well-known restaurant script details the events and agents involved in a typical restaurant visit, thus providing a 'holistic' structure that captures a piece of world knowledge. Generalizations also capture world knowledge but are typically more 'atomistic', for example, rules such as 'if you don't pay in a restaurant the owner will usually call the police'. As Bex, Bench-Capon and Verheij have argued recently (Bex et al. 2011), scripts and stories are essentially factual, non-legal cases in the vein of Aleven (1997), and the same argument moves available in CATO can be applied to cite and distinguish these factual stories. Furthermore, Bex and Prakken (2004) showed that Loui and Norman's compression rationales are not just relevant for legal rules but also for commonsense rules (i.e. generalizations). For example, the rule captured by 'if an eyewitness testifies P then usually P is the case' can be decompressed into three separate rules that capture the veracity of a witness (if a witness says 'I remember I saw P then P), her objectivity (if a witness remembers he saw P then P) and her observational sensitivity (if a witness saw P then P). Thus, if we extend the work by Loui and colleagues to apply to not just legal cases and rules but also to non-legal cases and rules, we have a flexible and powerful framework that encapsulates different ways of reasoning with commonsense knowledge as defeasible argumentation.

5.6 Trevor Bench-Capon (1993). Neural networks and open texture. *Commentary by Bart Verheij*

Is it possible to learn the rationale underlying decisions in an open textured domain, such as the law, given only a set of decided cases? It is this important question that Bench-Capon (1993) investigates in Bench-Capon (1993). He investigates the question using neural networks. At the time, neural networks were a popular research topic, as they had helped solve problems, e.g., recognition of hand-written characters, that seemed unsolvable using logic-based artificial intelligence. The paper is a beautifully written, self-contained essay, and contains lessons about the automatic learning of rules from cases that are still valuable.

The learning experiments are performed on cases concerning a welfare benefit to be paid to senior citizens visiting their hospitalized spouse. Past decisions are artificially generated (by a LISP program), constrained by six conditions. The setting is fictional, in the sense that it does not reflect an actual welfare policy, but the example has been carefully designed in order to be able to investigate different kinds of legal conditions. For instance, there is the Boolean condition that the two persons involved should be married. There is a threshold condition, namely that the couple's capital resources should not exceed some fixed amount. There are also dependencies between variables, such as between age and sex: the person should be of pensionable age, 60 for women, 65 for men.

Bench-Capon investigates two questions: Can the neural network be trained to decide cases on the basis of a given set of decided precedent cases? And can the decisions proposed be justified in terms of the conditional constraints used to generate the cases?[8]

[8] Bench-Capon discusses a third question, namely whether we can derive rules from the networks. Because of the way in which he addresses this question, it is not further discussed here.

Bench-Capon answers the *first question* about learning correct decisions with a resounding 'Yes'. Quoting the paper:

> Neural networks are capable of producing a high degree of success in classifying cases in domains where the factors involved in the classification are unknown. (p. 296)

He reports success rates of around 99 % for networks with one, two or three hidden layers, with networks trained on 2,400 generated cases and tested on 2,000 cases. He shows that in his setup irrelevant factors do not strongly reduce performance.

The *second question*, Can the network's decisions be justified?, is answered with interesting nuances. Bench-Capon's original training set shows that for the one and two hidden layer network designs *sex and age do not matter* for the decision: when the other constraints are satisfied, the network proposes the decision to pay the benefit, for almost all ages, even for ages close to 1. In the three hidden layer design, there is a relevant correlation between sex, age and the decision proposed, but the age difference associated with the dependency is 15, not 5. Also everyone above 40 is paid the benefit, instead of everyone above 65.

In order to explain this finding, Bench-Capon analyzes his material and finds that already four out of six factors can explain about 99 % of the cases. Bench-Capon continues his experiments on the basis of sets of generated cases with a more careful distribution over the satisfaction or non-satisfaction of the different conditions. He finds that he can hereby steer the network's performance to more closely reflect the sex-age constraint as it was actually used to generate the cases.[9]

He draws an insightful and important conclusion. That the two features were missed 'can only be detected if we have some prior knowledge of the domain which allows us to say this: otherwise we have no way of telling that the four conditions that were discovered were not in fact the whole story' (p. 294–295). This is almost like saying that the rationale shown by the neural networks can only be estimated by knowing about the rationale beforehand. Put yet another way: rules mimicking past decisions need not determine the rules used to make the decisions.

An interesting, and still relevant, discussion of Bench-Capon's paper can be found in the 1999 special issue of the journal *Artificial Intelligence and Law* (Philipps and Sartor 1999), devoted to neural networks (and fuzzy reasoning). In a position paper on the contemporary state of the art of applying neural networks to the law, Hunter (1999) claims that neural network applications to the law were in general flawed, in particular, because they were often based on inappropriate data. In contrast, Bench-Capon's paper is praised as being 'methodologically most appropriate', but, says Hunter, it is just 'about the use of neural networks to simulate necessary conditions, since the training set and the verification set were simply drawn from rules'. From this, Hunter concludes that Bench-Capon is not about *legal* neural networks at all, for they only model an already given rule-encoded doctrine—or, when learning fails, not even that.

[9] A question remains however: the network has 45 and 50 as significant ages, instead of 60 and 65. Bench-Capon gives no explanation for this oddity: is it a systematic consequence of the learning rule used, perhaps a small bug in the set up?

I think that Hunter misses the importance of that part of Bench-Capon's results, and of their relevance for law. Still, Hunter is right that some typically legal phenomena are not covered. Indeed, Bench-Capon's networks show no development in the light of new circumstances, cannot handle new insights, cannot incorporate landmark precedents, do not incorporate values that steer decisions. Interestingly, each of these themes is particularly addressed by one of Bench-Capon's other core research topics: case-based argumentation (with Ashley's 1990 a milestone).

I wonder what Bench-Capon's reaction at the time was when reading Hunter's paper. I believe that in his heart he agreed; his neural networks were not sufficiently 'legal'. For as we all know, in the years that followed, Bench-Capon devoted much of his time to what was missing, and developed dynamic case-based reasoning techniques, with the incorporation of values a key innovation (see, e.g., his influential Bench-Capon 2003).

At the same time, Bench-Capon (and the field) remains interested in the learning of rules underlying legal decisions, a recent contribution being the PADUA system (Wardeh et al. 2009). In this work, the same legal conditions are used as in the 1993 paper, this time using training sets that can contain errors.

Today Bench-Capon's 1993 conclusion seems to stand strong: The patterns in a given set of decisions do not determine the rules that led to the decisions. Perhaps the time has come to reconsider the dynamical relations between the logic of decision-making and the probabilities associated with data description?

5.7 Giovanni Sartor (1993). A simple computational model for nonmonotonic and adversarial legal reasoning. *Commentary by Guido Governatori*

Sartor (1993) addresses the issue of the representation of non-monotonicity in legal reasoning. He does so by providing a computationally oriented model in PROLOG. The nineteen-nineties can be considered the golden age for non-monotonic reasoning. A large amount of research in the general field of artificial intelligence was dedicated to this then emerging and promising area. Artificial Intelligence and Law was not immune from this trend and scholars like Sergot, Prakken, Gordon, Bench-Capon (and many others) were interested in exploiting these developments. A contribution of Sartor (1993) was an analysis of areas of law and legal reasoning requiring non-monotonic reasoning. A key observation was that resolution of conflicts in non-monotonic reasoning is based on the concept of preference ordering over elements in conflict, and this is no different from the resolution of what D. Ross called *prima facie duties* when they conflict, i.e., conflicting norms. However, normative systems are dynamic, there can be multiple concurrent legal sources, and legal languages often leave space for semantic indeterminacy (i.e., multiple interpretations are possible). These considerations led Giovanni Sartor to state that formalisations of legal reasoning need inference procedures taking into account an ordering relation, and that ordering relation should be obtained using many criteria that have to be "harmonised". The paper proposes a solution for this issue.

The technical solution developed in Sartor (1993) is that norms are represented by rules with the form $n : p_0 \leftarrow p_1 \wedge \ldots \wedge p_n$, where n is the name of the rule, where