich@gmail.com
<https://orcid.org/0000-0002-2062-2267>

Abstract:

This article explores the domain of legal analysis and its methodologies, emphasising the significance of generalisation in legal systems. It discusses the process of generalisation in relation to legal concepts and the development of ideal concepts that form the foundation of law. The article examines the role of logical induction and its similarities with semantic generalisation, highlighting their importance in legal decision-making. It also critiques the formal-deductive approach in legal practice and advocates for more adaptable models, incorporating fuzzy logic, non-monotonic defeasible reasoning, and artificial intelligence. The potential application of neural networks, specifically deep learning algorithms, in legal theory is also discussed. The article discusses how neural networks encode legal knowledge in their synaptic connections, while the syllogistic model condenses legal information into axioms. The article also highlights how neural networks assimilate novel experiences and exhibit evolutionary progression, unlike the deductive model of law. Additionally, the article examines the historical and theoretical foundations of jurisprudence that align with the basic principles of neural networks. It delves into the statistical analysis of legal phenomena and theories that view legal development as an evolutionary process. The article then explores Friedrich Hayek's theory of law as an autonomous self-organising system and its compatibility with neural network models. It concludes by discussing the implications of Hayek's theory on the role of a lawyer and the precision of neural networks.

Keywords: legal analysis, generalisation, legal concepts, logical induction, semantic generalisation, formal-deductive approach, fuzzy logic, non-monotonic defeasible reasoning, artificial intelligence, machine learning, neural networks, deep learning algorithms, legal theory.

1. Introduction: Problems and Limitations of Formal, Logical and Mathematical Methods in Legal Analysis

The domain of legal analysis encompasses the generation, scrutiny, and application of law in relation to particular cases. Thus, legal analysis entails employing diverse analytical methodologies,

such as categorising factual information into specific legal frameworks, employing legal reasoning, and engaging in the decision-making process. These methodologies are fundamental to the systematic study of law and, furthermore, hold great significance from the vantage points of artificial intelligence and machine learning. The utilisation of formal methods not only facilitates the processes of reasoning and decision-making, but also allows for the conceptualisation of legal analysis as a comprehensive *holistic* undertaking. Indeed, in essence, the analysis of legal systems encounters similar fundamental inquiries. These inquiries pertain to the collection, interpretation, retention, and identification of patterns, and the response to incoming, legally relevant information. Each of these questions holds significant importance and is inherently interconnected with the others.

Every interaction with the domain of legal regulations or the realm of actual juridical relations results in the accumulation of experiential legal knowledge. If discernible regularities exist within this wealth of real legal experience, then they can be identified, scrutinised and subsequently utilised. The identification of such regularities implies the presence of “shared recursive patterns in legal forms and judicial opinions, which enables the use of process and technology to routinize and scale very cheap and very high quality solutions to the myriad of legal needs” (Henderson, 2013, p. 479). Consequently, the isolation of these shared entities is commonly referred to as *generalisation*. Generalisation represents a pivotal task across all disciplines associated with data analysis (mathematical statistics, machine learning, etc.). Naturally, the human brain does not remain exempt from this cognitive pursuit, as evidenced by our occasional observations of its adeptness in generalisation.

The classical position holds that legal norms are written in reference to those cognitively shared patterns (shared, uniform and inflexible properties with clear distinct boundaries) which define *juridical categories*. According to Lakoff:

The classical view that categories are based on shared properties is not entirely wrong [w]e often do categorise things on that basis[...] it has become clear that categorisation is far more complex than that [...] new theory of categorisation, called prototype theory, has emerged. [...] shows that human categorisation is based on principles that extend far beyond those envisioned in the classical theory (Lakoff, 1987, p. 5).

More recent studies in the field of cognitive science and brain neurophysiology have demonstrated that categorisation, the study of generalisable representations, is a type of decision making and that categorisation learning research would benefit from approaches developed to study the neuroscience of decision making and generalisation (Seger & Peterson 2013). The multitude of approaches to generalisation implies that the generalisation procedure lacks a universal framework. Despite the ubiquity of generalisation, the task itself, when considered in its broadest form, remains somewhat ambiguous. The formulation of the generalisation problem can vary extensively, depending on the specific context in which it is required. Different problem formulations engender diverse and sometimes disparate methods of solution.

Legal analysis encompasses the examination of semantic constructions of specific legal concepts, wherein ideas are expressed and documented through natural language expressions in legal texts.¹ The philosophical-semantic approach to generalisation can be outlined as follows: when there are interconnected concepts sharing a common generic attribute, it becomes necessary to transition towards a new concept that offers a broader, albeit less specific, interpretation by eliminating the generic attribute.

Philosophy is concerned with the examination and interpretation of semantic constructions. To illustrate this point, we can consider the paradigmatic instance of a rule, namely, “no vehicles in the park.” Here, the term “a vehicle” can be defined as a mechanical device typically equipped with wheels and an engine, utilised for the conveyance of individuals or goods, particularly on land. As colloquially, “a vehicle” is understood to refer specifically to an automobile, the underlying core

meaning of this rule can be construed as “no automobiles in the park.” However, if we eliminate the specific attributes of “wheels” and “an engine,” we arrive at the more generalised concept of “a vehicle,” encompassing any machine employed for transportation purposes. In the example, the very term “a vehicle” itself indicates the process of generalisation. By simply discarding the extraneous terms, we attain a broader and consequently more ambiguous “penumbral” concept, thereby enabling the application of the rule to bicycles or chariots (Hart, 1958).

The process of generalisation yields the development of legal concepts that are employed in constructing further descriptions. Different viewpoints exist regarding the fundamental principle governing the identification of specific concepts. All the items enumerated in this inventory are directly relevant to this matter. During the process of generalisation, we acquire concepts that encompass numerous phenomena encountered in some manner beforehand. By isolating commonalities among these phenomena, we are able to describe the properties of ideal concepts, which are detached from the specific details of individual occurrences. Through the process of generalisation, it becomes possible to convey the outcome of such generalisation using a systematic framework of concepts. In this scenario, the generalised concepts are not merely a collection of unrelated elements, but rather assume an inherent structure composed of interrelationships.

It is these ideal concepts that form the foundation of law. All legal concepts, like “liability”, “contract”, “tort”, “crime”, etc., are idealisations of objects derived from our everyday experiences. Law introduces a formal system of rules for these concepts, enabling the construction, interpretation, and application of these rules. However, while these concepts are primary for law itself, they are connected to human experiences of their application. Consequently, legal scholars can employ a more targeted legal inquiry based on the experiential background associated with the ideal concepts, using logical devices such as induction.

Logical induction involves deriving general laws from a collection of specific cases. In the case of complete induction, set A comprises the elements $A_1, A_2, A_3, \dots, A_n$. If A_1 possesses attribute B and A_2 possesses attribute B , then all elements from A_3 to A_n also possess attribute B . Consequently, all elements of set A possess attribute B . In a case of incomplete induction, set A comprises elements $A_1, A_2, A_3, \dots, A_n$. If A_1 possesses attribute B and A_2 possesses attribute B , then all elements from A_3 to A_k also possess attribute B . Consequently, it is likely that A_{k+1} and the remaining elements of set A possess attribute B (incomplete induction pertains to probability and can be fallible). Induction addresses generalisation in two ways. First, when referring to a set of objects, it implies that something has previously served as the basis for combining these objects into a unified set. In other words, a mechanism has been identified that facilitated the preceding generalisation. Second, through induction, if we discover a characteristic peculiar to the elements of a particular group that describes a specific concept, we can employ this characteristic as a criterion for categorising it within that group. Logical induction shares several similarities with the semantic generalisation of concepts. However, the semantic approach places a slightly different emphasis, focusing on the features comprising the description of a concept and the possibility of discarding certain features to obtain a more general formulation. Nonetheless, the question remains open, regarding the source of such concept definitions that enable the process of generalisation through the act of discarding. Incomplete logical induction elucidates the way that descriptive features are formed.

On other hand, it is obvious that the logical (either formal-deductive or inductive) approach to legal decision-making comes close to what Perelman and Olbrechts-Tyteca describe as:

a strategy which may be called logical, is that in which the primary concern is to resolve beforehand all the difficulties and problems which can arise in the most varied situations, which one tries to imagine, by applying the rules, laws, and norms one is accepting [...] The logical approach assumes that one can clarify sufficiently the ideas one uses, make sufficiently clear the rules one invokes, so that practical problems can be resolved without difficulty by the simple process of deduction. This implies moreover

that the unforeseen has been eliminated, that the future has been mastered, that all problems have become technically soluble (Perelman & Olbrechts-Tyteca, 1971, pp. 197-198).

In the initial stages of formalist advancements, a notable inclination toward strong symbolism emerged, promoting the acquisition of legal knowledge through the quasi-algebraic manipulation of symbols. These symbols served as representations that conveyed precise legal meanings or events, enabling the deduction of rules based on their manipulation. Within this context, reasoning was understood to be the systematic manipulation of existing legal knowledge, employing algebraic techniques, with the aim of deciding whether established legal knowledge (a norm or a norm's interpretation) applies in a particular case. Such manipulation encompasses exploring an algebraic space encompassing various potential solutions.

This approach has proved to be a significant source of inspiration for the development of legal applications, encompassing traditional computer programs (e.g. payroll systems or social security payments) as well as legal expert systems (e.g. databases of legal norms). The operation of such expert systems draws predictable conclusions from a predetermined initial set of norms, concepts, and facts, predominantly according to the so-called 'syllogistic model of adjudication'. There is much to be said for this approach in a number of legal contexts. It is an obvious, and even recommended, choice when standardisation and efficiency are paramount, when individual cases do not merit special adaptation or when the relevant legal rules can be easily defined, formalised, and updated. Even theoretical issues that can be addressed within the deductive model of law (such as the formalisation of *deontic* and *normative* conditions) should not be underestimated.

Nevertheless, this formal-deductive approach fails to capture a number of central, socially significant, and theoretically interesting phenomena of legal practice, which include the ability of jurisprudence to use insufficient or contradictory information, draw analogies, learn from examples and experiences in applying vague and imprecise rules, etc. From a perspective of classical legal positivism, it is usually asserted that empirical statements, concepts or terms within norms are not inherently vague but, rather, exhibit open-textured characteristics. Open-textured concepts of language refer to those concepts whose extension are not predetermined for all instances before their application, and yet they possess a certain procedure (that is of judicial decision-making) for determining their applicability in specific cases and extending those concepts to 'undecided cases' outside their standard domain of application (Hart, 1994, pp. 123, 128-136). Within the realm of law, this procedure is known as judicial decision-making, through which the courts render a ruling on a case.

Henceforth, a multitude of legal theorists have advocated for the adoption of a more adaptable formal model, underpinned by fuzzy logic and non-monotonic defeasible reasoning (Hage, 2005). In addition, another faction of scholars has pushed the boundaries further by augmenting this model with the distinct characteristics of artificial intelligence and machine learning.² These learning mechanisms enable the legal practice to serve as an effective problem-solving process in contentious cases, wherein the application of legal rules paradoxically necessitates the alteration and transformation of the rules themselves. The following section aims to demonstrate the potential utilisation of neural networks, a prominent deep learning algorithm, in addressing analytical tasks within the realm of legal theory.

2. Neural Networks: Formal Description

Neural networks can be described as neurally inspired computational tools for modelling neurological and cognitive processes. The capacity for an artificial neural network to effectively process and generalise information from previously unseen data is commonly referred to as generalisation. In Frank Rosenblatt's formulation, the concept of pure generalisation encompasses the following scenario: "In a learning experiment, a perceptron is typically exposed to a sequence of

patterns containing representatives of each type or class which is to be distinguished, and the appropriate choice of response is ‘reinforced’ according to some rule for memory modification. The perceptron is then presented with a test stimulus, and the probability of giving the appropriate response for the class of the stimulus is ascertained. Different results will be obtained, depending on whether or not the test stimulus is chosen to correspond identically to one of the patterns which were used in the training sequence. If the test stimulus is not identical to any of the training stimuli, the experiment is not testing ‘pure discrimination’, but involves generalisation as well. If the test stimulus activates a set of sensory elements which are entirely distinct from those which were activated in previous exposures to stimuli of the same class, the experiment is a test of ‘pure generalisation’. The simplest of perceptrons, which will be considered initially, “have no capability for pure generalisation, but can be shown to perform quite respectably in discrimination experiments, particularly if the test stimulus is nearly identical to one of the patterns previously experienced” (Rosenblatt, 1962, p. 68).




Figure 1. Perceptron diagrams. Source: Rosenblatt, F. *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Washington Spartan Books (Rosenblatt, 1962, p. 86).

A neural network can be described as a collection, characterised by a specific arrangement, of interconnected neurons (Haykin, 2006, p. 32). In this context, neurons are regarded as individual entities responsible for the reception and transmission of information. In isolation, neurons do not possess significant individual significance; their relevance lies solely within the interconnected network they form. Upon receiving incoming signals, a neuron assigns a specific weight to each of them. Subsequently, the signal is multiplied by its corresponding weight, the resulting values are aggregated, and a singular numerical value is generated. This resultant value is then passed on to the activation function, which determines whether the signal should be propagated further along the neural pathway (Haykin, 2006, pp. 42-44).

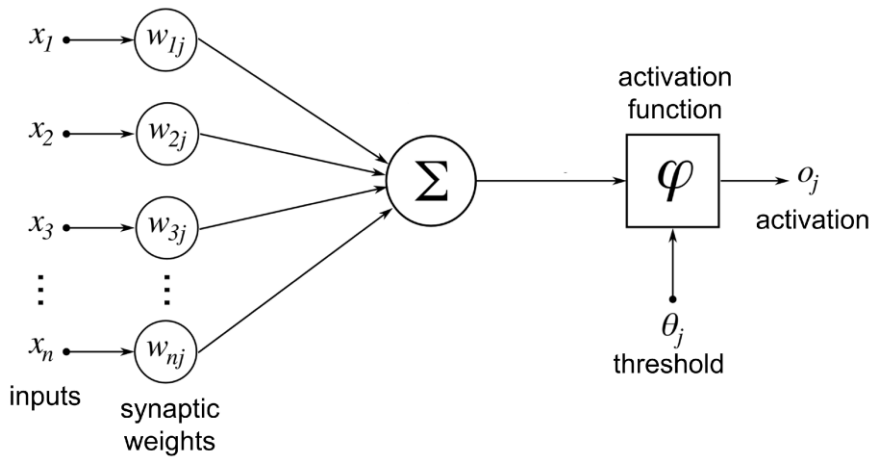


Figure 2. Diagram of an artificial neuron with n inputs and their corresponding synaptic weights. All weighted inputs are added and an activation function controls the generation of the output signal. Source: (Camuñas-Mesa et al., 2019).

An elementary neural network comprises three layers and facilitates unidirectional data transmission. The network encompasses input neurons, a concealed intermediary layer of neurons that remains imperceptible to external observation, and an output neuron.




Figure 3. The architecture under consideration is a fully-connected direct propagation neural network featuring a single hidden layer and a single output layer. Source: (EE 260, 2020).

Henceforth, it becomes evident that a neural network can be regarded as a form of mathematical function, essentially operating as a program. Instead of explicit programming, the neural network necessitates a process known as ‘training’ or adjustment. The process of training a neural network appears to be straightforward: by presenting a set of well-understood examples, we modify the coefficients of the underlying mathematical function, constituting the neural network, in adherence to specific rules, contingent upon the network’s responses, whether they are deemed to be correct or incorrect.

3. Neural Networks and Legal Theory

Neural networks inherently challenge fundamental principles of the syllogistic model, a widely employed formal approach in jurisprudence. This contradiction arises due to several key distinctions between these two paradigms:

- 1) In contrast to the syllogistic model, which condenses legal information into a collection of axioms, neural networks encode legal knowledge within the synaptic connections of their computational units.
- 2) The implementation of neural networks relies on reactive dispositions, whereas the syllogistic model operates on the basis of the logical validity of conclusions.
- 3) Neural networks possess the capability to assimilate novel experiences, whereas the deductive model of law can only be altered by the inclusion of overarching axioms.
- 4) Unlike intermittent updates commonly employed by the syllogistic model, neural networks exhibit evolutionary progression, enabling gradual refinement over time.

Obviously, these contradictions affect the two main trends in theoretical jurisprudence, namely natural law and legal positivism, which are united by the idea of law as a system of axioms (although with disagreements about the source and content of these axioms). Meanwhile, turning to the history of jurisprudence, we can find theories that are very close to the basic ideas underlying neural networks.

Statistical analysis of legal phenomena, developed within the framework of sociological and realistic approaches to the study of law, is the earliest and most obvious precedent for the use of ideas underlying the modern use of neural networks. Nevertheless, ideas close to the concept of neural networks are also found in other areas of legal research. In particular, we can recall a number of theories that, based on customary law, consider legal development as an evolutionary process, i.e. as a process of the selection and development of individual normative provisions, e.g. the historical school of Friedrich Carl von Savigny and Georg Friedrich Puchta in Germany, and the social philosophy of David Hume and Adam Smith in Scotland.

More recently, this approach to law has found its fullest formulation in the legal theory of the great economist Friedrich Hayek. First, Hayek offers us an unusual conception of the rule, which contrasts with the usual assumption that the rule is a linguistic entity, i.e. a statement or proposition. He views a rule as a special kind of disposition that “causes an organism to respond to stimuli of a certain class... with a response of a certain kind” (Hayek, 1977, p. 40). The imposition of numerous regulations (dispositions) upon a specific situation governs both our cognitive and practical behaviour. In order for these regulations to structure our experiences, it is not imperative that we possess a conscious awareness of them. Our unconscious cognitive tendencies are even more overarching and conceptual than our linguistic expressions. In fact, our language frequently proves inadequate in conveying the full extent of the mind’s capacity for considering the nature of required actions, and we often struggle to articulate, in words, what we inherently understand through practical knowledge. The intricate rules that govern our behaviour can only be acquired through emulation, whereby individuals learn to act in accordance with the same principles by imitating specific actions, although they can never assert those principles themselves.

Furthermore, this rule-based perspective of the human mind, as espoused by Hayek, is applicable to our sense of justice. Our ability to perceive the actions of others as meaningful and to evaluate our own or others’ actions as just or unjust must be grounded in the possession of highly abstract rules governing our behaviour, even if we remain oblivious to their existence and lack the means to articulate them verbally. The practical duty of a judge extends beyond adhering to these rules (thus safeguarding the expectations derived from them) to verbalising them in a manner accessible to the general public. This is a formidable task akin to the challenge of formulating scientific laws.

Lastly, Hayek presents us with an evolutionary model of the development of these rules (dispositions) in which they originate “from human action but not from human design.” According to Hayek, fundamental (and exceedingly complex) moral and legal principles do not arise from deliberate human choices but, rather, emerge unpredictably and spontaneously through social and cultural evolution (Hayek, 1976, p. 165). Evolution leads to the spread of behavioural dispositions that are best adapted through the persistence, expansion, and imitation of those groups that adopt them. Thus, the ‘best-adapted’ rules situation does not require anyone to know the reasons for the

success of those rules but depends only on the ‘greater success’ of the social order based on them. These reasons are implicitly included in the system of rules handed down by tradition, even when they are beyond the comprehension of individuals.

This leads us to Hayek’s central thesis, which posits that social order emerges as a result of an autonomous process of self-organisation governed by selective evolution. Consequently, the scope for legislator-initiated reforms within the realm of social order is relatively limited. While Hayek acknowledges the need for occasional improvements in established rules, he contends that the human mind is only capable of immanent criticism, which represents a constrained and partial attempt to enhance the internal coherence of the existing order. In other words, this form of criticism evaluates specific rules within a particular system by considering their consistency and compatibility with other recognised rules that shape a specific order of action. Since any established system of rules of conduct is founded on a partially understood body of experiential knowledge, and serves to guide actions in a manner that is only partially comprehended, it is impractical to aspire to its improvement (Hayek, 1976, p. 165). In this context, Hayek posits that the concept of consistency should not be misconstrued as logical consistency and it is unnecessary to reframe existing rules into a coherent set of axioms. Rather, achieving such consistency can be better accomplished by assigning priority to conflicting rules and establishing criteria for resolving and eliminating conflicts.

To encapsulate the main theses of Hayek’s theory in relation to the fundamental components of the neural network model, it is worth noting that Hayek’s proposition that norms are derived from “learning by example” aligns with the learning strategy commonly observed in neural networks. Moreover, the notion of the legal system as an autonomous self-organising system corresponds to the phenomenon of self-organisation exhibited by neural networks. Lastly, the task of a lawyer, as Hayek sees it (the arrangement and harmonisation of normative material), corresponds to the process of enhancing the precision of neural networks by generating and refining a precise formal representation of input data.

4. The Problem and the Critique

It is imperative to enumerate several foundational, unresolved inquiries in legal theory, which manifest as technological predicaments in the realm of neural networks. **The primary concern** pertains to the correlation between the ‘context of discovery’ and the ‘context of justification’, i.e. the manner in which a legal solution is attained and subsequently substantiated through appropriate reasoning. In neural networks, solutions are not derived through the formulation of reasons but through the application of analogies to previous instances of successful problem-solving (the notion of neural reasoning as an analogy is challenged in Dan Hunter’s (1994) article). This conveys the predicament of whether and how such solutions ought to be justified. Should some form of justification for the outcomes produced by neural networks not exist? Should it be a retrospective opportunistic rationalisation? (This notion finds support in the works of Andrew Stranieri and John Zeleznikow (2005), and Mark Gawler and Bryn Lewis (Stranieri et al., 1999), who acknowledge the realist perspective of perceiving decision motivation as a form of rationalisation.) If there is a genuine attempt to derive a coherent set of principles from the problem-solving patterns of the network, what are the principles that can influence problem-solving behaviour and even modify its ‘unconscious’ patterns? How does legal reasoning integrate symbolic and sub-symbolic processes and how can the deliberate utilisation of refined conceptual structures interact with the unconscious activation of parallel connections?

In 1994, Dan Hunter emphasised that one of the primary challenges with neural networks is their inability to comprehend the rationales underlying their decisions (Hunter 1994). Consequently, in the event of conflicts, the system will be incapable of providing a logical justification for its conclusions. Although it is presently feasible to partially comprehend certain conclusions through

reverse engineering and the algorithmic reporting process, understanding the mechanisms behind machine-generated results necessitates the development of diverse approaches.

Neural networks possess an inherent proficiency in pattern classification, rendering them seemingly promising candidates for emulating analogical reasoning processes. Hobson and Slee explored the utilisation of artificial neural networks to emulate this facet of reasoning in their work. They constructed a neural network ‘index’ of the 1968 Theft Act (England) (Hobson & Slee, 1994). In this index, the researchers analysed the factual circumstances to determine the presence or absence of various concepts, as defined by the wording of the Theft Act. The presence or absence of each concept was represented in the form of a state matrix, which subsequently served as input for their neural network. The ultimate verdict on whether a given situation constituted theft within the confines of the Act was employed as the desired outcome for the neural network. Based on this material, Hobson and Slee argued that a neural network could be trained to classify cases falling under the purview of the law. During the training process, the neural network autonomously groups the cases utilised for training into shared categories. Following training, new cases can be presented to the neural network. The network then classifies the cases into the general groups established during training to reach a verdict. Through this classification process, the neural network seemingly emulates similar reasoning, as comparable cases yield congruent verdicts.

The second theoretical inquiry pertains to the acquisition of legal knowledge. It has been observed that neural networks employ an example-based learning paradigm, rather than one based on legal theory. This approach to legal learning is not novel to legal professionals, as cases have always played a pivotal role in the instruction of common law. Even lawyers practicing within continental legal systems are progressively recognising the significance of case precedents (as some have historically done, exemplified by Friedrich Carl von Savigny’s recognition that mastery of Roman casuistry is essential for comprehending law). A fundamental question arises: must the selected learning patterns be pristine prototypes? Easily discoverable learning patterns that address specific problems can be found in textbooks. However, an alternative approach involves utilising all available court cases to train the neural network. Nevertheless, this approach carries the risk of incorporating examples shaped by secondary problems that may ‘contaminate’ the main problem.

An experiment conducted by Filipe Borges, Raoul Borges and Daniele Bourcier (2003) sought to model certain aspects of the French penal code using a neural network. A cursory examination of the hidden neurons’ activity revealed the emergence of distinct ‘tendencies’ or preferences in data processing. Notably, some neurons focused specifically on ‘murders’, while others focused on ‘sexual offences’, and others displayed either stoical or hyperactive behaviours. Despite such a latent specialisation of individual neurons, the overall decisions made by the neural network remained entirely pertinent. The development of specialised functions within the hidden neurons suggested the possibility of refining the representation of the legal dispute model within the neural network.

The **third theoretical concern** pertains to the concept of ‘self-organisation’, which holds significance, not only in the realm of neural networks, but also in diverse disciplines, such as biology, systems theory, economics, and computer science. Within the field of contemporary sociology of law, this concept is frequently invoked to construct abstract and comprehensive frameworks for understanding the development of legal systems, drawing upon ideas of evolution and self-organisation (as exemplified by Luhmann’s legal sociology). However, legal research has thus far encountered challenges in offering precise and, potentially, controlled explanations of specific facets of law based on self-organisation theories.

An argument has been posited suggesting that neural networks are incapable of modelling the process of legal decision-making due to their inability to apply norms. However, this claim is subject to debate. If the assertion implies that neural networks are unable to apply norms because of their normative content, it is inaccurate. If norms can be expressed in the form of cases or rules, a neural network can be employed to model them. In this context, the normative content within these cases or rules becomes irrelevant. In fact, the fundamental functioning of neural networks can be

viewed as the application of a norm, mandating that similar cases be resolved in a consistent manner. If the implication is that norms cannot be expressed in terms of cases but necessitate representation as rules, it is still premature to assume that neural networks cannot model legal decisions. There remains a possibility that neural networks can be utilised to model norms.

Moreover, the contention that neural networks are unable to apply norms because they lack normative content within the neural network itself is also a matter of debate. This question is intertwined with the broader inquiry of whether neural networks and computers can engage in thinking, which, though beyond the scope of this article, remains an open question. Nonetheless, it may be valid to assert that the outcomes generated by a neural network cannot determine value-based decisions. Simultaneously, when examining the classification of legal systems into continental and Anglo-Saxon (which constitute the two fundamental models today), the potential for employing neural networks in law enforcement practice can be assessed as follows.

In the continental legal system, the approach is apparent. Since the continental criminal process can be formally categorised into pre-trial and trial proceedings, the utilisation of a neural network to make decisions regarding the termination of criminal proceedings or sentencing presents certain challenges. This is due to the fact that the entire process is characterised by police (investigative) procedures and the introduction of neural network-derived sources as evidentiary materials may hinder the adherence to cornerstone principles. The Anglo-Saxon legal system is not entirely straightforward either. On one hand, it is a system grounded in fundamental principles concerning the safeguarding of human rights (e.g. the Human Rights Act of 1998), which precludes the possibility of employing a neural network. But, on the other hand, it is a precedent-based system that, from the perspective of the Borges theory mentioned earlier, can accommodate the utilisation of various electronic systems in practical applications. For instance, one could consider a form of collective program capable of incorporating precedent norms (thus serving as a custodian of legal information) that could be employed by demonstrating competence in reactive dispositions and leaving room for evolutionary development. Overall, these considerations shed light on the potential applications of neural networks in legal contexts, taking into account the distinctive features and requirements of different legal systems.

References

- Borges, F., Borges, R., & Bourcier, D. (2003). Artificial neural networks and legal categorisation, In *The 16th Annual Conference on Legal Knowledge and Information Systems (JURIX'03)*, The Netherlands, 11-12 December 2003, pp. 11-21.
- Camuñas-Mesa, L. A., Linares-Barranco, B., & Serrano-Go Gotarredona, T. (2019). Neuromorphic Spiking Neural Networks and Their Memristor-CMOS, *Hardware Implementations. Materials* 12, p. 2745.
- EE 260 (Spring 2020). Advanced VLSI Design for Machine Learning and AI. Available at: <https://vsclab.ece.ucr.edu/courses/2019/12/01/ee-260-spring-2020-advanced-vlsi-design-machine-learning-and-ai>.
- Hage, J. C. (2005). *Studies in Legal Logic*, Dordrecht: Springer.
- Hart, H. L. A. (1958). Positivism and the Separation of Law and Morals, *Harvard Law Review* 71/4, pp. 593–629.
- Hart, H. L. A. (1994). *The Concept of Law*, 2nd edition, Oxford: Oxford Press.
- Hayek, F. A. (1976). *Law Legislation and Liberty. Volume II. The Mirage of Social Justice*, London: Routledge and Kegan Paul.
- Hayek, F. A. (1977). *New Studies in Philosophy, Politics, Economics and the History of Ideas*, London: Routledge and Kegan Paul.
- Haykin, S. (2006). *Neural networks: a complete course*, Moscow: William Publishing.
- Henderson, W. D. (2013). A Blueprint for Change. *The Pepperdine Law Review* 40/2, Available at: <https://digitalcommons.pepperdine.edu/plr/vol40/iss2/7>.

- Hobson, J. B., & Slee, D. (1994). Indexing the Theft Act 1968 for case based reasoning and artificial neural networks, *In Proceedings of the Fourth National Conference on Law, Computers and Artificial Intelligence, unnumbered additions*, Exeter: Exeter University Centre for Legal Interdisciplinary Development.
- Hunter, D. (1994). Looking for law in all the wrong places: Legal theory and legal neural networks, In H. Prakken, A. J. Muntjewerff, A. Soeteman, and R. Winkels (eds.), *Legal Knowledge Based Systems: The Relation with Legal Theory*, Lelystad: Koninklijke Vermande, pp. 55–64.
- Lakoff, G. (1987). *Women, Fire And Dangerous Things: What Categories Tell Us About The Nature Of Thought*, Chicago: University of Chicago Press.
- Perelman, Ch. & Olbrechts-Tyteca, L. (1971). *The New Rhetoric: A Treatise on Argumentation*, Paris: Notre Dame Press.
- Rosenblatt, F. (1962). *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*, Washington: Spartan Books.
- Seger, C. A. & Peterson, E. J. (2013). Categorisation=Decision Making + Generalisation, *Neuroscience & Biobehavioural Reviews* 37/7, pp. 1187–1200.
- Stranieri, A., & Zeleznikow, J. (2005). Data Mining Using Neural Networks, In *Knowledge Discovery from Legal Databases. Law and Philosophy Library, vol. 69*, Dordrecht: Springer.
- Stranieri, A., Zeleznikow, J., Gawler, M., & Bryn, L. (1999). A Hybrid Rule – Neural Approach For The Automation of Legal Reasoning in The Discretionary Domain of Family Law in Australia, *Artificial Intelligence and Law* 7/2-3, pp. 153-183.

Notes

1. This semantic construction of concept should not be confused with statutory construction, i.e. the activity of translating the semantics of a legal text into legal rules.

2. Technically, one can consider machine learning to be a type of artificial intelligence, working by identifying patterns in data and then applying a learned model to new data.