Artificial Intelligence and Law **7:** 115–128, 1999.

Introduction: From legal theories to neural networks and fuzzy reasoning

LOTHAR PHILIPPS
Institute of Philosophy of Law and Computers and Law Ludwig-Maximilians-University, Professor-Huber-Platz 2, D-80539 München

GIOVANNI SARTOR
Law Faculty, Queen's University of Belfast, BT7 1NN, Northern Ireland

Abstract. Computational approaches to the law have frequently been characterized as being formalistic implementations of the syllogistic model of legal cognition: using insufficient or contradictory data, making analogies, learning through examples and experiences, applying vague and imprecise standards. We argue that, on the contrary, studies on neural networks and fuzzy reasoning show how AI & law research can go beyond syllogism, and, in doing that, can provide substantial contributions to the law.

Key words: analogy, fuzzy logic, learning, legal formalism, neural networks, vagueness

1. Mechanism and Formalism in the Law

Laypersons, and even lawyers, are likely to associate computational approaches to the law with a cluster of four adjectives: mechanical, deductive, formalistic, and formal. Those adjectives appear to be inextricably conjoined, and are indeed frequently used interchangeably. A legal practice is supposed to be approachable by computational methods and techniques (mechanical) as far as it is reducible to the valid derivation from a set of existing axioms (deductive). This, in turn, seems to presuppose that this practice is only concerned with the law as an autonomous "form", as a set of concepts and rules to be identified according to content-independent (i.e., "formal") criteria and to be construed and applied disregarding any psychological, ethical and social consideration (formalistic). A further feature is added to this cluster by linking those notions of "formalistic" and "formal" to the use of formal languages and logical-mathematical methods.

This assimilation of the mechanical, the deductive, the formalistic, and the formal pre-dates the invention of computers. Let us recall the criticisms that socially-oriented jurisprudes (from Interessenjurisprudenz, to Freirechtslehre, to sociological jurisprudence) brought against their "conceptualist" colleagues at the beginning of this century. The conceptualists, in their attempt to ensure the autonomy of the law and to insulate it (formalism) against moral judgements and political evaluations tried to limit legal decision-making to the application of previously defined rules and concepts (deduction). This view was rejected by the "socio-legal" scholars as being both descriptively inadequate (since legal decision often go beyond prima-facie interpretations of "formal" rules) and normatively incorrect (since lawyers should take an active role in social change). Anti-formalists used the word "mechanical" as a disparaging label evoking the rigid narrow-mindedness promoted by their academic adversaries. (The idea that a formalist judge would function as a machine occurred also to other thinkers of that time, equally impressed by the machinery of the industrial revolution, such as Max Weber.)

The term "mechanical jurisprudence" became, therefore, a metaphorical weapon in the struggle of legal methods (Methodenstreit) at the beginning of this century, when computers where still to be invented and even formal logic was at its beginnings. This use determined the cluster of concepts associated with the notion of "mechanical jurisprudence".

Today this notion is no more a pure metaphor: it has acquired a literal denotation, consisting of many legally-oriented computer programs and the corresponding computational models and theories. Can we still take for granted the idea that there is a necessary link between the four adjectives mentioned above (mechanical, deductive, formalistic and formal)? The use of computers (as machines guided by formal languages and mathematical models) seems to have made even stronger the link between the mechanical and the formal, but should a mechanical approach to the law still necessarily be a formalistic and a deductive one?

Let us first observe that a formalistic-deductive approach to legal decision-making comes close to what Perelman and Olbrechts-Tyteca (1969, pp. 197–8), critically describe as the "logical" strategy of problem solving (cf. Prakken 1997, p.19, who call this approach "naïve deductivist view"):

> The ... [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.

Such a strategy seems indeed to inspire many legal applications, both in conventional computer programs (e.g., payroll systems or calculations of social security entitlements) and in legal expert systems (e.g., in deductive databases of legal rules). The operation of those systems consists in the derivation of the consequences of a pre-determined stock of norms, concepts and facts, as required by the so-called syllogistic model of adjudication. There is much to be said in favour of such an approach in a number of legal contexts. It is the obvious and recommendable choice where standardisation and efficiency are of paramount importance, where individual cases do not deserve special adaptations, where the relevant legal rules can be easily identified, formalised, updated. Even the theoretical issues which may be addressed within a deductive model of the law, such as those of the formalisation of deontic and normative modalities, should not be underestimated.

However, such an approach falls short of capturing a number of central aspects of legal practice, which contribute to making it both socially important and theoretically interesting: its capacity to make use of insufficient or contradictory information, of performing analogies, of learning through examples and experience, of applying vague and imprecise standards, etc. It is thanks to these performances that this practice can be an effective problem-solving process in controversial cases, where paradoxically the legal order can only be applied by contributing to its development.

If the adjectives mechanical, deductive, formalistic and formal were an indivisible cluster (representing different facets of a unified "logical" approach to the law), this would imply a serious limitation to the objectives of AI & law research, both at the application level and at the theoretical one. Our discipline being both mechanical and formal (since it intends to model aspects of the law through computable formalisms) would be condemned to be similarly deductive and formalistic. AI & law could not contribute to, nor derive any suggestion from, non-formalistic views of the law and of legal reasoning and could not address those aspects of the legal practice just mentioned.

We hope that this special issue will contribute to show that there is no necessary connection between the mechanical and the formal, on the one hand, and the deductive and the formalistic, on the other. We believe that computable formalisms may successfully address various non-deductive and non-formalistic aspects of legal reasoning. In this way, they can both draw inspiration from, and contribute to the advance of, legal theories. We will not here describe all various directions in which AI & law has explored this perspective, by applying, adapting and refining different AI techniques (non-monotonic reasoning, argumentation, case-based reasoning, automatic induction, data mining, genetic algorithms, etc.). We will focus just on two formal techniques, neural networks and fuzzy logic, which most clearly exemplify this line of research.

2. Neural Networks

Neural networks contradict some basic tenets of the syllogistic model. A neural network stores legal knowledge in neuronal connections (rather then in a set of axioms), exhibits it competence in reactive dispositions (rather than in valid inferences), learns though experience (rather than through the acquisition of general axioms), and develops evolutionarily (rather than by way of discontinuous updates).

Such an approach to the law conflicts with both principal traditions of jurisprudence, that is natural law and legal positivism, which are united by the view of the law as an axiom system (although disagreeing on the source and the content of these axioms). In the history of jurisprudence, however, we can also find some theories that closely fit the basic ideas underlying neural networks. The statistical accounts of legal phenomena, developed in the framework of sociological and realist approaches to the law, is the first and most obvious precedent for the current use of neural networks. However, further links can be identified with other strands of legal research. Here we consider in particular those theories that focus on customary law and view legal development as an evolutionary process, which selects and spreads normative dispositions. Among them we may recall, for example, 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 (not to speak of the sparse but brilliant reflections of Edmund Burke in Ireland). More recently this approach to the law has found an admirable synthesis and a most comprehensive formulation in the legal theory of the great economist Friedrich Hayek, a few aspects of whose work we shall briefly recall.

Firstly, Hayek offers us a concept of a rule that contrasts with the usual assumption that a rule is a linguistic entity, i.e., a statement, a sentence or a proposition. He views a rule as a disposition or propensity which "makes an organism respond to stimuli of a certain class ... by a response of a certain kind" (Hayek 1977, p. 40). The superimposition of many rules (dispositions) govern both our cognitive and practical behaviour, "both the specification of a particular experienced event, and the specification of a particular response to it". For such rules to constitute our mind and organise our experience it is not necessary that we are conscious of them. We are indeed reminded that our unconscious dispositions are even more general and abstract than our linguistic formulations. Indeed, our "language is often insufficient to express what the mind is fully capable of taking into account in determining action" and "we will often not be able to communicate in words what we well know how to practice". Therefore, the highly complex rules which govern our behaviour can only be learned through example, i.e., "by imitating particular actions ... the individual acquires by analogy the capacity to act in other cases on the same principles which, however, he could never state as principles".

Secondly, this view of human rule-based intelligence is also applied by Hayek to our sense of justice: "our capacity to recognise other people's actions as meaningful, and the capacity to judge actions of our own or of others as just or unjust, must be based on the possession of highly abstract rules governing our actions, although we are not aware of their existence and even less capable of articulating them in words". The task of the judge, according to Hayek, is to act according to those rules (protecting the expectations based upon them), but also to articulate them in words. This is a difficult and fallible task, which can be assimilated to the problem of formulating scientific laws.

Thirdly, Hayek provides us with an evolutionary model of the development of those rules (dispositions), where they result "from human action but not from human design". In Hayek's account the basic (and more complex) moral and legal rules are not the result of a deliberate human choice, but emerge unpredictably from social and cultural evolution (cf. Hayek 1979, 165 ff). Evolution leads to the spreading of the behavioural dispositions which are best adapted, through the persistency, the expansion and the imitation of those groups which adopt them. The emergency of the best adapted rules, therefore, does not require anybody's knowledge of the reasons explaining the success of those rules, but only the superior performances of the social order based upon them. Those reasons survive encoded, incorporated in the system of rules transmitted by tradition, even though they are not accessible to any individual intellect.

This leads us to the fourth and politically most controversial aspect of Hayek's model, that is the idea that, since the social order is the result of an autonomous process of self-organisation, governed by selective evolution, there is little space for deliberate reform. Hayek admits that often a refinement of the traditional rules is needed, so that it is worthwhile to try to articulate them in language. Nevertheless, what human reason can achieve is only an immanent criticism, i.e., a limited and piecemeal attempt to increase the internal consistency of the existing order (i.e., the compatibility of the actions that it governs).

> Since any established system of rules of conduct will be based on experience which we only partly know, and will serve an order of action in a manner which we only partially understand, we cannot hope to improve it by reconstructing anew the whole of it . . . . . We shall call 'immanent criticism' this sort of criticism which moves within a given system of rules and judges particular rules in terms of their consistency or compatibility with all other recognised rules in inducing the formation of a certain kind of order of actions. (Hayek 1976, p. 24)

In this connection, Hayek observes that such consistency is not to be understood as logical non-contradiction and that there is no need to rephrase the existing rules into a consistent axiomatic base. Rather, such consistency may better be achieved by prioritising the conflicting rules, and providing criteria for solving or avoiding their conflicts. The conscious manipulation of explicit rules does not imply the elimination of their sub-symbolic counterparts. These continue to provide the background which supports, integrates and validates linguistic abstractions.

Let us now move from this jurisprudential discussion to its technological counterpart, by briefly examining some features of neural networks (a detailed example, which complements the following description is provided in the appendix). A neural network links "input cells" to "output cells" (cf. Figure 1). In a legal interpretation of such a neural network "input cells" represent the norm's conditions, while "output cells" represent possible legal consequences. Between input and output units there are usually "hidden units". (It has not yet been clarified whether such hidden units can be viewed in legal applications as intermediate terms of a

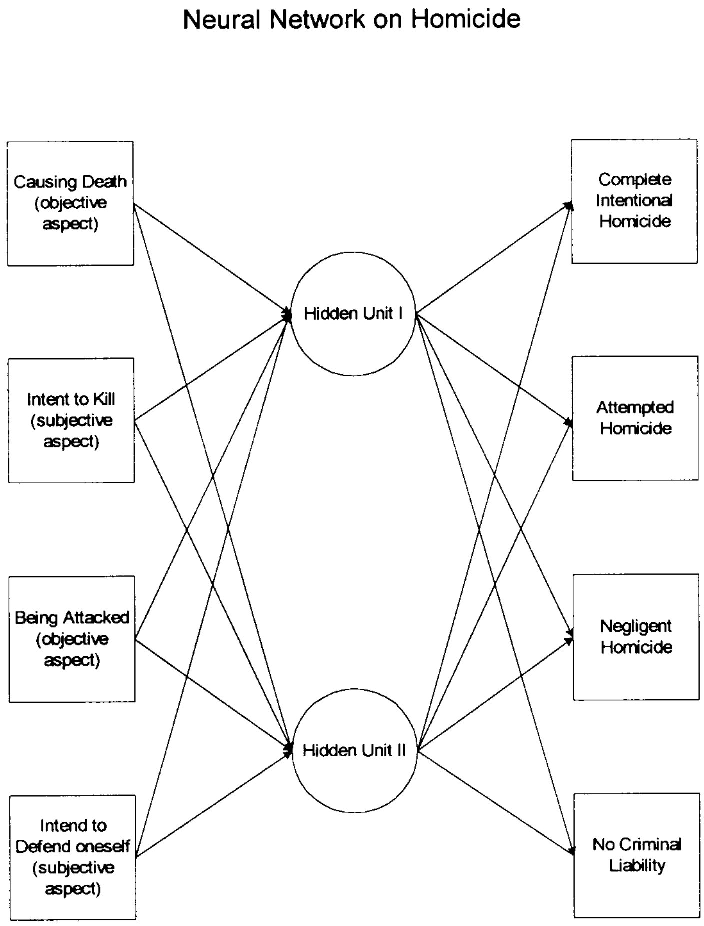

Figure 1. A neural network.

general kind.) The activity of input units is transmitted to hidden units which in turn transmit it to the output units. Some hidden units will not be involved: they will come into play with other input combinations. The behaviour of hidden units is not determined by a computer program, rather the network itself organises the flow of activity.

The connection between input and output is learned by the net, by practising learning patterns, i.e., precedents which provide models for later decisions. The learning process requires that a path is established from a certain input combination to the corresponding output. The process is iterative: incorrect results are fed back into the network and corrected gradually, until the learning patterns are "mastered" and have become habitual for the network. (We limit our discussion to feed-forward networks with back-propagation.)

What if the trained network is given a new input combination, rather than a familiar one? The network will associate the new combination with the input pattern which comes the closest to an already familiar learning pattern. Thus reasoning by analogy is achieved naturally without having been programmed into the network.

It is not difficult to link Hayek's observations to neural networks. So, his account of learning through examples matches the learning strategy usually adopted for neural networks; his conception of the articulation of rule-statements as the attempt to capture unconscious dispositions reflects the problem of giving a linguistic formulation to the knowledge embedded in such networks; his idea of the legal system as the expression of a self-organising (autonomous) order is linked to the phenomenon of self-organisation (and reciprocal adaptation) of neural networks; his idea of the lawyer's work as driven by consistency (or coherency) addresses the problem of improving the behaviour of a network by producing and refining an explicit representation of its knowledge. It seems to us that even the few connections just described may suffice to show the jurisprudential relevance of research on neural networks. This research on the one hand addresses problems which are crucial for legal theory, on the other hand may draw inspiration from some important trends of legal thinking.

Let us just mention three fundamental open issues in legal theory, which are reflected in technological problems in the area of neural networks.

The first concerns the link between the "context of discovery" and the "context of justification", i.e., between the way in which a legal solution is found, and the way in which it is supported by giving appropriate reasons.

In neural networks, a solution is not found by articulating reasons, but rather by drawing an analogy from previous successful problem-solving (the view of neural reasoning as analogising is contested in Dan Hunter's article). This originates the problem of whether and how this solution should be justified. Should no justification at all be provided for the results of a neural network (as Hunter argues)? Should it be an opportunistic rationalisation, given ex post? (This idea is developed in the contribution of Andrew Stranieri, John Zeleznikow, Mark Gawler and Bryn Lewis who accept the realists' idea of justification as rationalisation.) Should it express the serious attempt to induce a set of coherent principles from the network's problem-solving patterns, principles which may affect the problem-solving behaviour and even modify its "unconscious" patterns? This issue is linked to even broader concerns: how does legal thinking combine symbolic and sub-symbolic processes, how can the controlled use of refined conceptual structures interact with the unconscious firing of parallel connections?

The second theoretical issue concerns learning the law. As we have seen, neural networks adopt an example-based (rather than a theory-based) model of learning. Such an approach to legal learning is not new to lawyers: cases have always had a paramount importance in the teaching of common law and civil lawyers too are increasingly appreciating the importance of cases. (As some among them have always done: Friedrich Carl von Savigny already observed that the law can only be learned by mastering the casuistry of the Roman sources.) However, example-based learning can be developed in very different ways, in particular, according to different criteria for choosing the examples.

Should the selected learning patterns be pure prototypes? Relatively few of such learning patterns would be required, each dealing with a single problem. These examples could be found in textbooks. Alternatively, should all available court cases be used for training the network? Does this engender the risk that examples are flawed by secondary problems, which might "contaminate" the main problem? (This aspect is addressed in Hunter's contribution.)

The third theoretical issue concerns "self-organisation", which is a key term not only for neural networks, but also for many other disciplines, from biology, to system theory, to economics, to computer science. Modern sociology of law has frequently appealed to this concept and has build very abstract and wide-scoped pictures of the development of legal systems on the basis of ideas of evolution and self-organisation (Luhmann's legal sociology exemplifies this trend.) However, legal studies have so far met little success in providing precise (and possibly controllable) accounts of specific aspects of the law on the basis of theories of self-organisation. Possibly, the vantage point of neural networks (as Danièle Bourcier and Gérard Clergue observe) will allow this development to take place.

3. Fuzzy Logic

Let us now move to the second feature of legal reasoning which is addressed in this special issue, that is dealing with fuzziness, or vagueness. The jurisprudential problem which fuzzy logic undertakes to solve is old, even ancient (cf. Ron A. Shapira's contribution to this issue): A judge must decide on legally relevant situations, which can often be described only in indeterminate terms. The decisions must be determinate, however, and often even expressed as a numerical quantity. Let us shortly consider what indeterminacy means.

Indeterminacy is not uncertainty (as Jacky Legrand observes). "Mater semper certa est, pater semper incertus", the Romans said, meaning that one could never be sure that a man was the real father of a child, even if he was the mother's husband. What a "father" was, however, was regarded as perfectly determinate.

Indeterminacy also does not mean ambiguity. The term "mother" has become ambiguous, thanks to modern reproduction technology: mother might be the woman whose egg was fertilised in vitro, or the woman who received the fertilised egg in her uterus and bore the child. The term of (biological) mother has thus split and has become ambiguous, but each of the two meanings can be clearly defined and is therefore determinate. On the contrary, a genuine example of indeterminacy is provided by the legal term "negligence", and the attribute "gross" which might qualify negligence is in turn indeterminate.

The most famous attempt to analyse indeterminacy was undertaken by the Hart (1961). He distinguishes a term's "core of certainty" and its "penumbra of doubt". Hart's example of "vehicles" which are not permitted in a park has become the paradigmatic example of an indeterminate legal term. While a regular car would lie at the core of certainty, a skateboard (which had not been invented when Hart published his book) would fall into the penumbra of doubt, and a toy car would clearly lie outside the term vehicle. Hart's metaphor is not bad, but not great either: there are definitely phenomena to which a term applies clearly, some that give reason for doubt, and others which are certainly outside the term's scope. However, at what point may one begin to reasonably doubt the term's applicability and at what point is such application out of the question? This remains in the realm of metaphor and cannot be fixed and criticised. Most importantly: "penumbra" implies a uniform distribution, when in reality there is a curve of decreasing applicability.

A better metaphor was used much earlier by the German legal scholar Philip Heck (1914, p. 107): "A nucleus of certain meaning is surrounded by a gradually fading halo of meaning". Gradual dimming has replaced the uniform penumbra. Heck's image allows us to ask whether a term is almost applicable. Even this metaphor, however, leaves the gradual dimming non-quantified.

A different aspect of indeterminacy was emphasized by the legal scholar Adolf Merkel in the 19th century: "Many legal terms have a blurry quality. Their areas of application are not delineated by insurmountable fences, but rather they spill over into neighbouring areas". (Cited after Grünhut 1926, p. 16.) This metaphor anticipates the conceptualisation of modern fuzzy logic. Thus the skateboard might be considered a "vehicle" to a certain degree, but also to a certain degree an "aid to pedestrians". However, to what degree still cannot be determined. From a legal standpoint Merkel's metaphor is so appropriate because the conflict of two legal terms will usually represent the conflict of two underlying interests. Terminological limits are lines of demarcation in conflicts of interests, and a term's indeterminacy makes it the subject of legal argument. (For this reason one might expect Merkel's metaphor to have been coined by Heck, the founder of "Interessenjurisprudenz".) The decision, for instance, whether a skateboarder is using a "vehicle" in a park depends, on the one hand, on the interest in a leisure activity, and on the other hand on the interest in peaceful and tranquil enjoyment of the park without fear of being disturbed or injured.

On comparing Merkel's, Heck's, and Hart's metaphors the impression is that they become less acute the more recently they were coined. Perhaps this is no coincidence but a result of the growing awareness of logic and set theory. Hart's metaphor is compatible with two-valued logic, in which an element either belongs to a set or does not. The only concession necessary is that there are cases in which the question of set membership is (at least temporarily) "open". Heck's metaphor is more problematic: it requires that the degree of an element's membership may be greater or lesser. Merkel's metaphor finally, if taken seriously, requires us to accept that the law of contradiction does not apply, or at least not completely.

This is exactly the concession fuzzy logic requires us to make. How absurd this must have seemed a few decades ago is reflected in the words of abusive criticism

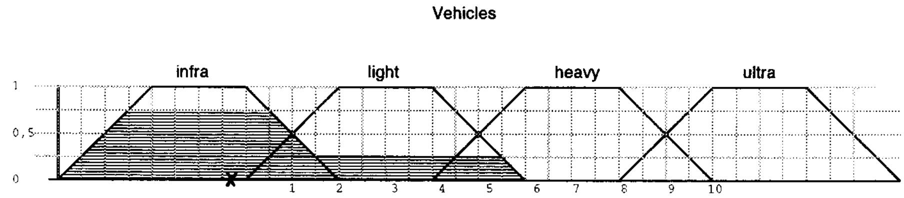

Figure 2. A representation of fuzzy concepts.

of leading scientists, which were directed against Lofti Zadeh, the pioneer of fuzzy logic, and which Bart Kosko (1993), one of Zadeh's followers, gleefully cites today.

Zadeh invented a mode of representation which bridges the gap between the indeterminacy of natural languages and the numerical determinacy of mathematics and which combines compelling imagery and mathematical precision. His method allows us to link the determinacy of decisions to fact patterns described in indeterminate terms. (For a pictorial representation of Zadeh's method see Figure 2 below.) Zadeh represents terms using trapezoids (or triangles) which in their lower, broader area overlap with neighbouring terms, but are separate from their neighbours in their upper, narrower area. To represent concrete cases, these trapezoids are filled to the level which corresponds to the degree to which the term applies to the facts. Filling the trapezoid completely represents cases which clearly fall within the term's meaning. Where the trapezoid's sides begin to sink down, drawing further apart, the area begins in which doubt is possible, and the doubt grows stronger as you go on. As soon as a side intersects with the side of the neighbouring trapezoid, the neighbouring term will also be applicable, again to a certain degree.

In this manner an estimate can be expressed that the skateboard is an "aid to pedestrians" up to a degree of 0.75 and a "vehicle" up to a degree of 0.25. We should not worry if this assessment is not fully satisfactory: the estimate can be corrected, we will return to this point. However, is the skateboard a vehicle or not? The either-or of the principle of the excluded middle is as dispensable to fuzzy logic as the not-both-at-once of the law of contradiction. The law however, requires a clear decision. Fuzzy logic may help to prepare the decision. As Zadeh took the metaphor of overlapping terms seriously, so he also took another metaphor seriously: it is the focal point, the centre of gravity which matters. One has to determine the centre of gravity of the area that is filled by the overlapping trapezoids. This is a matter of integral calculus. The result is easy to visualise, however, by imagining that the filled area of the overlapping terms is drawn on cardboard, cut out, and balanced on the finger tip until the equilibrium is found: that point marks the centre of gravity. For our skateboard example, the centre of gravity plausibly falls in the trapezoid (the term) of "aid to pedestrians" (cf. Figure 2.)

All this can be specified numerically. The distance between a filled trapezoid's lower and upper borders, which expresses the degree up to which the represented term is satisfied, can be assigned values ranging from 0 to 1 on the (vertical) y-axis. Values on the (horizontal) x-axis, on which the trapezoids' bases rest, represent the increasing strength of a decision: in the legal realm this might be the amount of damages in Euros or length of prison sentence in months. By marking off the centre of gravity on the x-axis, the measure of the resulting decision can be determined with numerical exactness. The system's numerical quantification allows primarily this: the fuzzy logic model can be checked and modified if the decision found in a first run seems unsatisfactory. This applies to the trapezoid's size and the placement of the intersections. In these modifications the terms' proportions must naturally remain plausible. Often a painstaking process of reducing and increasing, of shifting back and forth is required before a model is satisfactorily adjusted. This requires much patience and it is perhaps for this reason that fuzzy logic is especially successful in Asia, in Japan and now also in China, in technical applications (fuzzy control), but increasingly also in information technology (see the contribution of MinQiang Xu, Kaoru Hirota, and Hajime Yoshino.)

A note on terminology: The terms of natural language which are represented by trapezoids are regarded as values of a "linguistic variable". The variable in our example could be termed "means of propulsion" with the values "vehicle" and "aid to pedestrians". At the trapezoids' base, a "numeric variable" can be marked off, such as "prison sentence in months" or "damages in Euros". Fuzzy logic aims at facilitating the transition from natural language terms to numeric expressions. The first step in the process, capturing the facts in overlapping terms and trapezoids, is called "fuzzification". The second step is "defuzzification": determining the terms' (trapezoids') centre of gravity to reach a numerically quantifiable decision.

Furthermore, just as a single statement need not be completely true or false, but may approach truth without reaching it, a logical deduction may approach strict validity more or less. Lofti A. Zadeh was among the pioneers also in this area of research, giving it the name of "approximate reasoning". It is obvious that the phenomenon of approximate but nevertheless rational deduction, the rules which govern it and the criteria to judge its validity, are of importance to legal argumentation (this issue is addressed in the contributions of Jacky Legrand and Lothar Philipps).

4. Conclusion

There are obvious similarities between neural networks and fuzzy logic. An input might stimulate neighbouring hidden cells, and the network will decide according to the "peak" stimulus; in fuzzy logic models this corresponds to the "centre of gravity". In addition, both systems use an iterative process: the results will be used to correct the preconditions. It is thus no wonder that fuzzy logic and neural networks are well suited to each other. A neural network may be used to fine-tune a fuzzy logic model or both systems may be linked in a hybrid solution. Zadeh once told one of the authors of the present introduction (Lothar Philipps): "Fuzzy logic and neural networks are like two sides of a coin, or – in other words – they relate to each other as psychology and physiology do".

If fuzzy logic models and corresponding neural networks might be interchangeable in principle, there may be reasons for preferring the ones to the others, in specific contexts. Let us consider some reasons which may apply in legal domains. When the central problem lies in the conditions of a norm's application a fuzzy logic model will be more suitable. This is the case, for example, when terms, which have been established by legislation or legal scholarship, must be distinguished from one another (i.e., when the problem lies in a semantic indeterminacy of the antecedent of a norm). When, on the other hand, the legal consequence of a norm must be finely and harmoniously differentiated between several similar cases, a neural network might be considered first. Examples would be the determination of immaterial damage compensation, estate divisions and alimony payments in family law (as in the contribution by Stranieri et all.) However, these are tendencies, not necessities. On the one hand the terms in a norm's condition need to be fine-tuned to obtain a just and differentiated legal result. On the other hand the grading of legal consequences (e.g., criminal sentences) depends on conditions which are represented in overlapping terms (this aspect is considered in István Burgulya's article)

On a more personal note, we must add two statements to this introduction. The first is a request for indulgence which is directed mostly to the authors. Illness and other misfortunes have unduly delayed the publication of this special issue. The other is an expression of heartfelt gratitude to Cees Groendijk who bore the greatest share in the labour of organisation.

Appendix: Two Examples

To clarify and explain the notions presented above (and to make them more accessible to the reader approaching the subject matter of this volume for the first time) we provide two full-fledged examples complemented by pictorial representations.

The first (Figure 1 above) is a neural network concerning a potential criminal offence, homicide. The input units on the left show conditions that may or may not be fulfilled, while the output units on the right display legal categories the act might fall under. In the middle there are two hidden units (Philipps 1991).

There are four possible categories of homicide:

(1) intentional and completed homicide,
(2) attempted homicide,
(3) negligent homicide, and
(4) no criminal liability (at least as far as homicide is concerned).

The conditions under which an act can be assigned to one of those categories depend on four components, which can be present or not:

(1) the objective element of causing death,
(2) the subjective element of intent to kill,

(3) a second objective element: a situation of possible self-defence (an attack by the victim to-be), and
(4) a second subjective element: the intent to defend oneself.

Obviously there are 16 combinations for the presence or absence of the four components. For each of the four categories, the paradigmatic combination of components is used as a learning pattern. (Admittedly, there are some simplifications in this system.)

(1) For the intentional completed killing, both the objective and subjective elements of homicide are regarded as present, but not those of self-defence.
(2) For attempted homicide, only the subjective element of intent to kill is present.
(3) For negligence, only the objective element of death caused is present.
(4) For no liability, none of the four elements are present.

To these four learning patterns, two more were added:

(5) another typical case of no liability ("justification"), in which someone kills another intentionally in a situation of self-defence, and
(6) a variant of the negligent crime, in which the perpetrator erroneously believes himself to be attacked. (This result is accepted by the mainstream in many criminal law systems; those lawyers who follow a minority opinion and would punish the perpetrator for intentional homicide in this case must leave out this learning pattern.)

Given six learning patterns for 16 combinations, ten constellations remain open. The network is able to determine the ten missing solutions in a way that is legally consistent. It should be noted that the conclusions are achieved by real analogy and not by "analogy" simulated by means of formal logic; from a purely logical standpoint, different outcomes would be equally possible.

Not all the network's proposals are unanimously accepted in jurisprudence. The following case is especially controversial: one kills another without realising that he is already being attacked by the other. The perpetrator acts (objectively) in a situation that allows for self-defence, but (subjectively) not in order to defend himself. The neural network presented here suggests that this case should be viewed as an attempted homicide, a solution which is mainstream in the jurisprudence of many European, Latin American, and Asian countries. Nevertheless, some authors suggest a completed homicide, while a small minority see no criminal liability at all. It is fascinating that these results can also be achieved by the neural network. Its structure will have to be modified: the hidden units removed. The network will then suggest completed homicide (with greater weight) or no liability at all (with lesser weight).

All controversial solutions proposed by a neural network will remain within a particular doctrinal framework. This lends support to Hayek's hypothesis of structural assumptions. Legal scholars engaged in doctrinal debates follow certain structural assumptions of which they may not be completely aware. However, light can be shed on those assumptions by experimenting with neural networks.

Let us finally illustrate our second picture (Figure 2 above), which is dedicated to fuzzy logic. It develops the classical example of H.L.A. Hart (vehicles in the park) to show the link between fuzzy representation and the "binary" logic of legal decision making.

As the picture shows, fuzzy logic enables a transition from vague terms to precise decisions. In our example, anyone driving "vehicles" in a park – "light" or "heavy" ones – is liable to be fined. (The "money units" on the y-axis, multiplied by 100, could mean Euros, Dollars, etc.) There is an overlapping area between any bordering terms: heavy motorcycles could be considered to be "heavy" light vehicles, or small motorcars regarded as "light" heavy vehicles. Note that the range of decisions can be utilised fully only if the extremes of the terms' scale partially overlap with terms beyond – infra or ultra. "Infra" terms could be skates, skateboards etc; "ultra" terms could be tanks or tractors. By counting the number of squares that cover the filled areas of the trapezoids (terms) one can easily verify our conjecture that as for skateboards (our example), the centre of gravity lies outside of the area of "real" vehicles – and so skateboards are not to be fined. (It is presupposed here that a skateboard is a vehicle only up to one-quarter, and an "aid to pedestrians" up to three-quarters.) By the way, this grid technique, and other methods of simplification as well, is often preferred to the integral calculus in fuzzy control applications, especially when it is more important to optimise speed instead of precision.

References

Grünhut, M. 1926. *Begriffsbildung und Rechtsanwendung im Strafrecht.* Tübingen: Mohr.

Hart, H.L.A. 1961. *The Concept of Law.* Oxford: Oxford University Press.

Hayek, F.A. 1973. *Law Legislation and Liberty. Volume I. Rules and Order.* London: Routledge and Kegan Paul.

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.

Hayek, F.A. 1979. *Law Legislation and Liberty. Volume III. The Political Order of a Free People.* London: Routledge and Kegan Paul.

Heck , Ph. 1914. *Gesetzesauslegung und Interessenjurisprudenz*. Tübingen: Mohr.

Kosko, B. 1993. *Fuzzy Thinking, the New Science of Fuzzy Logic.* New York: Hyperion.

Perelman, Ch. and Olbrechts-Tyteca, L. 1969. *The New Rhetoric. A Treatise on Argumentation*. Notre Dame (Indiana): University of Notre Dame Press.

Prakken, H. 1997. *Logical Tools for Modelling Legal Argument. A Study of Defeasible Reasoning in Law.* Dordrecht: Kluwer.

Philipps, L. 1991. Tû-Tû 2. Von Rechtsbegriffen und neuronalen Netzen. In L. Philipps and R. Wittmann (eds), *Rechtsentstehung und Rechtskultur. Festschrift für Heinrich Scholler*, Heidelberg: Decker & Müller, pp. 179–189.