

The Law as a Learning System¹

Edwina L. Rissland
Robert T. Collins²

Department of Computer and Information Science
University of Massachusetts

ABSTRACT

In this paper, we present three examples of the development of legal concepts and rules and analyze them in terms of AI learning methods, in particular, the candidate elimination algorithm. The three areas of legal doctrine are : (1) the doctrine of the liability of manufacturers to third party buyers; (2) the Fourth Amendment doctrine relating to stop-and-frisk searches; and (3) the doctrine of attractive nuisance. For each of the legal areas, we give a synopsis of the legal developments and show how it can be viewed from an AI point of view.

BACKGROUND

Work on machine learning [Carbonell et al., 1983] has typically concentrated on well-established, well-structured problem areas, like methods of integral calculus. In such domains, concepts and their interrelationships are already known ahead of time to the learning program. Thus the typical concept or rule refinement learning program already knows a concept description language over which it is going to form its generalizations and specializations.

There is nothing wrong with knowing the conceptual landscape before learning new material. However, there are other more fluid situations resting upon less finalized concept spaces in which a lot of interesting learning takes place. This is even true in traditional "neat" domains like mathematics [see Lakatos, 1976]. Furthermore, even in the case of learning in well-established and well-structured domains, the concept spaces are almost never truly static, almost always have some "bias", and are amenable to re-definition and re-structuring [Mitchell, 1983; Utgoff, 1983].

The law provides an example of a domain which is both fluid and "messy". Legal concepts and rules can change, and the concepts are "open-textured," that is, have no clearcut boundaries or hard-and-fast definitions. Sometimes both the doctrine (rules) and the underlying conceptual spaces are in flux at the same time.

In the sense that the legal system as a whole responds to new training instances (i.e., cases) and changes the performance of its adjudication task, we will think of the law as a learning system. The most obvious repository of its knowledge is its case base, which is published and accessible to all. In this paper we concentrate on how the law learns new rules and concepts as evidenced through discussion in published opinions.

¹This work was supported in part by the Advanced Research Projects Agency of the Department of Defense, monitored by the Office of Naval Research under contract no. N00014-84-K-0017.

²Copyright ©1986. Edwina L. Rissland & Robert T. Collins. All rights reserved

The learning algorithm which we use to analyze the evolution of legal knowledge is Mitchell's candidate elimination algorithm (CEA) [Mitchell, 1983; *The Handbook of Artificial Intelligence, Vol 3*]. In essence the algorithm works to narrow the definition of a concept or range of applicability of a rule by squeezing in from very general and very specific instances. Positive instances cause CEA to generalize and negative ones cause CEA to specialize.

The CEA algorithm works by maintaining a "version space" of all descriptions which are consistent with the examples seen so far. The version space is represented by maintaining a maximally specific set S and a maximally general set G of concept descriptions. All other descriptions in the version space are implicitly contained within these two bounds. To refine the version space, CEA examines positive and negative "training instances". Given a positive instance of the concept, the set S is made more general to include descriptions matching it; set G is also made more specific to exclude descriptions covering negative instances. These positive and negative refinements push S and G closer together until they converge on the correct description.

The idea of describing the behavior of the legal system as a learning system using a CEA-like algorithm can be seen in an analysis by the legal scholar Max Radin :

"The single instance is capable of generalization, and the generalization will not stop at any particular place, unless by negative decision, by a statement that a given situation is outside the genus, a subsequent court has deliberately attempted to stop it. Then the process begins all over again, because the excluded situation is itself capable of successive generalizations and we must know whether a large or small genus is to be excluded from (A). A Buick car is not in (A). Is an automobile truck? An electric machine? A hypothetical new type of car driven by more explosive mixtures than gasoline, and so on?"

– Radin, 1933, at p. 209

One difference between application of CEA to the legal domain and the more usual ones is that in the law, the hierarchy of concepts and the range of a rule usually evolve at the same time.³ This is so because of the non-static, open-textured nature of the law. In particular, one can change the effect of a rule by changing either (1) the rule itself such as through the addition of conjuncts to the pre-conditions or listing of exceptions to the rule, or (2) changing the meaning of the ingredient concepts. Thus in our analysis we interleave discussion of the evolution of a legal rule and its group of ingredient concepts. From an implementation standpoint, that means we have two CEA processes running cooperatively. We illustrate the development of rules and concepts in three example doctrines—products liability, search and seizure, and attractive nuisance. There is nothing special about these three; they are indicative of the growth of legal knowledge in general.

THE INHERENTLY DANGEROUS EXAMPLE *LEARNING A LEGAL CONCEPT*

Our first example is a famous one that has been written of several times by legal scholars [Levi, 1949; Radin, 1933]. It concerns the development of the idea that a manufacturer of goods is liable (in

³Note: Even though we concentrate on "rules," Anglo-American law is based on cases and not rules. This is so even in statute law, for here the ingredient rules and concepts in the statute must be interpreted in light of specific cases (see [Levi, 1949] for an example).

case	date	item	finding
<i>Thomas v. Winchester</i>	1852	belladonna	liable
<i>Loop v. Litchfield</i>	1870	balance wheel	not liable
<i>Losee v. Clute</i>	1873	steam boiler	not liable
<i>Devlin v. Smith</i>	1882	painter's scaffold	liable
<i>Torgesun v. Schultz</i>	1908	bottle of aerated water	liable
<i>Statler v. Ray Mfg. Co.</i>	1909	coffee urn	liable

Table 1: Summary of "Inherently Dangerous" Cases

tort) to third parties injured by a defective product. For instance, if you buy a defective Buick and a front wheel falls off and you are hurt, the manufacturer of the car, Buick Motor Company, must compensate you for your injury even though you bought the car from an intermediate vendor like a car dealership (hence, you are called a "third" party), and not directly from Buick Motor Company. A closely related doctrine, which developed later and in a parallel way is the (contracts) doctrine of "implied warranty" by which is meant that there is an implicit warranty that goes along with a product to the effect that the product is warranted to safely do what it is supposed to do as long as it not abused. These doctrines are at the basis of our current protection of consumers.

This was not always the case. Originally, there was no relationship of liability between manufacturers and third parties : the law required "privity" of contract, that is, an immediate seller-purchaser relationship between manufacturer and injured. This could be summed up by the rule of thumb : *no privity, no liability*.

Our learning episode begins in 1852 with the case *Thomas and Wife v. Winchester* decided in the New York Court of Appeals. In this case, Mr. Thomas had his local druggist fill a prescription for dandelion extract for his ailing wife. The druggist filled the prescription from a jar which he had bought from another druggist who had bought it from the manufacturer, Winchester. Unfortunately, what was actually in the jar was belladonna, a substance looking somewhat similar but causing very different effects. As a result, Mrs. Thomas got quite sick and suffered "derangements of the mind"; she eventually recovered. Mr. and Mrs. Thomas sued Winchester, the manufacturer and packager of the mislabelled poison.

The Thomases won. The court justified making an exception to the "privity" rule because it believed that misbranded poisons like belladonna were so "imminently" or "inherently dangerous" that they deserved an exception. The court likened belladonna to a loaded gun in the hands of a child, an item so "inherently" dangerous that it was sure to cause injury to unsuspecting and innocent victims. Thus this landmark case established an exception to the privity rule. Subsequent cases would establish its scope.

For the period from 1852 until 1916, the New York Court of Appeals handed down decisions (see Table 1) concerning further exceptions to the privity rule. The exception started off with belladonna and ended up including painter's scaffolds, bottles of aerated water and coffee urns. It was not extended to cover items such as carriages or steam boilers which become dangerous only if defective.

In effect, the Court was defining the *inherently dangerous* category in a CEA-like manner by including or excluding examples on a case-by-case basis. That is, if an item allowed recovery for third person injury under the rubric of the *inherently dangerous* exception, it was in the class—i.e., was a *positive instance*— and if it didn't, it wasn't—a *negative instance*. Positive instances served both to flesh out the concept of *inherently dangerous* and added pressure to generalize it; negative instances narrowed it.

In this middle period of about fifty years, the Court was also steadfastly refusing to generalize the *inherently dangerous* exception to a more general class like *dangerous if defective*. However, in this period it was simultaneously pushing the boundaries of the *inherently dangerous* class (and thus the exception) to include some items like coffee urns and hairwash which seem to be more of the class of item that are *dangerous only if defective*. Nevertheless, the Court spoke as if it wasn't and summed up its position on the privity rule as follows :

“One who manufactures articles inherently dangerous, e.g., poisons, dynamite, gun-powder, torpedoes, bottles of aerated water under pressure, is liable in tort to third parties... On the other hand, one who manufactures articles dangerous only if defectively made, or installed, e.g., tables, chairs, pictures or mirrors hung on the walls, carriages, automobiles, and so on is not liable to third parties for injuries caused by them, except in cases of willful injury or fraud.”

– *Cadillac v. Johnson*, 221 F. 801 (1915), at p. 803

This case represents the last gasp of a court trying to hold onto its “old” interpretation of the concepts and the scope of the exception. Note it explicitly rules out carriages and automobiles. In the very next year these very things were let in.

In 1916, the New York Court of Appeals decided the landmark case of *MacPherson v. Buick* 217 N.Y. 382. In this case, MacPherson, a third party, was allowed to recover for injury caused by a defective Buick, an item *dangerous only if defective*. This is exactly the case that the Court had persistently (e.g., in *Cadillac* case above) said it wasn't going to let the exception cover. The case of a car or carriage had now flip-flopped from being a negative to being a positive exemplar for the privity exception. As a positive exemplar, it greatly expanded the exception. *MacPherson* also ended the unreasonable stretching of the *inherently dangerous* class and the need to talk fictions about being *inherently dangerous* to win recovery for an item *dangerous only if defective*.

As Judge (later to become Supreme Court Justice) Cardozo stated in this landmark opinion :

“We hold then, that the principle of *Thomas v. Winchester* is not limited to poisons, explosives, and things of like nature, to things which in their normal operation are implements of destruction. If the nature of a thing is such that it is reasonably certain to place life and limb in peril when negligently made, it is then a thing of danger. Its nature gives warning of the consequences to be expected. If to the element of danger there is added knowledge that the thing will be used by persons other than the purchaser, and used without new tests, then, irrespective of contract, the manufacturer of this thing of danger is under duty to make it carefully. That is as far as we are required to go in this case. There must be knowledge of a danger, not merely possible, but probable. It is *possible* to use almost anything in a way that will make it dangerous

if defective ... where danger is to be foreseen, a liability will follow ... We have put aside the notion that the duty to safeguard life and limb, when the consequences of negligence may be foreseen, grows out of contract and nothing else. We have put the source of the obligation where it ought to be. We have put its source in law.”

This opinion places automobiles within the class of *dangerous only if defective* and grants recovery for injury by them. Thus the exception to the privity rule has been generalized. *MacPherson* also creates a somewhat finer concept hierarchy : it speaks of the class *thing of danger*, which includes both things *inherently dangerous* and things *foreseeably dangerous if defective*. See Figure 1.

Thus, after the episode beginning with *Thomas v. Winchester* decided in 1852 and ending with *MacPherson v. Buick* decided in 1916, the New York Court of Appeals has learned new concepts (e.g., *inherently dangerous*, *dangerous only if defective*), placed them into a hierarchy of concepts, articulated an exception to a well-established rule, and re-classified a training instance for the rule from negative to positive (i.e., the car in *Cadillac* and that in *MacPherson*). Said another way, the law has learned the new *liability to third party* rule and its underlying concepts.

THE STOP AND FRISK RULE

Another example of the law elaborating an exception to a rule, and the concepts underlying it, is the development of the stop-and-frisk doctrine. This area of the law involves Constitutional issues surrounding the Fourth Amendment which states :

The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no Warrants shall issue, but upon probable cause, supported by Oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized.

The Amendment states that in order to search and seize a person or property (NB an arrest is a seizure of person), the police must have “probable cause” and a warrant.

Prior to the first stop-and-frisk case which came before the Supreme Court in 1968, there already were well-established exceptions to the necessity of a warrant for search and seizure but none to the necessity of probable cause. Examples of situations where police can perform a warrantless search are : (1) if a person consents, (2) if the search is incident to a person’s arrest, and (3) if the police feel that evidence might be destroyed if the search is postponed. An example of a constitutionally approved warrantless seizure is the arrest of a suspect running away from the scene of a crime.

Thus, until this case, the requirements for a search were (1) probable cause, and (2) a search warrant or one of the recognized exceptions. The way that police are encouraged to behave according to these Constitutional standards is through the “exclusionary rule” which excludes illegally obtained evidence from being included at trial, or if admitted at trial, it permits a guilty verdict to be overturned.

The constitutionality of stop-and-frisk episodes, which we have all grown accustomed to seeing in TV renditions of police work, was laid down in the case of *Terry v. Ohio* 392 US 1 (1968). In this case, a police officer named McFadden observed three men, one of whom was Terry, apparently

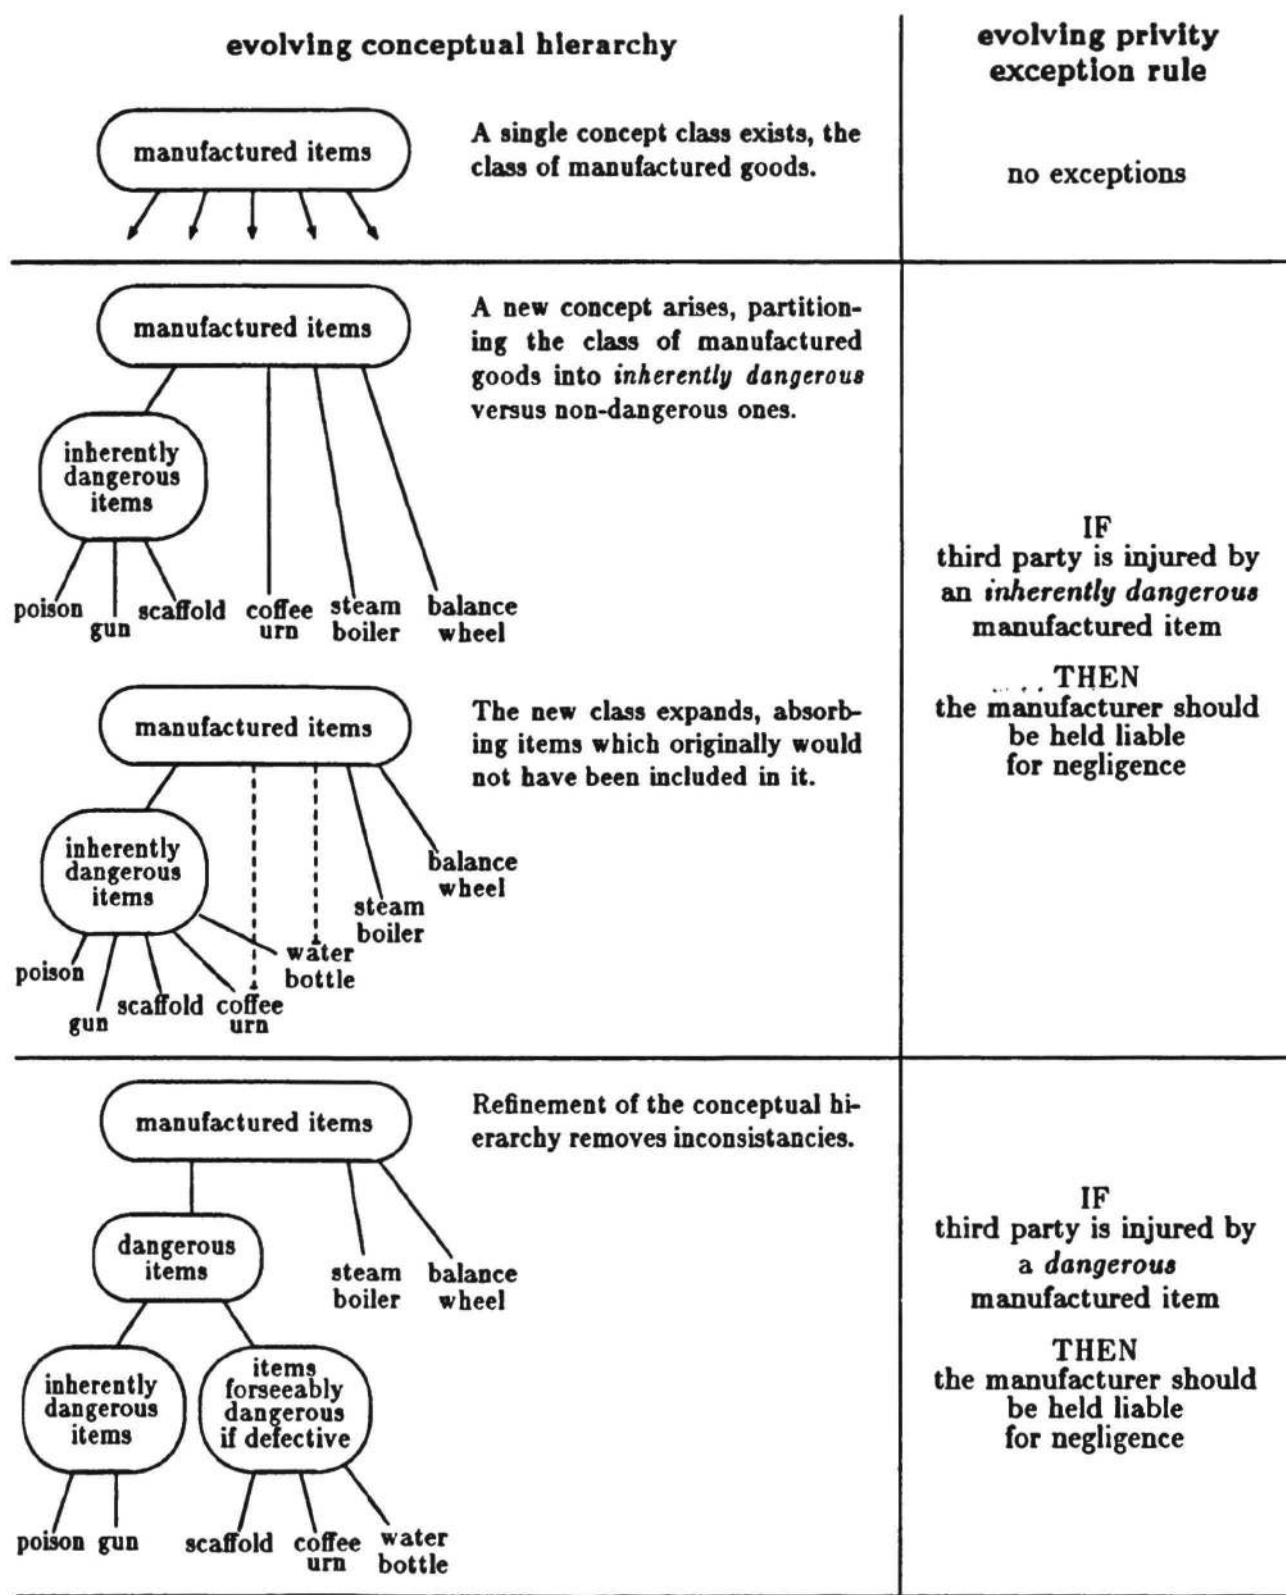


Figure 1: The "Inherently Dangerous" Example

“casing” a store in broad daylight in downtown Cleveland. McFadden’s suspicions were aroused when Terry and his friends studied and walked past a store and then conferred with each other some two dozen times. McFadden suspected they were planning a robbery and since it was daytime, he feared that it would be an armed robbery and that the three men were armed. McFadden accosted and stopped them. When Terry mumbled something unintelligible to McFadden’s inquiries, he grabbed Terry and patted down the outside of his clothing, felt a pistol, and subsequently removed it. He did the same to the second man, with the same result of finding a loaded weapon. The third man was “clean”. Terry and his co-defendant were convicted of carrying a concealed weapon. They appealed, and the Supreme Court held it was an unreasonable search and seizure.

In the Supreme Court opinion, Chief Justice Warren immediately dismisses the prosecution’s argument that Officer McFadden had probable cause (and thus its conclusion that the search and seizure was constitutionally correct according to the doctrine as it now stood). What he had was something less. Thus the constitutional question was whether one can perform a stop-and-frisk (a kind of arrest and search) without probable cause and under what conditions.

The Court said :

“We merely hold today that where a police officer observes unusual conduct which leads him reasonably to conclude in the light of his experience that criminal activity may be afoot and that the persons with whom he is dealing may be armed and presently dangerous, where in the course of investigating this behavior he identifies himself as a policeman and makes reasonable inquiries, and where nothing in the initial stages of the encounter serves to dispel his reasonable fear for his own or others’ safety, he is entitled for the protection of himself and others in the area to conduct a carefully limited search of the outer clothing of such persons in an attempt to discover weapons which might be used to assault him.”

– *Terry v. Ohio*, 392 U.S. 1 (1968), at p. 443

In other words, to reach the conclusion that *P can STOP and FRISK X*, five preconditions must be satisfied :

1. there exists *reasonable suspicion* of X’s criminal activity;
2. there exists *reasonable belief* that X is armed and *dangerous*;
3. P *identifies* himself as a police officer;
4. P makes *reasonable inquiries*; and
5. there does not exist *dispelling evidence*.

This landmark decision for the first time recognized an exception to the requirement that Fourth Amendment searches and seizures must be based on probable cause.

What followed *Terry* were cases fleshing out the scope of the stop-and-frisk rule by discussing its ingredient concepts and the necessity of its preconditions. For instance, four years later in *Adams v. Williams* (1972), the Court allowed the stop-and-frisk exception to cover the somewhat dubious instance of an officer reaching through the driver’s window to reach into the waistband of the driver and seize a gun—not exactly a classic example of stop-and-frisk. In *Adams*, the Court also seems to be trying to stretch the concept of *reasonable suspicion* in order to invoke the rule of *Terry*. In *Terry* the suspicion was based on the officer’s own observations and his own experience from more than 30 years of detective work; In *Adams*, it is not based on the officer’s own observation but on

a tip from an informant who had previously furnished unreliable information. In any case, *Adams* makes quite clear that the conditions necessary for a *Terry* stop are indeed less than probable cause.

Stretching the *reasonable suspicion* concept raises the specter of the stop-and-frisk exception swallowing up the established rules (especially vis-a-vis probable cause) of the Fourth Amendment. This is addressed by Justice Brennan in his dissent where he quotes the lower Appeals Court :

“*Terry v. Ohio* was intended to free a police officer from the rigidity of a rule that would prevent his doing anything to a man reasonably suspected (our emphasis added) of being about to commit or having just committed a crime of violence, no matter how grave the problem or impelling the need for swift action, unless the officer had what a court would later determine to be probable cause for arrest. It was meant for the imminent danger or of harm recently perpetrated to persons or property, not the conventional ones of possessory offenses. If it is to be extended to the later at all, this should be only where observation by the officer himself or well authenticated information shows “that criminal activity may be afoot.” ... I greatly fear that if the contrary view should be followed, *Terry* will have opened the sluiceways for serious and unintended erosion of the Fourth Amendment.”

– 407 U.S. 143 (1968), at p. 147

Justice Brennan emphasizes that *reasonable suspicion* is not the only condition necessary for a *Terry* stop and that for a suspicion to be reasonable it must be based on “good” information. He is also emphasizing that the second condition requiring the existence of *danger* is also necessary. It is well to note that much police work involves crimes of violence or possession, especially of drugs, and that the *Terry* exception was meant to handle certain situations growing out of the former and not as an investigatory tool—that is exactly what warrants are for.

In a separate dissent, Justice Marshall, with whom Justice Douglas joined, makes some meta-level comments on the evolution of this body of law as well as laying out his exception to it :

“Four years have passed since we decided *Terry v. Ohio*, 392 U.S. 1 (1968), and its companion cases, *Sibron v. New York* and *Peters v. New York*, 392 U.S. 40 (1968). They were the first cases in which this Court explicitly recognized the concept of stop-and-frisk and squarely held that police officers may, under appropriate circumstances, stop and frisk persons suspected of criminal activity even though there is less than probable cause. This case marks our first opportunity to give some flesh to the bones of *Terry et al.* Unfortunately, the flesh provided by today’s decision cannot possibly be made to fit on *Terry*’s skeletal framework.”

– 407 U.S. 143 (1968), at p. 148

To recap, what we have seen in *Terry* and *Adams* (and also of *Sibron* and *Peters* mentioned by Marshall), is the hewing out (in *Terry*) of a sizable exception to the Fourth Amendment requirement of probable cause, and essentially a carving out of an exception to the exception of *Terry* (in *Adams*) despite the dissenters’ complaints. That is, after allowing stop-and-frisk in *Terry* when five preconditions are met, the Court then weakens and elevates the first *reasonable suspicion*

condition. So now the antecedent to the action stop-and-frisk rests primarily on the predicate of *reasonable suspicion* whose force has been weakened by stretching its meaning.

This second example of a learning episode is interesting to compare with our first. First of all, the stop-and-frisk episode feels like a more compact learning episode where the first did not. Whereas the courts took fifty or so years to spell out the inherently dangerous exception, the stop-and-frisk exception was well-articulated in one case. The emphasis in the inherently dangerous episode was on progressively refining the concept *inherently dangerous* and the exception to the privity rule, and the situation only gets well-settled in the final case. The emphasis in the stop-and-frisk episode is on arguing the importance of the predicate *reasonable suspicion* and spawning off a discussion much like the inherently dangerous one for the meaning of the concept of *reasonable suspicion*. The next stage of the stop-and-frisk story, which we don't have space to recount, is to work out what is meant by a "stop", especially the parameters of its duration and place of occurrence. Thus both examples illustrate how a rule and the concepts underlying it evolve jointly. The tension in interpretation—and hence learning—seesaws between formulation of the rule and definition of its ingredient concepts.

THE ATTRACTIVE NUISANCE DOCTRINE *LEARNING THE RULE OF THE LEAD CASE*

In our third and last example, we examine an episode from an area of trespass law : the doctrine of attractive nuisance. A longer exposition of this area can be found in [Collins, 1986]. This is another example of the creation of an exception to a well-settled rule. However, instead of focusing on the evolution of the doctrinal rule, we focus on the evolution of the rule of the lead case.

Doctrinal rules are what we have examined in the first two examples. They are rules which summarize the state of the case law as a whole. The rule of a case, on the other hand, refers to what an individual case stands for in the law. The difference is most apparent at the end of a learning episode; initially they may be identical. As example of the difference, the final doctrinal rule for liability of manufacturers to third parties concerns items which are dangerous if defective. The rule of the lead case, *Thomas v. Winchester*, says that only in the case of an item inherently dangerous is there liability.

Attractive nuisance doctrine deals with the liability of a landowner to young children who receive injuries while trespassing. Note this is an exception to the usual situation where a landowner owes no duties to trespassers, except to refrain from intentionally injuring them either in person or through the use of traps. This doctrine evolved through a line of cases in the manner we have already seen in our first two examples. Its lead case is *Sioux City and Pacific Railroad Company v. Stout* 84 U.S. 657 (1873).

What we want to concentrate on in this section is describing how the legal system comes to learn what *Stout* stands for, that is, what the rule of *Stout* is. In its most specific interpretation, the rule of *Stout* is a rule with antecedent conditions identical to the legal facts of the case, and with consequent identical to its legal conclusion. In discussing the rule of a case one does not usually want to state the rule in this most specific form for then it is sure to match almost no other case that comes before a court and thus cannot be applied to these other cases. On the other hand, if one jumps to generalities, especially in the antecedents, one might easily over-generalize. So the

Rissland, Collins

problem becomes just what range of possibilities is to be stood for by the case. This is another CEA-like task. Again, subsequent cases will be the source of positive and negative exemplars.

This phenomenon of learning what a case stands for by examining successor cases was described by Karl Llewellyn :

“Applying this two-faced doctrine of precedent to your work in a case class you get, it seems to me, some result as this : You read each case from the angle of its maximum value as precedent ... Contrariwise, you will also read each case for its minimum value as precedent, to set against the maximum ... The first question is, how much can this case be fairly made to stand for by a later court to whom the precedent is welcome? ... The second question is, how much is there in this case that cannot be got around, even by a later court that wishes to avoid it?”

– Llewellyn, 1930, at p. 69

Llewellyn is clearly speaking of several things here such as the nature of precedent. He previously spoke of the nature of “loose” versus “strict” precedent; the former is reading a rule broadly to cover a large range of cases, the limiting maximally general situation being coverage of all cases (i.e., a rule with no antecedents), and the latter is reading narrowly to cover a small range of cases, the limiting maximally specific situation being coverage of only the case where the rule originally came from. That is, the rule of precedent can allow for a range of interpretations just as we see in version spaces.

In order to see how this process of interpretation through subsequent cases leads to the rule of a case, we look at how the rule of the *Stout* case evolved. A summary of the facts of *Stout* are :

Henry Stout, who was a child of six years, was injured on a turntable belonging to the Sioux City and Pacific Railroad Company. The turntable was in an open space near two traveled roads but with few houses in its vicinity; it was about a quarter mile from the railroad’s station house and was not enclosed or visibly separated from adjoining properties. Stout went to play on the turntable on the suggestion of two older boys, aged nine and ten, one of whom had previously played there. As the boys began to turn the turntable, Henry’s foot was caught between the end of the rail on the turntable as it was revolving, and the end of the iron rail on the main track of the road, and was crushed.

Railroad workers had previously seen boys playing on the turntable and had forbidden them from playing there. This was Stout’s first visit to play there. The turntable had a broken latch at the time of the accident, although this latch was easily undone and whose main purpose was not to lock or guard the turntable but to keep it in position while being used.

The Supreme Court upheld the lower court’s verdict for the plaintiff. It decided that the plaintiff was not himself negligent, either directly or through his trespass, because his tender age exempted him from the same degree of care required of an adult. Thus, the fact that Henry Stout was six years old is most relevant to the case and the language of the opinion suggests that this generalizes to all children of tender years. Further legal conclusions were : (1) the turntable was dangerous and likely to cause injury to children (evidenced by Stout’s injury); (2) the defendant should have anticipated

children would resort to it (evidenced by the fact that even though it was located in a remote place, children had previously played there and had been observed by railroad employees); and (3) the defendant could have taken minor precautions to prevent the accident (e.g., fixed the broken latch, erected a fence).

To use CEA to learn the *rule of Stout*, we need to initialize our version space. See Figure 2. The most specific instance is just the facts and conclusion of *Stout* :

Antecedent : Henry Stout

who was six years old
received a crushed foot from
a dangerous railroad turntable
on SC&PRR property
on which he trespassed
and which SC&PRR could have anticipated
and whose accident was easily preventable

Consequent : SC&PRR is negligent and should have anticipated and prevented this incident.

We also need a most general formulation of the rule of *Stout*. To do this we will simply set all the attributes mentioned in the most specific rule to be the upper or most general numeric or symbolic value. The value hierarchies we use are somewhat arbitrary and are for illustrative purposes only. Thus from subsequent cases, what is to be learned is not only the ingredient conjuncts but also their ranges of values. This sort of top down learning has recently been studied by Fu and Buchanan [1985] in their *RL* program. In particular, all of our subsequent cases will be negative training instances and will serve to specialize the rule.

- Training Instance 1: *Smith v. Hopkins*, 120 F 921 (1902):

Facts: Plaintiff was an 18 year old high school student who was killed while on a picnic with his friends when he walked up on a pile of railroad ties placed on defendant railway's right of way and which collapsed.

Holding: Plaintiff lost. Court reasoned *Stout* did not apply here because plaintiff was 18 and of an understanding mind.

Rule Update: Specialize range of plaintiff's age from *any age* to *young*, the next most specific category excluding an 18 year old.⁴

- Training Instance 2: *Reidel et al. v. West Jersey RR Co.*, 177 F 374 (1910)

Facts: Plaintiff, seven or eight years old, was playing with other children near a railroad track and saw some flowers on the other side of the track. Access to the track was barred with a fence and a bolted gate. The children unbolted the gate and started across. While crossing, plaintiff fell on the "third rail" and was burned by the current.

Holding: For the defendant since, unlike *Stout*, the property was adequately guarded by a fence.

⁴We could also try to specialize *Stout* by adding a conjunct "not including railroad ties". We don't because in its opinion, the court pointed to the difference in the plaintiffs' ages as the basis to distinguish *Stout*.




Figure 1: Initial Version Space for Antecedent of Stout Rule




Figure 1: Most General Final Antecedent of Stout Rule

Update: Specialize prevention predicate to exclude above average care.

Training Instance 3: *Heller v. N.Y., N.H. & H.R. Co.*, 265 F 192 (1920).

Facts: Defendant owned and maintained an electrical wire passing close to a bridge abutment that was covered with a screen. The plaintiff, an eleven year old boy, climbed the abutment some twenty feet or more, climbed through a hole in the screen, reached out with a paint can some 29 inches, touched the wire, and was electrocuted.

Holding: For the defendant since the plaintiff's extraordinary act was not foreseeable.

Update: Specialize the range of anticipation predicate to exclude extraordinary unforeseeable situations.

Given these three training instances, the most general interpretation for the rule of *Stout* is now as shown in Figure 3 (the most specific rule is unchanged).

In summary, what we have tried to illustrate with our third example is how the possible interpretation of the rule of a case is refined through subsequent cases. In particular, we have shown how the rule of the case is specialized through subsequent, negative training instances.

CONCLUSION

What we have shown in this paper is how the law can be viewed as a system that learns from experience in a CEA-like manner. We presented three episodes of such learning. The first, the inherently dangerous example, shows how the law learns both a rule (actually an exception to a rule) and the underlying concept hierarchy simultaneously. The second, stop-and-frisk example,

shows how the law refines a rule (again, an exception to a rule) through discussing ingredient preconditions as well as (recursively) learning the concepts used in the rule. The third example shows how future (negative) exemplars are used to learn the rule of a case.

REFERENCES

- Carbonell, J.G., Michalski, R.S., & Mitchell, T.M. (1983), "Machine Learning : A Historical and Methodological Analysis," *AI Magazine*, Fall.
- Collins, R.T. (1986), "LEX Justice : A Version Space Approach to Legal Reasoning," M.S. Project, Department of Computer and Information Science, University of Massachusetts.
- Fu, L., & Buchanan, B.G. (1985), "Learning Intermediate Concepts in Constructing a Hierarchical Knowledge Base," in *Proceedings IJCAI-9*, Los Angeles, California, August.
- Lakatos, I. (1976), *Proofs and Refutations*, Cambridge University Press, London.
- "Learning and Inductive Inference (Chapter 14)," *The Handbook of Artificial Intelligence, vol. III*, Cohen & Feigenbaum (Eds.), William Kaufmann, Inc. 1982.
- Levi, E.H. (1949), *An Introduction to Legal Reasoning*, University of Chicago Press, Chicago.
- Mitchell, T.M. (1983), "Learning and Problem Solving," in *Proceedings IJCAI-8*, Karlsruhe, West Germany, August.
- Radin, M. (1933), "Case Law and Stare Decisis : Concerning *Präjudizienrecht in Amerika*," *Columbia Law Review*, vol. 33, p. 199.
- Rissland, E.L. (1984), "Argument Moves and Hypotheticals," in *Proceedings First Annual Conference on Law and Technology*, University of Houston, Houston, Texas.
- Rissland, E.L., & Ashley, K.D. (1984), "Hypotheticals as Heuristic Device," in *Proceedings AAAI-86*, Philadelphia, PA., August.
- Utgoff, P.E. (1983), "Adjusting Bias in Concept Learning," in *Proceedings IJCAI-8*, Karlsruhe, West Germany, August.