Of Artificial Intelligence and Legal Reasoning

Citation

Cass R. Sunstein, Of Artificial Intelligence and Legal Reasoning (Public Law & Legal Theory Working Papers No. 18, 2001).

Published Version

http://chicagounbound.uchicago.edu/public_law_and_legal_theory/207/

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:12795544>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

2001

Of Artificial Intelligence and Legal Reasoning

Cass R. Sunstein

Follow this and additional works at: [http://chicagounbound.uchicago.edu/
public_law_and_legal_theory](http://chicagounbound.uchicago.edu/public_law_and_legal_theory)

Part of the [Law Commons](#)

Recommended Citation

Cass R. Sunstein, "Of Artificial Intelligence and Legal Reasoning" (Public Law & Legal Theory Working Papers No. 18, 2001).

This Working Paper is brought to you for free and open access by the Working Papers at Chicago Unbound. It has been accepted for inclusion in Public Law and Legal Theory Working Papers by an authorized administrator of Chicago Unbound. For more information, please contact unbound@law.uchicago.edu.

CHICAGO

PUBLIC LAW AND LEGAL THEORY WORKING PAPER NO. 18

OF ARTIFICIAL INTELLIGENCE AND LEGAL REASONING

Cass R. Sunstein

**THE LAW SCHOOL
THE UNIVERSITY OF CHICAGO**

This paper can be downloaded without charge at:

The Social Science Research Network Electronic Paper Collection:
http://papers.ssrn.com/paper.taf?abstract_id=289789

Of Artificial Intelligence and Legal Reasoning

Cass R. Sunstein*

Abstract

Can computers, or artificial intelligence, reason by analogy? This essay urges that they cannot, because they are unable to engage in the crucial task of identifying the normative principle that links or separates cases. Current claims, about the ability of artificial intelligence to reason analogically, rest on an inadequate picture of what legal reasoning actually is. For the most part, artificial intelligence now operates as a kind of advanced version of LEXIS, offering research assistance rather than analogical reasoning. But this is a claim about current technology, not about inevitable limitations of artificial intelligence; things might change in the future.

I. HYPO and Analogy

The computer on which I am now writing is capable of many impressive feats. Sometimes it talks to me. It can recognize spelling errors and point them out to me. It is astonishing how many words it seems to know. My computer can also find (some) bad writing, and it lets me know when I should rewrite (some) bad sentences. Everyone also knows that the best computer chess player can beat the best human chess player. Fewer people know that an onboard computer system from Carnegie Mellon University has driven a ban almost all of the 2849 miles from Washington, DC to San Diego, California, both day and night, in the rain, and with an average of 63 miles per hour.¹ But this is only the barest tip of the iceberg.

* Karl N. Llewellyn Distinguished Service Professor, University of Chicago, Law School and Department of Political Science.

¹ See David Waltz, Artificial Intelligence: Realizing the Ultimate Promises of Computing (2000), <http://www.cs.washington.edu/homes/lazowska/cra/ai.html>.

Can computers engage in legal reasoning too? Can they do it well? Even better than people? Some grounds for an affirmative answer might emerge from the simple observation that much of legal reasoning is analogical in nature.² In ordinary life, analogical reasoning often takes the form, *White House* is to *President* as *X* is to *Congress*, with the solution consisting of a judgment that *X* is the Capitol Building. The task of identifying good analogies - the kind of task imposed on high school students - seems to be the sort of thing on which computers can excel. If this is right, perhaps computers can do well in law too, simply because legal reasoning is pervasively analogical and based on close attention to past cases. An understanding of the relationship between artificial intelligence and legal reasoning might well illuminate both of these endeavors.

It is best to anchor the discussion in an illustration. Suppose that the rule in state A is that employers can discharge employees “at will,” that is, for any reason or for no reason at all. Suppose that an airline then discharges a copilot for refusing to fly a plane that the copilot believes to be unsafe to fly.³ Is the latter discharge lawful?

Let us assume that there are many analogies in the relevant jurisdiction. Suppose that the courts in state A have created a series of public policy exceptions to the at will rule – that they have said that an employer cannot be discharged for refusing to commit a crime, or for obtaining workers’ compensation benefits, or for cooperating with the police about potential criminal activity on the part of the employer. Suppose too that courts have limited by the reach of the public policy argument by allowing employers to discharge employees for smoking on the premises, for reporting to the Community Credit Bureau about possible regulatory violations by a bank, and for engaging in political activity, outside of the workplace, on behalf of candidates of whom the employers disapproves. Might it be possible for a computer to find, or show, which cases are “most” analogous to the discharge of the copilot, and which cases are “least” analogous to it?

A number of people have attempted to answer this question in the affirmative -- to show the potential role of artificial intelligence in assisting lawyers, and perhaps even in engaging in legal reasoning. I will use as an illustration an extremely interesting book by Kevin Ashley, which makes some striking claims about the role of computer programs in analogical reasoning in

² See Edward Levi, *An Introduction to Legal Reasoning* (1949).

³ See *Bueth v. Britt Airlines*, 787 F.2d 1194 (7th Cir 1986).

law.⁴ Ashley has created a computer program, HYPO, which appears to excel at providing assistance in trade secrets cases. If a HYPO is told about a case, HYPO will, among other things, draw up a set of analogous cases; tell you how they are similar and how they might be distinguished; rank them in order of analogousness; and even give you arguments about how to meet the claim that the cases are different from the case at hand, with citations. Ashley suggests that HYPO is far more useful, in many ways, than Lexis and Westlaw, insofar as the latter simply rely on “keywords” in past cases.

More strikingly, he shows that HYPO’s performance, when confronted with a fact pattern, is not so different from the performance of actual judges. HYPO tends to refer to the same cases and to make the same arguments about how they are similar and different; HYPO even make similar responses to claims that cases are similar and different. But Ashley’s conclusion is still more ambitious: “If lawyers argue with precedents precisely because it is not feasible to prove the right answer by deductive logic, then the goal of a theory of analogical legal argument should not be to explain what the right answer is. Precedential reasoning is interesting precisely because, even without logical necessity, there still may be an ordering to the persuasiveness of arguments. The appropriate goal for a theory of arguing from precedents is to describe that order accurately. . . . Hypo is a step toward such a theory.”⁵

How does HYPO provide “a step” toward a theory of accurately describing the “order” of the persuasiveness of arguments? How would we know if artificial intelligence is actually engaging in legal reasoning?

II. Weak and Strong

A. Hypotheses

What I am going to urge here is that there is a weak and strong version of the claims for artificial intelligence in legal reasoning; that we should accept the weak version; and that we should reject the strong version, because it is based on an inadequate account of what legal reasoning is. We should reject the strong version not because artificial intelligence is, in principle, incapable of doing what the strong version requires (there is no way to answer that question, in principle), but because there is no evidence that, at the present time, any

⁴ Kevin D. Ashley, *Modeling Legal Argument* (1990).

⁵ *Id.* at 254.

computer program is in a position to do what is necessary. To the question, can computer programs engage in legal reasoning, the best answer is therefore: Not yet.

According to the weak version, artificial intelligence can serve as a large improvement on existing computerized services such as Lexis and Westlaw, because well-designed programs are able to assemble an array of relevant cases, to suggest similarities and differences, and to sketch arguments and counterarguments. This is a true and important point. On the strong version, artificial intelligence can now engage in legal reasoning, because a well-designed program can tell a lawyer, or even a judge, what cases are really closest to the case at hand, and what cases are properly distinguished from it. I believe that the strong version is wrong, because it misses a central point about analogical reasoning: its inevitably evaluative, value-driven character.

What is legal reasoning? Let us agree that it is often analogical. In his classic discussion of legal reasoning, Edward Levi rightly emphasizes this point.⁶ But in doing so, Levi makes a serious mistake: He suggests that when engaged in reasoning by example, courts ask what case is “more” similar to the case at hand.⁷ It is much more accurate to say that analogizers in law have to ask which case has relevant similarities to the case at hand. It is more accurate still to say that whether a case has relevant similarities to the case at hand depends on the principle for which the initial case is said, on reflection, to stand. It follows that the crucial step in analogical reasoning consists, not in a finding of “more” similarities, not in establishing “many” distinctions, and not even showing “relevant” similarities and differences, but instead in the identification of a principle that justifies a claim of similarity or difference. Because the identification of that principle is a matter of evaluation, and not of finding or counting something, artificial intelligence is able to engage in analogical reasoning only to the extent that it is capable of making good evaluative judgments.

The point is illuminated by Ronald Dworkin’s influential work on legal reasoning.⁸ Dworkin says that “analogy without theory is blind. An analogy is a way of stating a conclusion, not a way of reaching one, and theory must do the

⁶ See Levi, *supra* note 1.

⁷ *Id.* at 3, note 8.

⁸ Ronald Dworkin, *Law’s Empire* (1985).

real work.”⁹ I think that this view is too simple; an analogy is partly a way of reaching a conclusion, because it helps people to understand and to assess the principles to which they are actually committed. But Dworkin is right to say that analogical thinking cannot get off the ground without some kind of theory or principle, helping to unify or divide the case at hand and the cases that have come before.

We can therefore venture a hypothesis: Since HYPO can only retrieve cases, and identify similarities and differences, HYPO cannot really reason analogically. The reason is that HYPO has no special expertise in making good evaluative judgments. Indeed, there is no reason to think that HYPO can make evaluative judgments at all.

B. An Example

Consider the problem with which I began. Is an airline permitted to discharge a copilot who refuses to fly a plane on the ground that it is unsafe to fly? Let us see how HYPO might be helpful on this question. In a way, HYPO might show, this case like a case in which an employer discharges someone for refusing to commit perjury. In both cases, the employer’s action threatens to injure third parties. On the other hand, HYPO might add, the cases are distinguishable: The discharge by the airplane does not threaten to produce a crime, and in any case the airplane seems to have a legitimate interest in ensuring that safety judgments are made by pilots rather than copilots. Perhaps HYPO will note that in a way, the airplane case is “most” like the decision allowing employees to be fired for reporting possible regulatory violations by a bank. In the airplane case, however, the discharge would have more serious consequences, including massive deaths. Doesn’t this distinction make a difference?

The only way to answer these questions, and to come to terms with the universe of analogies, is to settle on a principle that explains why the case at hand should fall on one or another side of the line. We might say, for example, that an employer is never permitted to discharge an employee as a result of an objectively reasonable judgment, by the employee, that a certain course of action is necessary to save lives. This principle does not conflict with any of the precedents. Or we might say that an employer is always permitted to discharge an employee when the employee has refused to accept a reasonable order from a

⁹ Ronald Dworkin, *In Praise of Theory*, *Ariz. State L. J.* (1997).

hierarchical superior, if that order (a) is job-related and (b) would not require the employee to commit a crime. This principle does not conflict with any of the precedents.

How should a court choose among the two possible principles? How should a lawyer persuade a court to make that choice? It is not helpful to say that the question is which precedent is “closer” to the case at hand. Whether a precedent is closer depends not on a factual inquiry, but on identification of a (normative) principle by which “closeness” can be established. It is more helpful to proceed by asking which principle is actually better. How can we figure that out? An important question is whether the pro-employee principle, in the airline case, would actually improve safety on balance (or instead perhaps impair, as the court of appeals suggested in the case). Another important question is whether the pro-employee principle would disrupt airplane operations, by giving co-pilots a right to veto flights when safety is not much of an issue. It is worthwhile to note that these are empirical issues. Judges may not know how to answer them. But my guess is that HYPO, with its admittedly excellent database, knows even less.

There is yet another avenue for progress, involving an assessment of the proposed principle by seeing if it is inconsistent, from the normative point of view, with anything else that we believe, or to which the legal system would likely commit itself. Here HYPO is not entirely unhelpful, but it can hardly do what needs to be done. I think that Dworkin is correct to suggest that legal reasoning often consists of an effort to make best constructive sense out of past legal events.¹⁰ If analogical reasoning is understood in this light, the analogizer attempts to make best constructive sense out of a past decision by generating a principle that best justifies it, and by bringing that principle to bear on the case at hand. Why should we think that HYPO has any skill at that endeavor?

My conclusion is that artificial intelligence is, in the domain of legal reasoning, a kind of upscale LEXIS or WESTLAW—bearing, perhaps, the same relationship to these services as LEXIS and WESTLAW have to Shepherd’s. A terrific advantage is that the relevant programs can assemble a wide range of relevant cases without turning up so much that does not bear on the problem at hand. But the more extravagant claims on behalf of artificial intelligence in law are based on a crude picture of legal reasoning, one that disregards the need to

¹⁰ See *id.*

root judgments of analogousness, or disanalogousness, in judgments of principle and policy.

III. Three Qualifications

There are three qualifications to what I have said thus far. First, precedents will sometimes sharply constrain the law's room to maneuver. Assume, for example, that an employee alleges that she was discharged for cooperating with the authorities about apparent tax fraud by her employer, and that a previous case says that an employer may not discharge an employee for cooperating with the authorities about apparent drug use by her employer. Sometimes the case at hand cannot plausibly be distinguished from previous cases, because there is no principle that can support the precedent without also producing a certain result in the case at hand. An upscale version of LEXIS, one that has a full stock of precedents on hand, should be able to identify and resolve problems of this kind.

The second qualification is that we cannot exclude the possibility that eventually, computer programs will be able both to generate competing principles for analogical reasoning and to give grounds for thinking that one or another principle is best. Perhaps computers will be able to engage in the kind of empirical testing that is often a crucial (though overlooked) basis for good legal outcomes. Perhaps computers will be able to say whether a particular normative principle fits well with the normative commitments of most people in the relevant community. I have hardly suggested that these are unimaginable possibilities. The possibilities for growth, in the domain of artificial intelligence, cannot be predicted at this exceptionally early stage.

The third qualification is that the weak and strong versions of the claims for artificial intelligence in law, as I have described them, are really poles on a continuum, not a dichotomy, and there is reason to hope for movement from the weak in the direction of the strong. In fact Ashley moves in this direction insofar as he attempts to order cases by determining the strength, or weakness, of one or another connection between the case at hand and the analogies. An effort to specify relevant factors, and to order their importance, is a step in the direction of producing analogy-warranting principles.¹¹ If artificial intelligence is not now able to engage in legal reasoning, it does not follow that it cannot get closer to

¹¹ See Kevin Ashley, *An AI Model of Case-Based Legal Argument from a Jurisprudential Viewpoint* (forthcoming).

doing exactly that. At this stage, there are promising experiments, ones that could be quite helpful to lawyers.

I have emphasized that those who cannot make evaluative arguments cannot engage in analogical reasoning as it occurs in law. Computer programs do not yet reason analogically. But this proposition should not be confused with the suggestion that in the nature of things, evaluative arguments are uniquely the province of human beings, or that computer programs will never be able to help human beings with it, or even to engage on it on their own.

Readers with comments should address them to:

Cass R. Sunstein
University of Chicago Law School
1111 East 60th Street
Chicago, IL 60637
773.702.9498
csunstei@midway.uchicago.edu

University of Chicago Law School
Public Law and Legal Theory Working Paper Series

1. Cass R. Sunstein and Edna Ullmann-Margalit, Second-Order Decisions (November 1999; *Ethics*, v. 110, no. 1).
2. Joseph Isenbergh, Impeachment and Presidential Immunity from Judicial Process (November 1999; *Yale Law and Policy Review* v.18 #1, forthcoming) .
3. Cass R. Sunstein, Is the Clean Air Act Unconstitutional? (August 1999; *Michigan Law Review* #3).
4. Elizabeth Garrett, The Law and Economics of “Informed Voter” Ballot Notations (November 1999, *University of Virginia Law Review*, v. 85).
5. David A. Strauss, Do Constitutional Amendments Matter? (November 1999)
6. Cass R. Sunstein, Standing for Animals (November 1999)
7. Cass R. Sunstein, Culture and Government Money: A Guide for the Perplexed (April 2000).
8. Emily Buss, Without Peers? The Blind Spot in the Debate over How to Allocate Educational Control between Parent and State (April 2000).
9. David A. Strauss, Common Law, Common Ground, and Jefferson’s Principle (June 2000).
10. Curtis A. Bradley and Jack L. Goldsmith, Treaties, Human Rights, and Conditional Consent (May 2000; *Pennsylvania Law Review* v. 149).
11. Mary Ann Case, Lessons for the Future of Affirmative Action from the Past of the Religion Clauses? (May 2001, *Supreme Court Review*, 2000)
12. Cass R. Sunstein, Social and Economic Rights? Lessons from South Africa (May, 2000).
13. Jill Elaine Hasday, Parenthood Divided: A Legal History of the Bifurcated Law of Parental Relations
14. Elizabeth Garrett, Institutional Lessons from the 2000 Presidential Election (May 2001).
15. Richard A. Epstein, The Allocation of the Commons: Parking and Stopping on the Commons (August 2001).
16. Jack Goldsmith, The Internet and the Legitimacy of Remote Cross-Border Searches (October 2001).
17. Adrian Vermeule, Does Commerce Clause Review Have Perverse Effects? (October 2001).
18. Cass R. Sunstein, Of Artificial Intelligence and Legal Reasoning (November 2001).