2004

Capturing the Dialectic between Principles and Cases

Kevin D. Ashley
University of Pittsburgh School of Law, ashley@pitt.edu

CAPTURING THE DIALECTIC BETWEEN PRINCIPLES AND CASES

Kevin D. Ashley*

ABSTRACT: Theorists in ethics and law posit a dialectical relationship between principles and cases; abstract principles both inform and are informed by the decisions of specific cases. Until recently, however, it has not been possible to investigate or confirm this relationship empirically. This work involves a systematic study of a set of ethics cases written by a professional association's board of ethical review. Like judges, the board explains its decisions in opinions. It applies normative standards, namely principles from a code of ethics, and cites past cases. We hypothesized that the board's explanations of its decisions elaborated upon the meaning and applicability of the abstract code principles and past cases. In effect, the board operationalizes the principles and cases. We hypothesized further that this operationalization could be captured computationally and used to improve automated information retrieval. A computer program was designed to retrieve from the on-line database those ethics code principles and past cases that are relevant to analyzing new problems. In an experiment, we used the computer program to test the hypotheses. The experiment demonstrated that the dialectical relationship between principles and cases exists and that the associated operationalization information improves the program's ability to assess which codes and cases are relevant to analyzing new problems. The results have significance both to the study of legal reasoning and improvement of legal information retrieval.

CITATION: Kevin D. Ashley, Capturing the Dialectic Between Principles and Cases, 44 Jurimetrics J. 229–279 (2004).

Theorists in ethics and law have posited a dialectical relationship between abstract normative standards and the concrete cases in which they are applied. The abstract principles inform the analysis and decision of the specific cases, but

---

*Professor, University of Pittsburgh School of Law. The research described here has been supported by Grant No. 9720341 from the National Science Foundation. The author gratefully acknowledges this support. He also thanks Bruce McLaren and Bryan Seigworth for reading drafts of this text and making valuable suggestions.

the decisions of specific cases can specify the meanings of the abstract principles and change the ways in which they are applied in future cases. This work empirically investigates an aspect of the dialectic between principles and cases.[1]

I. PRINCIPLES, CASES, AND THE DIALECTIC

A. Ethics Codes, Principles, and Cases

Many professions, including law, medicine, and engineering, employ professional codes of ethics. The codes help to inform professional role morality, "moral standards and principles that apply especially to one's role as a professional."[2] Although neither always adequate nor complete, such codes are "formal expressions of the ethical norms of the professional . . . community" and "articulate shared standards of professional ethics."[3] This work focuses on the National Society of Professional Engineers (NSPE) Code of Ethics.[4] Each of its seventy-five provisions states a general normative principle governing engineers' professional behavior.

In addition to role morality, professionals employ common and personal morality, "the set of moral ideas shared by most members of a culture or society" or that "we accept as individuals but which are not necessarily accepted by others."[5] The code provisions often incorporate common moral standards of honesty and fairness, but sometimes these different sources of moral obligations conflict. "[T]here are situations in which professional standards may differ from those of personal morality and even the usual standards of common morality."[6]

The ways in which ethics codes inform professional morality and moral decision making are complex. The codes do not specify a simple formula or procedure for making the right decision. "No code, in fact, could supply an algorithm that would give an automatic solution to all ethical problems faced by engineers."[7]

---

1. The author and Bruce McLaren have collaborated on the development of computational models of reasoning with cases in professional ethics. As his Ph.D. dissertation project in the University of Pittsburgh Graduate Program in Intelligent Systems (ISP), McLaren designed and built the SIROCCO (*S*ystem for *I*ntelligent *R*etrieval of *O*perationalized *C*ases and *CO*des) program. *See* Bruce McLaren, Assessing the Relevance of Cases and Principles Using Operationalization Techniques (1999) (unpublished Ph.D. dissertation, University of Pittsburgh Graduate Program in Intelligent Systems) (on file with author) [hereinafter McLaren, *Ph.D. dissertation*], *available at* http://www.pitt.edu/~bmclaren/dissertation.ps.zip (last visited Feb. 26, 2004). This dissertation provides the most comprehensive account of the SIROCCO program and the experiments performed with it.

2. CHARLES E. HARRIS, JR. ET AL., ENGINEERING ETHICS: CONCEPTS AND CASES 33–34 (2d ed. 2000).

3. *Id.* at 8, 14.

4. For cases decided based on the Code, see National Society of Professional Engineers, *Opinions of the Board of Ethical Review*, vol. I–VIII (1958–1998) [hereinafter NSPE BER, Opinions]. Cases from 1976–2001 illustrate provisions of the Code. National Society of Professional Engineers, Cases from the National Society of Professional Engineers Board of Ethical Review, http://www.niee.org/cases/cases_by_year.htm (last visited Feb. 15, 2004).

5. HARRIS ET AL., *supra* note 2, at 8, 14.

6. *Id.* at 14.

7. *Id.* at 11.

Code principles are too abstract to be applied mechanically or deductively. Consider the sample NSPE codes in Figure 1.[8] The provisions deal with engineers' obligations to protect public safety and maintain confidences, but typically, they do not define many of the key terms. "It would be nice to have precise definitions of all these terms [that is, "public health, safety, and welfare, conflict of interest, bribery, extortion, confidentiality, trade secret, and loyalty"] but like most terms in ethics, their meanings are somewhat open-ended."[9]

Moreover, provisions may apply in particular circumstances that appear to recommend conflicting obligations. What happens, for instance, when an engineer, retained by a city to research a dumpsite for possible development, agrees to maintain confidentiality of her results but discovers that improperly disposed toxic wastes could pose a public health risk for certain uses of the site? Even if she advises the city of the need to remediate, if the city balks and decides not to develop the site, must she inform regulators or can she rely on the restrictions of the confidentiality agreement? All three of the code principles in Figure 1 would appear to apply with possibly conflicting recommendations. Her primary obligation is to protect public safety, but suppose the risk is relatively low. Does the code require her to disclose to regulators even remote risks over her client's objections? Given the confidentiality agreement she signed, does the law require her not to disclose the information? How do the engineer's legal obligations under the confidentiality agreement square with her ethical obligations?

Harris et al. describe a multi-stage, casuistic process to guide analysis of ethics problems like this one. Casuistry is "a method for determining the proper moral evaluation of actions in a given case or cases by comparison with reference cases."[10] Casuistic theories of applied ethics focus on comparing an ethical dilemma to paradigms, or cases in which a moral principle is clearly applicable.[11] In an iterative process, a moral actor is led to identify morally relevant facts, questions, and concepts, frame ethical issues, suggest resolutions, identify the consequences of those resolutions for the various participants, and conclude by justifying a resolution based on the consequences the actor identified.[12] Code provisions help identify and frame conceptual and factual issues, such as the degree of risk to public safety or the necessity of seeking consent to disclose a confidence, and focus on their moral implications. In analyzing the ethical issues, moral actors may compare the fact situation to paradigm case examples in a process of line-drawing (that is, locating the problem on a spectrum of scenarios from clearly right to clearly wrong). Alternatively, moral actors may engage in a process of considering alternative ways to resolve conflicts.[13] Presumably, the engineer in the previous example would find it instructive to see and compare

---

8. *See* Figures section at the end of this article.

9. HARRIS ET AL., *supra* note 2, at 45 (emphasis omitted).

10. *Id.* at 74 n.3.

11. *See, e.g.*, ALBERT R. JONSEN & STEPHEN TOULMIN, THE ABUSE OF CASUISTRY: A HISTORY OF MORAL REASONING (1990).

12. HARRIS ET AL., *supra* note 2, at 56.

13. *Id.* at 60–69.

other cases involving engineers who discovered conditions that threatened public safety but had entered into nondisclosure agreements or had been told to keep silent.

Another role that ethics code principles play is as rationalizations to explain and justify the decisions of moral decision makers. "Professional codes of ethics can provide a rationale for professionals to adhere to professional standards even when pressured by others to violate them."[14] The over 400 ethics case opinions involved in this work, available on the Internet,[15] were issued by the NSPE Board of Ethical Review between 1958 and 1998. Following a standard format in each published opinion, the Board recited the facts of the case and the question presented, explicitly cited applicable code provisions, announced its decision, and provided a detailed explanation of how the cited code provisions apply to the facts of the case and justify the Board's stated conclusion.

An example of an NSPE BER opinion is Case 97-5, "Signing a Confidentiality Agreement—Duty to Disclose Danger to the Public Health." The facts are presented in Figure 2. Excerpts of the Board's analysis are shown in Figure 3. In this case, a city had retained an engineer to research a dumpsite, and the engineer had agreed to maintain confidentiality of his results. Subsequently, the engineer discovered toxic wastes that could threaten public safety. The engineer advised the city of the need to remediate, but when the city balked, he decided *not* to inform regulators, relying on the restrictions of the confidentiality agreement. The Board addressed two issues. First, it found that it was unethical for the engineer not to inform the appropriate regulatory agencies of the potential dangers to public health. Second, it found that the engineer did not behave ethically in signing the confidentiality clause, after having been informed by the city that there was a possibility that the site contained hazardous and toxic wastes.

As illustrated in Figure 3, a key feature of the NSPE BER opinions is the use of citation. Like judges, the Board regularly cites, interprets, and applies the relevant code principles and explains their significance to the case at hand. Indeed, the Board publishes these opinions especially to elucidate the Ethics Code.[16]

Also, like judges, the Board regularly cites its own past cases. However, these cases have no formal *stare decisis* effect. The past case illustrates how a cited ethics code principle applied in similar circumstances. The Board usually explains the analogy between a past and current case, states that the analogous case guides the analysis of the new one or distinguishes the current from the past case, and occasionally recommends that a code principle should be changed.

The Board members are not judges; the Board consists of five to seven professional engineers appointed annually. Unlike judges' decisions of litigated cases, the Board's decisions are merely advisory. The Board accepts descriptions

---

14. *Id.* at 15.

15. NSPE BER, *Opinions*, *supra* note 4.

16. The NSPE Board of Directors approved the publication of the BER opinions "with the hope that they will serve to make the profession's ethical principles a living and dynamic force." *Id.* at vol. II, p. iii.

of fact situations as submitted by interested persons and provides an opinion as to whether the conduct described was ethical or not. The NSPE BER cases are also almost always consensus reports; the commentaries only rarely give minority dissenting opinions.[17]

B. Dialectic of Principles and Cases in Ethics and Law

Professional codes of ethics are not immutable. Indeed, they "provide a focus for debate on how professional ethics should be modified. . . .They have been modified in a number of areas as a result of changing perceptions of professional obligations."[18]

There are two ways in which code principles and their meanings change. First, codes may be modified in an intensional way, that is, by explicitly modifying the language of the code to add to or change its moral concepts. For instance, a code provision recognizing the concept of an engineer's duty to the public, as distinguished from his duties to clients or employers, was first introduced into an influential engineering code in 1947. In 1974, that code provision was further revised to require engineers to "hold paramount the safety, health and welfare of the public."[19]

Second, the meanings of code principles may be changed extensionally, that is, by example. An authoritative decision maker like the Board of Ethical Review may determine that particular conduct under specified circumstances does or does not constitute an example of the kind of activity that a code principle proscribes. In explaining why the provision applies to the fact situation and is violated (or not), the Board implicitly specifies conditions for the code principle's applicability. Gradually, this information elaborates upon the meaning of the code principle.

A similar phenomenon may be observed regarding past cases; in explaining why the past case is analogous to or distinguishable from the new problem, the Board implicitly supplies conditions on the applicability of the past case as an example of the application of a code principle.

Both modes of change, but especially extensional change, illustrate a dialectical relationship of abstract normative principles and specific cases, a phenomenon noted by casuistic theorists. In the dialectical relationship, principles inform the decision of specific cases, and cases over time enable ethical principles to evolve in meaning. As explained by noted bioethicist John Arras,

> [W]hile principles may continue to exercise normative force over judgments in particular cases, those very judgments can serve to test, specify, and even disprove particular formulations of principle and theory.[20]
>
> For the casuists, cases provide the considered judgments and paradigms from which moral principles ultimately are derived and to which they must remain

---

17. HARRIS ET. AL., *supra* note 2, at 17.
18. *Id.* at 15.
19. *Id.*
20. John D. Arras, *Principles and Particularity: The Role of Cases in Bioethics*, 69 IND. L.J. 983, 1006 (1994).

> faithful within the creative tension of reflective equilibrium. Cases also put flesh on abstract moral principles, giving them concrete meaning, weight, and specificity.[21]

Similarly, some jurisprudential theories posit a dialectical relationship between legal cases and normative principles, in which the principles are changed as they are applied in specific cases. The legal principles play a role in deciding cases, but the cases also lead to changes in the meaning of the abstract principles. In their models of analogical legal reasoning, Cass Sunstein and Scott Brewer both emphasize the role of comparing a problem and specific cases in discovering, formulating, and testing normative principles for justifying a decision of the problem.[22] According to Sunstein,

> Principles are thus both generated and tested through confrontation with particular cases.

> [The] meaning [of analogies] lies in their use. They are not simply unanalyzed fact patterns; they are used to help people think through contested cases and to generate low-level principles. In this way they have a constitutive dimension, for the patterns we see are a product not simply of preexisting reality, but of our cognitive structures and our principles as well. The principles and patterns we develop and describe are in turn brought to bear on, and tested through confrontation with, other cases.[23]

Brewer describes this process of reflective adjustment between specific examples (as expressed in exemplary propositions) and general normative principles as a common and vitally important instance of example-based reasoning.[24]

> [T]he reasoner faces an active decision about whether to "hold on" to the prima facie judgment that the exemplary proposition is true, and thus make some change to the tentatively held principle with which it is inconsistent, or instead "hold on" to the prima facie judgment that the principle is true, and accordingly declare that the exemplary proposition is false.[25]

The dialectic between general rules and specific cases is characteristic of common law practice.

---

21. *Id.* at 1011.

22. The dialectic relationship of abstract principles and concrete cases is probably consistent with hermeneutic models of legal interpretation. "We simply do not know what we think about a text until we grapple with a specific application of it. If a text is problematic, the best way to test it is by application to a specific problem." William N. Eskridge, Jr., *Gadamer /Statutory Interpretation*, 90 COLUM. L. REV. 609, 676 (1990).

For one of the earliest recognitions in the AI literature of a dialectic in legal reasoning between "proposing of a doctrinal proposition and the probing of it with hypothetical cases" and examples, see Edwina L. Rissland, *The Ubiquitous Dialectic*, in ADVANCES IN ARTIFICIAL INTELLIGENCE, PROC. SIXTH EUR. CONF. ON ARTIFICIAL INTELLIGENCE (ECIA-1984) 367–72 (Tim O'Shea ed., 1985).

23. Cass R. Sunstein, *On Analogical Reasoning*, 106 HARV. L. REV. 741, 775, 779 (1993) (footnote omitted).

24. Scott Brewer, *Exemplary Reasoning: Semantics, Pragmatics, and the Rational Force of Legal Argument by Analogy*, 109 HARV. L. REV. 925, 939 (1996).

25. *Id.* at 940 n.48.

> [R]ules in the common law are contextualized within and emerge from fact situations and fact patterns. These fact situations and fact patterns, as well as the verbal formulations of the rules, play a major role in shaping the scope of common law rules, as applied by subsequent courts.[26]

C. Investigating the Dialectic with AI

The experiment described here focuses on the dialectical effect of systematic efforts of authoritative decision makers, like judges, to justify and explain their decisions over time. In their opinions, the decision makers cite the principles, rules, and past cases, relate them to the factual context of the problem, and explain the inferences they warrant. This paper characterizes that explanatory information as *operationalizing* the meaning of the principles and cited cases.[27] Operationalization supplies conditions of relevance and applicability that can inform future users—including a computer program—how the principles and past cases should be applied in future cases.

We hypothesized that the Board's operationalization information[28] could be both captured and used to improve automated information retrieval. In this experiment, computational techniques in Artificial Intelligence (AI) were employed for this purpose.[29] The AI used here was an information retrieval program, SIROCCO, whose behavior is to retrieve code provisions and cases that are relevant to new ethics problems.

The intelligence of SIROCCO's behavior can be judged in comparison to that of human experts. Specifically, for ethics cases the Board has decided, one can compare the code provisions and cases SIROCCO deems relevant with those that the Board deemed relevant.

A central technique of AI is to represent the kinds of knowledge required to achieve intelligent performance in a form that the computer program can use. By representing the operationalization information the Board supplies in its explanations, SIROCCO "captures" to some extent the dialectical relation between principles and cases. Specifically, SIROCCO represents important aspects of the Board's opinion, including the facts of the case, the Board's decision, and its citations of relevant code principles and past cases. It also represents the Board's operationalization information, including the facts that the Board deemed important in explaining the citation, groupings of principles and cases, and certain other information described below.

---

26. INTRODUCTION TO INTERPRETING PRECEDENTS: A COMPARATIVE STUDY 537 (D. Neil MacCormick & Robert S. Summers eds., 1997).

27. McLaren, *Ph.D. dissertation, supra* note 1, at 8.

28. Operationalization information is the contribution the Board's explanations make over time to the meaning of abstract normative principles and past cases.

29. As a sub-field of computer science, AI endeavors to create computer programs whose behavior, if performed by a human being, would be regarded as intelligent. SEMANTIC INFORMATION PROCESSING, at *v* (Marvin L. Minsky ed., 1968). For an overview of work applying AI in the legal domain, see Edwina L. Rissland, *Artificial Intelligence and Law: Stepping Stones to a Model of Legal Reasoning*, 99 YALE L.J. 1957 (1990).

Once captured, the operationalization information can be turned on and off at will, enabling an experiment to test the hypothesis. Because SIROCCO's success as a retrieval program can be assessed empirically, an experiment was conducted to measure the extent to which operationalization information, contributed over time in the Board's decisions and rationales citing principles and cases, improved the program's information retrieval ability. The degree of improvement corresponds to at least one effect of the dialectic between principles and cases.

As noted, SIROCCO's task is retrieving codes and cases relevant to performing an ethical analysis. The program does not formulate the analysis, make a decision, or compose the rationale. As a result, the experiment does not address all aspects of the dialectical relationship of principles and cases in the full, normative sense described by Arras, Sunstein, or Brewer. Instead, the experiment demonstrates an epistemological contribution the Board's explanations make to assessing relevance. As a by-product of the Board's normative decision making, its explanations elaborate the meanings of abstract normative principles; extensionally, they define and refine principles and cases that may be relevant to the normative analysis of new problems.

D. Guide to the Paper

The remainder of this paper explains SIROCCO's methods for representing operationalization information about principles and cases (Part II), how the program uses operationalizations to retrieve relevant information (Part III), experiments comparing SIROCCO to a full-text retrieval system and to a version of itself lacking the operationalization information (Part IV), the results of the experiments (Part V) and why they have theoretical and practical significance in ethics and law (Part VI).

An extended example, based on NSPE BER case 97-5, shows how the facts of the case, the Board's citations of code principles and cases, and other operationalization information in its opinions are represented in a form SIROCCO can reuse. Then a summary of the program's database of source cases is described. Part III continues the example by comparing the Board's analysis of case 97-5 with SIROCCO's output for the same case when submitted as a target problem. Then, the paper contrasts how a full-text retrieval system like WESTLAW or LEXIS retrieves information about the target problem with SIROCCO's approach of mapping important narrative structures among the target problem and source cases.

In the experiments (Part IV), SIROCCO is compared to the full-text retrieval program and to a version of itself in which the operationalization information is turned-off or "ablated." The results are presented in Part V and interpreted in Part VI. The ablation experiment results are key to testing the hypothesis concerning the dialectic between principles and cases and have significance to legal reasoning. SIROCCO's approach is compared briefly to other work using AI to model legal reasoning. Lastly, the program's potential for normative information

retrieval and as a component of a tutoring system to teach engineering ethics is discussed.

II. REPRESENTING CASES AND CODE PRINCIPLES

A major challenge for an empirical investigation like this one is to develop a way to represent the BER opinions in a form that a computer program can interpret and that is reasonably convenient for human users. The most natural choice of representation, of course, would be simply to employ the natural language texts of the opinions.

Unfortunately, computer programs cannot yet "read" and understand texts of the complexity of opinions like that in Figures 2 and 3, especially over the range of engineering scenarios such cases deal with. One needs to represent fact situations involving engineers, their employers, clients, and regulators in scenarios raising issues such as public safety, confidential information, duties to employers, credit for engineering work, proprietary interests, and honesty in reports.

A. Representing Case Facts

To deal with the representation problem, Bruce McLaren developed a standardized language for representing the cases and created a website where human case enterers could read the case opinions and summarize the factual scenarios and the Board's rationales for decision.[30]

Using the standardized language and website, the main events of a case are represented in a narrative form called a "fact chronology." This is a list of numbered sentences, each describing an event or step in the narrative (that is, a "fact") in roughly chronological order.[31] The case enterer identifies the main actors and objects of the case and specifies the types of actors and objects using a standardized hierarchy for the engineering domain (that is, the "actors and objects hierarchy").[32] The enterer then lists each step in the story (that is, each fact) using a large but limited set of "fact primitives," each of which contains a verb.[33] Finally, the enterer arranges the facts in chronological order and specifies some chronological relationships using a set of ten "time qualifiers."[34] The qualifiers cover the gamut of temporal relationships that can occur between events, including before, after, during, and others.

The fact chronology of Case 97-5 is shown in Figure 4; it summarizes the facts reported in Figure 2. The main actors are Engineer-A, a Principal-Engineer; ABC-Engineering, an Engineering-Firm; and Municipality-M, a Governmental-Body. The case involves a Dumpsite-Reclamation, a kind of Engineering-Project. The chronology comprises ten facts involving Engineer-A and Municipality-M.

---

30. SIROCCO, *at* http://sirocco.lrdc.pitt.edu/index.html (last visited Feb. 14, 2004).
31. McLaren, *Ph.D. dissertation*, *supra* note 1, at 23.
32. *Id.* at 219.
33. *Id.* at 239.
34. *Id.* at 73.

Fact 2, for example, means "Engineer-A **submits a proposal to** Municipality-M **for** Dumpsite-Reclamation." The fact primitives are shown in boldface. In Fact 2, "submits-a-proposal-to-for" is the fact primitive. Facts can be nested. That is, one fact primitive may incorporate another fact primitive as a kind of subordinate clause. In Fact 3, for instance, one fact primitive is nested in another. Fact 3 means, "Municipality-M **informs** Engineer-A **that** Dumpsite-Reclamation **may be hazardous to safety**." The time qualifiers for Case 97-5 are indicated in the right column of Figure 4. Based on the time qualifiers, SIROCCO can draw certain inferences about the ordering of events, for instance, that Municipality-M informed Engineer-A of the potential hazards to public safety *before* Engineer A entered into the confidentiality agreement.

The website provides over two hundred fact primitives plus variations covering such concepts as **recommends products for, accepts a gift of ... from..., provides expert testimony for,** and **signs the plan or report**. For each fact primitive, the site provides templates, constraints, and examples.[35]

Fact chronologies provide a means for linking important events in the case narrative to features of the Board's explanations in a way that the computer program can "understand" and apply. For instance, after the case enterer lists the facts, he specifies which facts correspond to the ethical questions explicitly raised by the Board.[36] Here, Facts 5 and 10 are the "questioned facts" corresponding to the two questions the NSPE BER addresses for Case 97-5 in Figure 3. These questioned facts are indicated in the right column of Figure 4. As discussed below, case enterers can also link the Board's citations of code principles and cases to important specific events.

B. Representing Board's Citations and Explanations

Using the fact chronology, a case enterer summarizes succinctly the information that the Board considered most important in determining which ethics code principles apply, how to resolve conflicts among applicable code principles, and which past cases are relevant.[37] In essence, the case enterer fills out two tables on the website like those in Figures 5 and 6. The first deals with code principles cited by the Board. The second deals with past cases.

For each code principle the Board cites, the case enterer records the facts of the fact chronology that the Board appears to regard as the basis for determining that the code principle is relevant to the case's facts. In Figure 5, for instance, the Board cited the six code principles shown in the left column as relevant to case 97-5. The case enterer notes which code principles support the Board's ultimate conclusions, which code principles the Board held were violated, which facts in the fact chronology—including questioned facts—best explained why the Board

35. *See* Engineering Ethics Transcription Exercise: Fact Primitives, *at* http://www.pitt.edu/AFShome/b/m/bmclaren/public/html/ethics/reference/factframes.html (last visited Feb. 14, 2004).

36. McLaren, *Ph.D. dissertation*, *supra* note 1, at 245–46.

37. *Id.* at 250–59.

regarded the code principle as relevant, and the textual basis for the case enterer's determination.[38]

Similarly, case enterers record information about why the Board regarded past cases as relevant. For each of the four past cases cited by the Board in Case 97-5, as shown in Figure 6, a case enterer noted whether the Board regarded the case as analogous or distinguishable, the facts in the fact chronology that reflect why the Board regarded the case as relevant, and the textual basis for the case enterer's determination.[39]

The critical facts in the Facts columns of the tables in Figures 5 and 6 include the questioned facts (in boldface) relevant to the cited code principle or past case and certain underlined facts. The latter are particular facts that serve to explain why a relevant code provision was violated or not. With respect to cases, the underlined facts are key to the analogue between the case and problem or are the main facts distinguishing the case from the problem.

As noted above, the information collected in Figures 5 and 6 is significant as a record of the particular facts the Board regarded as important in explaining why the cited code principles and past cases were relevant. This information is said to *instantiate* the cited code principle or past case in the context of the problem facts. It explicitly relates the code principle or case to a questioned fact, other critical facts, and the temporal sequence of those facts in the citing case.[40] The temporal sequence comes from the facts' time qualifier information in the fact chronology.

As a result of the instantiation, SIROCCO can tell, for instance, that a code principle is related extensionally (that is, by example) to a real case's relevant facts and chronology in a way that the program can reuse.[41] Specifically, in Figure 5 code provision II-1-a is recorded as having been violated in case 97-5 and has been associated with Facts 6, 8, 9, and 10 and their accompanying temporal relations. Interpolating from the fact chronology, Figure 4, and translating the facts into English, that means that code II-1-a was violated where:

6. Engineer A discovers that the dumpsite reclamation is a safety hazard.
8. Engineer A asks Municipality M to inform the public about the safety hazard.
9. Municipality M disagrees with Engineer A about informing the public about the safety hazard. (Occurs after Fact 8)
10. Engineer A does not inform the public that the dumpsite reclamation is a safety hazard. (***The Questioned-Fact***) (Occurs after Fact 9)

C. Representing Operationalizing Information

Instantiations are important because they can be used extensionally to elaborate the meanings of the code principles and past cases for purposes of retrieving relevant information when analyzing new problems. Using instantiations, the Board, in effect, operationalizes the code principles and past cases by

38. *Id.* at 251–55.
39. *Id.* at 256–59.
40. *Id.* at 19, 27.
41. *Id.* at 83.

providing information with which one can learn how to apply those principles and cases to new fact situations.[42]

The Board employs other techniques to operationalize principles and cases. They are all techniques with which the Board draws various useful conceptual linkages among critical facts, ethics code principles, and past cases. For instance, in the context of a specific factual scenario, an opinion of the Board may group together code provisions or cases of similar import, note conflicts among principles and how they are resolved, and hypothesize factual variations that affect the applicability of code principles.

In his systematic investigation of the NSPE BER's opinions, Bruce McLaren identified nine operationalization techniques in all, listed in Figure 7.[43] As indicated in footnotes 94 through 99, examples of a number of the techniques appear in the NSPE BER opinion for Case 97-5, Figure 3.

Judges appear to use similar techniques in writing legal opinions. In citing statutes and cases, judges often include parentheticals explaining why the statutory standard or precedent is relevant to the fact situation in question. In singling out these critical facts that trigger a standard or form the basis of the analogy to a precedent, judges instantiate the standard or precedent (see techniques 1 and 6). Judges also group multiple legal authorities (for example, statutes or cases) into a particular citation sentence or clause introduced by a citation signal characterizing the extent to which they support a proposition. Signals such as "*accord*" and "*see also*" are used to introduce multiple supporting authorities in a manner similar to grouping code principles (technique 5) and cases (technique 8).

Judges often pose hypothetical variations of facts that would affect the application of a legal standard (technique 2), explicitly revise legal standards from time to time (technique 3), resolve conflicts among legal standards in specific contexts (technique 4), and use past cases to introduce, define, or elaborate upon issues or standards raised in a new case (technique 7). Finally, it is not uncommon for judges to lift applications of these citation techniques from the texts of prior cases and reapply them in the current case (technique 9).

In designing SIROCCO, the operationalization techniques of instantiation, grouping, and reuse (techniques 1, 5, 6, 8, and 9 in Figure 7) contributed most directly to retrieving and selecting relevant code principles and past cases. We treated them as a core set of techniques for purposes of improving retrieval. Each of the five techniques in the core set was implemented computationally in a way that directly affected how SIROCCO retrieves cases and code provisions. The remaining four techniques were also implemented but only provide additional textual information in SIROCCO's explanations of the cases and codes it retrieves.

Case enterers identified information associated with the operationalization techniques in the Board's opinions and represented it for SIROCCO's use in tables like those in Figures 5 and 6. As noted above, the Facts columns identify

---

42. *Id.* at 21–22.
43. *Id.* at 123; *see infra* fig. 7.

the critical facts that instantiate the cited code principles and past cases. These tables also show that the NSPE BER has grouped selected code principles or past cases (techniques 5 and 8). For instance, code principles I-1, II-1, and II-1-a all support the Board's conclusion, as do the three analogous past cases. In reusing operationalizations (technique 9) information is also recorded that the Board has reused a case or code instantiation from a cited case.

The remaining operationalization techniques (2, 3, 4, and 7) are not involved directly in the retrieval-selection task but help the program to explain the significance of retrieved code principles and past cases. Resolving conflicting principles (technique 4) enables the program to explain that the conflicting principles in the problem situation could be resolved in the same way as in a past case. It uses the information in Figure 5 indicating that certain code principles apply, some of which support and some of which conflict with the Board's conclusions. For instance, the instantiation information noted above and recorded in Figure 5 shows that Code II-1-a overrode Code III-4 in the fact situation of case 97-5.

Hypothesizing facts (technique 2) enables the program to point out factual changes that could affect how a code principle would apply. Revising a principle (technique 3) reflects the practice of the Board sometimes to recommend changing the wording of a code principle in light of a case's facts or changing social mores. Also, the Board sometimes explicitly defines a concept or elaborates upon a general issue in the context of a cited case (technique 7). Case enterers quote and store textual excerpts of the Board's comments regarding hypothesizing facts, revising principles and defining concepts (not shown); the program may splice these quotations into the program's explanations.

D. SIROCCO's Case Base

For the work described here, twelve case enterers used the case acquisition web site to represent 184 NSPE cases decided by the Board between 1958 and 1992.[44] These cases constitute SIROCCO's foundational case base. Of the foundational cases, 135 cited at least one code provision dealing with public safety, confidential information, duty to employer, credit for engineering work, proprietary interests, and honesty in reports and public statements (the "selected topics"). The remaining 49 cases did not cite any code principle dealing with a selected topic, but dealt with other code topics including conflicts of interest, honesty in advertising, misrepresentation and omission of facts, criticizing other engineers, and competence and qualifications (the "non-selected topics").[45]

The cases in SIROCCO's foundational case base are heterogeneous in the sense that they are spread across the topics and tend to cite different sets of codes. Code principles for which instantiations have been recorded in past cases range across the selected and non-selected topics including such provisions as I-4,

44. Engineering Ethics Transcription Exercise, *at* http://www.pitt.edu/~bmclaren/ethics/ (last visited Feb. 14, 2004).

45. McLaren, *Ph.D. dissertation*, *supra* note 1, at 26.

requiring engineers to be faithful agents and trustees for employers and clients; II-4-d, restricting participation by engineers in public service in certain types of decision making; III-5, prohibiting engineers from being influenced by conflicting interests; and II-1-c, restricting engineers from taking outside employment.

The case enterers were mostly graduate students specifically hired for the task of case entry. None had any prior familiarity with the SIROCCO program. They reported that it took an average of 2 to 3 hours to transcribe a source case including both the facts of the case (as in Figure 4) and the Board's analysis (as in the tables of Figures 5 and 6). Most of that time was spent representing the Board's analysis.[46]

III. RETRIEVING RELEVANT CASES AND CODES

Given a new case, SIROCCO is designed to retrieve cases and code provisions from its case base that are relevant to analyzing the new case. Its algorithms work with the representations of cases and code principles described above.

Cases must be inputted to SIROCCO in the form of a fact chronology like that of Figure 4. In other words, in order to use the system, one must first use the case acquisition website to create a fact chronology to summarize the narrative of the new case. The fact chronologies are then translated automatically into a form usable by SIROCCO (that is, a program translates the case enterers' inputs into schemas.)

A. Sample Output from SIROCCO

Alternatively, one can submit a fact chronology from an existing case to see how SIROCCO analyzes it.[47] For instance, Figure 8 shows SIROCCO's output upon inputting the fact chronology of Case 97-5. The program generates a list of possibly relevant codes and past cases and a report of the data on which its relevance determinations are based. In order to ensure that SIROCCO's analysis is independent and objective, in generating this output, the program ignores all information (if any) stored in its database of cases concerning the Board's analysis of the submitted case (that is, with respect to Case 97-5, SIROCCO ignores the information in Figures 5 and 6.)

Comparisons of the code provisions and past cases that SIROCCO regards as relevant to the ones cited by the Board in its analysis form the basis of the experiments evaluating SIROCCO's output described in the next section. The NSPE provides a reference list of subjects covered by the code provisions.[48] This information is useful in determining whether different code provisions cited by

46. Kevin D. Ashley & Bruce M. McLaren, *An AI Investigation of Citation's Epistemological Role*, PROC. EIGHTH INT'L CONF. ON ARTIFICIAL INTELLIGENCE & L. 35 (2001).

47. To run the SIROCCO program, visit the SIROCCO website, at http://sirocco.lrdc.pitt.edu/index.html. Although it does not yet support submitting new fact scenarios, one may submit existing NSPE BER fact situations for SIROCCO's analysis.

48. NAT'L SOC'Y OF PROF. ENGRS., THE NSPE ETHICS REFERENCE GUIDE 8 (1996).

the Board or by SIROCCO, in fact, deal with the same subject matter. For convenience, this information was summarized in a "code hierarchy."[49]

A comparison of the Board's citations for case 97-5 versus those of SIROCCO is shown in Figure 9. Significantly, SIROCCO captures the basic conflict, noted by the Board, between code provisions that require an engineer to maintain confidentiality concerning its client's information (that is, code provisions II-1-c, III-4) versus those that require an engineer to act in the interests of public safety (code provisions I-1, II-1-a). SIROCCO's additional suggestions in notes 4 and 5 of Figure 8 call attention to these conflicting obligations. As shown in Figure 8, SIROCCO also identifies case 89-7-1 as relevant, a case the Board said was "most probably closest to the situation faced by Engineer A." In note 6, the program also suggests checking out the cases the Board cited in case 89-7-1.

SIROCCO cites some code provisions and cases not cited by the Board. Perusal of these indicates that one code provision, III-2-b, is nearly equivalent to II-1-b, cited by the Board.[50] According to the code hierarchy, both relate to the same subject matter, a duty to public safety. Each imposes an obligation on engineers to protect public safety when signing or approving engineering documents. Also, the code hierarchy indicates that code provision III-4-a deals with the duty to protect confidential information, the same subject matter as codes II-1-c and III-4, both cited by the Board.

Two other codes cited by SIROCCO, I-4 and II-1-e, are arguably quite relevant to the analysis of case 97-5, even though, according to the code hierarchy, which is based on the NSPE's subject matter index, they do not deal with the same issues as any of the codes cited by the Board.[51] These provisions are applied in two cases that are very similar to case 97-5, which SIROCCO cites but which were not cited by the Board. Case 76-4 applied code provision I-4 (actually a nearly identical predecessor) and case 96-8 cited code provision II-1-e in contexts, like that of case 97-5, posing risks to public safety. Case 93-3-1 also appears to be relevant even though not cited by the Board. The case is distinguishable, and, indeed, SIROCCO does distinguish it in note 3.

This underscores a general problem in attempting to assess the quality of SIROCCO's output by comparing it to that of the Board; the Board simply does not cite all codes and cases that appear to be relevant. Since the Board's output

---

49. McLaren, *Ph.D. dissertation*, *supra* note 1, at 71.

50. Code provision II-1-b states, "Engineers shall approve only those engineering documents which are safe for public health, property, and welfare in conformity with accepted standards." NSPE Code of Ethics from 1981 to Present, http://www.pitt.edu/~bmclaren/ethics/codes/1981-present/index.html (last visited Feb. 14, 2004).

Code provision III-2-b states, "Engineers shall not complete, sign or seal plans and/or specifications that are not of a design safe to the public health and welfare and in conformity with accepted engineering standards. If the client or employer insists on such unprofessional conduct, they shall notify the proper authorities and withdraw from further service on the project." *Id.*

51. Code provision I-4 states, "Engineers, in the fulfillment of their professional duties, shall: Act in professional matters for each employer or client as faithful agents or trustees." *Id.* Provision II-1-e states, "Engineers having knowledge of any alleged violation of this Code shall cooperate with the proper authorities in furnishing such information or assistance as may be required." *Id.*

is the only objective standard conveniently available, however, we use it in the experiments described below.

SIROCCO missed two codes cited by the Board, II-1 and II-1-b, but both of these are nearly equivalent to code provisions SIROCCO does cite, I-1 and III-2-b, respectively, all of which deal with a duty to public safety.

The Board cites three cases that SIROCCO missed, namely cases 89-6, 90-5, and 92-6. Interestingly, the citation to case 89-6 appears to be a misprint in the text of the Board's opinion. From the context of the citation and the fact that case 89-6 appears to deal with entirely different issues, the Board must have meant to cite case 89-7. (Typos in the Board's opinions can also make evaluation of SIROCCO's outputs difficult!) Cases 90-5 and 92-6, on the other hand, are both clearly relevant; SIROCCO should have cited them but did not.

In this particular example, the comparison provides favorable, if not unequivocal, evidence that SIROCCO can retrieve relevant code provisions and cases in an intelligent way. The program succeeded in citing some important, relevant information but also missed some relevant information and cited some irrelevant information. The more systematic experiment discussed below tests SIROCCO's retrieval abilities against the Board's results over *many* problems and compares those abilities to the reasonable alternative methods. This experiment provides much more convincing evidence of the value of SIROCCO's approach. Before getting to the experiment, however, the remainder of this section describes SIROCCO's approach in more detail. It explains how the program performs retrieval and generates this kind of output.

B. How a Full-Text IR System Retrieves

Before describing how SIROCCO works, for purposes of contrast it may be useful to consider how a full-text information retrieval (IR) program would address the task. Managing Gigabytes or MG is a full-text retrieval program, not unlike the WESTLAW or LEXIS systems familiar to the legal profession.[52] At its heart is an inverted index. This database indexes every word appearing in any of the texts in the database after removing stopwords (that is, common words like "the," "a," and "and,") and stemming (that is, removing endings like "ing" or "es"). The full texts of the NSPE ethics code provisions and BER cases would be included. For each word, the inverted index records all of the codes and BER cases in which the word appears, the number of times the word appears in the text, and its frequency of appearance in the text corpus as a whole. A new problem, like the fact situation of Case 97-5, Figure 2, would be inputted as a raw text. MG strips away the stopwords, stems, counts the times that each word appears in the text, and then using the inverted index retrieves all of the codes or cases whose texts contain that word. The program then compares the retrieved texts to the input text and ranks them according to similarity.

---

52. *See generally* IAN H. WITTEN ET AL., MANAGING GIGABYTES: COMPRESSING AND INDEXING DOCUMENTS AND IMAGES (Morgan Kaufmann, 2d ed. 1999) (1994).

In a full-text retrieval program like MG, a new text is compared to existing texts using a trigonometric measure. Each case text is represented as a point in a very large dimensional space. Each different word in the full corpus of texts corresponds to another dimension (that is, x, y, z, a, b, c,...n—the total number of different words in the corpus). A particular text is located as a point in this space and can be specified by the distance along each dimension to get to that point. The magnitude or distance along each dimension is the word's TF/IDF weight, a measure proportional to how many times the word appears in the text (TF) and inversely related to the number of times the word appears in the corpus (IDF). If a text does not have a word, the distance along that dimension is 0.

A "term vector" is an arrow from the origin (0,0,0 ...0) to the point representing the word in this large dimensional space. Figure 10 shows excerpts of a term vector for the fact description of Case 97-5, Figure 2, including the TF/IDF weights for each word. The corpus for which the frequency information was computed contains about 400 NSPE cases. The weights have all been normalized so that the length of the term vector is one. The most distinctive words are "citi" and "dump." The least distinctive words are "firm" and "engin," not unexpected in a corpus of cases involving engineering firms. Some of the words that intuitively capture the import of Case 97-5 are italicized. Such words as "hazard," "confidenti," "waste," "toxic," and "inform" are among the top twelve. Other intuitively relevant words like "contamin," "disclose," and "public" are further down.

The new case text is compared to all of the existing case texts by computing the cosine of the angle between their corresponding term vectors, a straightforward trigonometric calculation. The smaller the cosine, the smaller the angle between the corresponding term vectors and, the MG model assumes, the more similar the texts represented by the vectors. MG outputs a ranked list of the retrieved case texts according to this measure of similarity.

Full-text retrieval systems are easy to set up. Given the MG program, one simply scans the texts of the codes and cases into the system. The inverted index is constructed automatically. No one need read or represent the texts. Nevertheless, they have certain weaknesses. Significantly, term vectors do not represent information about interactions among words in a text.[53] For instance, although combinations of terms like "disclose hazard" are an especially relevant feature of the fact situation of Case 97-5, this distinctive combination does not appear as a feature in the term vector of Figure 10.

C. How SIROCCO Retrieves: Mapping Narrative Structures

By contrast, SIROCCO compares the new problem (that is, the target case) to existing cases in the database (that is, source cases) by comparing the structures of their narratives as summarized in the fact chronologies. We hypothesized that comparing ethics cases as structured fact chronologies is better than comparing

53. Howard Turtle, *Text Retrieval in the Legal World,* 3 ARTIFICIAL INTELLIGENCE & L. 27 (1995).

them as term vectors in so far as the quality of the codes and cases retrieved. The fact chronologies record who did what to whom and when, considerations that are relevant to determining ethical obligations.

Since structure mapping is an expensive operation in terms of time and computational resources, SIROCCO proceeds in two stages, focusing its efforts on the most promising source cases and then on the most important parts of those source cases, the facts instantiating code provisions, and past cases. More specifically, as shown in Figure 11, in Stage 1 SIROCCO retrieves the *N* cases that appear to be most similar to the problem at a surface level. Then in Stage 2, for each of the applicable code provisions in these candidate cases, it attempts to construct a detailed mapping of the instantiation of that code provision in the candidate case to the facts of the problem. Finally, it applies heuristic rules to analyze the quality of these structure mappings and to make judgments about the most relevant codes and cases.

In order to select the best surface-matching source cases in Stage 1, the target and source cases are compared in terms of how well the fact primitives in their respective fact chronologies match (that is, the similarity of standardized verb-like concepts). The cases are represented as points not in a term vector space, but in a different kind of space using "content vectors."[54] In a content vector space, there are as many dimensions as there are fact primitives in the representation language. A particular case is a point in that space and is represented as the distance along each dimension to get to that point. The distance or magnitude along each dimension is the number of times the fact primitive appears in the fact chronology.

Content vectors summarize certain important features of a case. Figure 12 shows the content vector for Case 97-5. The magnitudes along three dimensions are higher ($n = 2$) for fact primitives "may-be-hazardous-to-safety," "is-a-safety-hazard," and "should-be-informed-about-the-hazard-or-potential-hazard." Intuitively, these features capture important features of the case. The last one, "should-be-informed-about-the-hazard-or-potential-hazard," seems particularly important because it combines the concepts of "informing" and "hazard." Term vectors do not capture conceptual combinations like that.

Content vectors enable a computationally inexpensive comparison of the cases for purposes of selecting candidates for the more expensive structure mapping. It is less accurate than structure mapping, but it is quick and helps to focus structure mapping on the most promising candidates.

Similarity is determined by computing the cosines of the angles between the content vectors, the same trigonometric calculation as with term vectors. Various weighting schemes are applied in Stage 1 to favor the source cases whose questioned facts and critical facts (that is, the facts associated with instantiations in the source case of its applicable code provisions) are matched in the target case and that are described at the same level of specificity as the target case (as measured with an Action/Event Hierarchy, described below).

---

54. *See* Kenneth D. Forbus et al., *MAC/FAC: A Model of Similarity-based Retrieval*, 19 COGNITIVE SCI. 141, 162 (1994).

As noted above, in the Stage 2 structure matching, the program focuses on drawing a detailed match between the target case's facts and those facts of the *N* best source cases which instantiated their applicable code provisions.[55] The instantiation information is recorded for each case in the Facts columns in Figure 5.

The program treats the problem of finding the best match as a search through a space of nodes, each of which represents a different possible match of a fact in the instantiation with one in the target case as well as other information including the previous fact matches and their temporal relations. The search proceeds using the A* algorithm, a heuristic search algorithm that evaluates each node in terms of the quality of the match up to that point and an estimate of the cost of achieving a solution from that point. As the search proceeds, the program attempts to match each of the facts of the source instantiation to a corresponding fact in the target case. At the same time, it must maintain a one-to-one and consistent mapping between the Actors and Objects of the source and target.[56]

The program uses the actors and objects hierarchy for this purpose. The matching is also sensitive to whether the temporal ordering of the facts is preserved. Match scores are assigned to reflect the degree of mismatch between the problem and case, and a ranked list of the mappings is passed along to the analyzer.

One of the challenges of using structural mapping is that surface variations in the way facts are expressed could defeat matching even though the underlying facts expressed are the same. SIROCCO embodies two techniques to address this challenge. First, the case entry web site and its limited language of fact primitives, the verb-like concepts, help to standardize the ways in which facts can be expressed. For each fact primitive, the web site offers examples illustrating its use in real cases. Second, SIROCCO addresses the problem of inconsistent descriptions by allowing more generalized matching.[57] An "action-event" hierarchy organizes fact primitives into clusters and abstracts more general characterizations of the primitives.[58] This enables fact primitives in different cases to match at more abstract levels. For instance, according to the hierarchy, "discovers-that" in the fact chronology of case 97-5 is a kind of "knowing-or-believing-something;" it will match fact primitives like "knowing" or "believing" at the more abstract level. The matching algorithm assigns somewhat lower significance or weight to a match at a higher level of abstraction.

The final stage is to analyze the outputs of the first and second stages and organize the program's output. Having generated the structural mappings to the target problem from the *N* best-ranked source case instantiations, the program applies heuristic rules to analyze the quality of the matches, select codes and cases, and explain the output. The heuristics are rules of thumb, not guaranteed

---

55. Since 1958, there have been three versions of the NSPE Code of Ethics with different numbering schemes and somewhat different language. SIROCCO employs a table to relate case references to older and newer code versions. *See* McLaren, *Ph.D. dissertation*, *supra* note 1, at 84–85.

56. Ashley & McLaren, *supra* note 46, at 36.

57. McLaren, *Ph.D. dissertation*, *supra* note 1, at 181.

58. *Id.* at 82–85.

to work but likely to produce a useful result. Using the heuristics, in addition to listing possibly relevant codes and cases as in Figure 8, the program can also explain the basis for its selections. Figure 13 shows excerpts of SIROCCO's detailed explanation of why it recommends as possibly relevant code provision II-1-a and cases 96-8-1 and 89-7-1. The "reasons" show the results of various heuristic selection rules. The "mappings" show the structural matches between the target problem and instantiations associated with the various code provisions and cases.

The selection heuristics recommend code provisions that, for instance, (1) occur more frequently in the top-ranked cases of Stage 1 (that is, "best surface matching cases"), (2) match a high percentage of critical facts in cases citing the code, or (3) are grouped with other codes cited in those cases.

The explanation heuristics enable SIROCCO to construct the list of additional suggestions shown in Figure 8. They find information in the source cases that may be useful for applying the possibly relevant codes and source cases in the target's circumstances. For example, SIROCCO suggests that the codes dealing with public safety (I-1 and II-1-a) may override the code dealing with confidentiality (III-4) in the circumstances of the target case 97-5-1.[59] It makes this suggestion because case 96-8-1 is an example of such a conflict involving the same codes, and there is a good structural match on the facts as shown in Figure 13. In Figure 3, the Board actually did employ such a "resolving conflicting principles" operationalization in its analysis of case 97-5-1, where it resolved competing obligations to public safety and to preserving a client's confidentiality.[60]

The heuristic rules are another way in which SIROCCO implements the operationalization techniques of Figure 7 and makes use of the operationalization information stored in Figures 5 and 6. The selection heuristics correspond to the core operationalization techniques: instantiating codes (1), grouping codes (5), instantiating cases (6), grouping cases (8), and reusing operationalizations (9). The explanation heuristics implement the remaining operationalization techniques: hypothesizing facts (2), revising codes (3), resolving code conflicts (4), and elaborating issues (7).

IV. AN EXPERIMENT TO ASSESS THE USE OF THE DIALECTIC BETWEEN CODE PRINCIPLES AND CASES

As part of his Ph.D. dissertation project, Bruce McLaren conducted an experiment to assess how well SIROCCO retrieved codes and cases and how much of that success could be attributed to its use of the operationalization information embodied in its case representation.[61]

59. *See infra* fig. 8, item 4.
60. *See infra* note 101.
61. McLaren, *Ph.D. dissertation, supra* note 1, at 19–23.

The experiment tested SIROCCO's performance on a test set of cases, all of which were decided after the cases in its database had been decided. As a result, in deciding any of the test cases, it would have been possible for the Board to have cited in its opinions some of the cases in SIROCCO's database. More specifically, SIROCCO's case database comprised the 184 foundational cases, all of which were decided in or before 1993. A set of 58 cases, used as the test set, were selected from a total of 77 cases decided by the Board after 1993 up to the time when the experiment was performed. Forty-four test cases were chosen at random from the 52 cases dealing with the selected topics. Fourteen test cases were chosen at random from the 25 cases dealing with non-selected topics. Two independent case enterers, who were not familiar with the SIROCCO project, transcribed all of the cases into the representation of facts and reasons described above and illustrated in Figures 4, 5 and 6. Their transcriptions were not edited, but were submitted directly to SIROCCO.

The experiment used a set of test cases to compare the performance of six methods:

1. **SIROCCO**: This is the program as described in the previous sections.
2. **Random**: This method selected code provisions and cases at random for each input case.
3. **Informed-Random**: This method selected code provisions and cases at random, but preferred codes and cases that appeared more frequently in the Board's citations in the entire NSPE BER corpus.
4. **MG (Managing Gigabytes)**: This full-text retrieval method, described in the previous section, translated the input cases into term vectors and compared them to the term vectors of code provisions and other cases in a database comprising the texts of the NSPE Code of Ethics and BER cases.
5. **Extended-MG**: This method was like MG, but preferred code provisions cited more frequently in the top X selected cases.
6. **Non-Op SIROCCO**: This version of SIROCCO used content vectors, but all contributions related to the core set of operationalization techniques were turned-off, including instantiation, grouping, and reuse (that is, techniques 1, 5, 6, 8, and 9 in Figure 7).

Random and Informed-Random were intended to provide a baseline for comparing all of the other approaches. The comparisons of SIROCCO to MG and Extended-MG would indicate how well a more traditional approach to textual information retrieval works for retrieving ethics code provisions and cases. Inputs to these programs were textual descriptions of the input case's facts.

The comparison of SIROCCO to Non-Op SIROCCO would show whether and how well SIROCCO makes use of the operationalization information generated as the Board decided cases over time. This is the information most closely associated with the dialectical relationship between codes and cases.

A comparison of this kind is called an ablation experiment[62] because certain knowledge is turned-off or ablated in order to test its contribution to the

62. For a description of a similar experiment, see Edwina L. Rissland et al., *Evaluating a Legal Argument Program: The BankXX Experiments*, 5 ARTIFICIAL INTELLIGENCE & L. 1 (1997).

program's success. Non-Op SIROCCO employed a modified first-stage retrieval process with content vectors, but it did not use operationalization information including code instantiations, case instantiations, grouping codes, grouping cases, or reusing past operationalizations. Specifically, in the first stage, it did not focus on questioned and critical facts in comparing content vectors; it did not use the second-stage structural mapping and thus did not use code and case instantiations to focus the structural mapping; nor did it use the selection heuristics in the analysis phase. Like Extended-MG, Non-Op SIROCCO did prefer to cite codes that appeared more frequently in the list of the *N* top-rated cases. Because this information reflects something about how the Board has applied the codes in past cases, it can be considered a very weak kind of operationalization information, but it is the only such information Non-Op SIROCCO used.

The experiments reported here focused on evaluating the contribution of the core set of operationalization techniques. SIROCCO's selection heuristics use these techniques to generate citations of codes and cases, which can be compared objectively to the citations of the Board. No tests were conducted of the effects of turning off the information associated with the remaining operationalization techniques (that is, 2, 3, 4, and 7, in Figure 7), the ones that help SIROCCO explain the significance of retrieved information. These explanations are expressed in SIROCCO's Additional Suggestions. While an experiment could be designed to evaluate their quality in comparison to the Board's explanations, for instance, by a panel of ethicists blinded as to the source of the explanations, we did not conduct that experiment.

In the experiment, each method processed the test cases one at a time. The code provisions and cases cited by the Board in the test case opinion were deemed to be relevant. For each test case we compared how well the six methods' sets of recommended relevant code provisions and cases covered those cited by the Board. As shown in Figure 14, in comparing a method's citations to the Board's, we computed the magnitudes of the following sets:

- True Positives (TP): The relevant retrieved cases (that is, the overlap between the method's and the Board's citations).
- False Negatives (FN): The relevant cases not retrieved (that is, the Board's citations missed by the method).
- False Positives (FP): The retrieved cases that were not relevant (that is, the method's additional citations beyond those recommended by the Board).

We used the magnitudes of these sets to compute three information retrieval metrics for comparing SIROCCO's and the Board's citations: the method's recall, precision, and F-measure. Recall is the percentage of the relevant cases the method retrieved. Precision is the percentage of the retrieved cases that were relevant. The F-measure is a heuristic combination of precision and recall; it quantifies the degree of citation overlap between the method and the Board.[63] In the F-measure formula, beta (β) represents the relative importance of precision

63. David D. Lewis et al., *Training Algorithms for Linear Text Classifiers*, PROC. 19TH ANN. INT'L ACM SIGIR CONF. ON RES. & DEV. INFO. RETRIEVAL 298 (1996).

versus recall. The value was set at 1.0, assigning them equal importance. With respect to the task of retrieving relevant ethics code and case provisions, it seems as though missing a relevant source and retrieving an irrelevant one are equally serious mistakes.

In comparing SIROCCO's and the Board's citations, certain adjustments should be made. SIROCCO may cite different code provisions that are identical or quite similar to one cited by the Board. In addition, SIROCCO may cite cases that are relevant even if not cited by the Board.[64]

In order to make such adjustments as objectively as possible, the following technique was developed. In computing F-measures for each trial case, an "exact-match" and an "inexact-match" F-measure were calculated. The former indicated the extent of exact matches of codes and cases between the method's and the Board's results. The latter employed criteria for counting certain matches of codes and cases even though they were not exact matches.

The code hierarchy groups "related codes together according to similarity of the issues they address."[65] As noted above, the hierarchy is based on the NSPE's own subject reference list.[66] With respect to code provisions, even if a method-cited code did not exactly match one cited by the Board, if it shared an abstract category with one cited by the Board, it was treated for purposes of the inexact-match F-measure as a match.

With respect to cases, an inexact match was scored according to a citation overlap metric. The score increased with the overlap of code provisions cited in common in the two cases and with the inverse of the length of the citation path between two cases. For instance, if a case directly cited another, the path length is 1. If two cases share a citation to a third case, the path length is 2.[67]

Some of these computations and adjustments can be illustrated in the context of Figure 9, the comparison of citations by the Board versus SIROCCO for case 97-5. The three sets for comparing citations, *TP*, *FP* and *FN*, are as follows:

TP = {I-1, II-1-a, II-1-c, III-4, 89-7-1}.
FP = {I-4, II-1-e, III-2-b, III-4-a, 76-4-1, 93-3-1, 96-8-1}.
FN = {II-1, II-1-b, 89-6, 90-5, 92-6}.

For purposes of computing the exact-match F-measure, the magnitudes were computed as follows: $|TP| = 5$. $|FP| = 7$. $|FN| = 5$. $R = 5 / 10 = 50\%$. $P = 5 / 12 = 42\%$. $F = (2 * .50 * .42) / .50 + .42 = .46$. For purposes of computing the inexact-match F-measure, using the code hierarchy, the code provisions in *FP* and *FN* were each checked for an abstract match with codes in *TP*. Abstract matches were found for II-1, II-1-b, III-2-b, and III-4-a, but not for I-4 and II-1-e. As a result, *FP* and *FN* are each reduced by 2. In addition, each case in *FP* and *FN* is checked for an abstract match (that is, case or code citation overlap) with the one case in *TP*, 89-7-1 and further reductions are made as appropriate. For instance, cases 76-

64. McLaren, *Ph.D. dissertation*, *supra* note 1, at 153.
65. *Id.* at 83.
66. NAT'L SOC'Y OF PROF. ENGRS, *supra* note 48, at 8.
67. Ashley & McLaren, *supra* note 46, at 38.

4-1 and 89-7-1 both cite some code provisions in common. Each reduction to *FP* and *FN* increases *R*, *P*, and *F*.

In comparing the six methods against the Board's citations using the F-measure, the data generated turned out not to be a bell-shaped distribution. Because the Board's citations were relatively sparse, the recalls, precisions, and F-measures of all the data sets were highly variable. Accordingly, we applied a "nonparametric bootstrap procedure" to compare the six methods. This technique involves "resampl[ing] from the original data [using random sampling with replacement]...to create replicate datasets" with which to make the comparisons.[68] This approach is appropriate because the data observations (that is, F-measures) were independent and the likelihood of seeing any particular F-measure would not vary from observation to observation.[69] In this procedure, each data set comprised 58 F-measure readings (that is, one reading for each test case). Each of the six methods yielded two data sets, one for exact and one for inexact matching. For each data set, 58 readings were drawn at random and replaced into the data set and the mean was calculated. This step was repeated 100,000 times. Summary mean F-measures were obtained for the 2.5 and 97.5 percentiles. Significance values (p-values) were computed "by doubling the proportion of the 100,000 mean differences between two methods that are less than or equal to zero."[70]

V. RESULTS OF THE EXPERIMENT

The results of the comparison of the mean F-measures for all six methods over all of the trial cases are shown in Figure 15.[71]

SIROCCO's performance was better than any of the other five methods using either exact-match or inexact-match. All of these differences were statistically significant except with respect to the inexact-match using Extended-MG. That is, if we reject the null hypothesis that SIROCCO's performance was the same as that of Random, Informed-Random, MG, and Non-op SIROCCO using either exact- or inexact-match the probability is less than .05 that such a conclusion would be in error. The same was true regarding Extended-MG using exact-match. SIROCCO performed better than Extended-MG with inexact-match at $p < .052$, very nearly significant at the conventional level.

We conducted an additional experiment to further compare SIROCCO and Extended-MG. Two graduate students in ethics examined the extra code and case citations suggested by either SIROCCO or Extended-MG for the test cases. They were blinded as to the source of the suggestions. For each extra citation they opined as to whether it was reasonably relevant to analyzing the test case. Written instructions to the evaluators defined "reasonably relevant" as whether it would "be reasonable for an experienced ethical reasoner to reference the cited item in an argument answering the question raised by the case."

---

68. A.C. DAVISON & D.V. HINKLEY, BOOTSTRAP METHODS AND THEIR APPLICATION 2, 22 (1997).

69. McLaren, *Ph.D. dissertation*, *supra* note 1, at 143.

70. *Id.* at 144.

71. *Id.* at 146; *infra* fig. 15.

The recall, precision, and F-measures for SIROCCO and Extended-MG were revised accordingly and compared. As shown in Figure 16, SIROCCO's performance turned out to be significantly better than that of Extended-MG using either exact-match ($p < .001$) or inexact-match ($p = .01$).

VI. DISCUSSION AND CONCLUSIONS

A. Ablation Experiment Result

SIROCCO outperformed Non-Op SIROCCO in the ablation experiment. SIROCCO's core operationalization techniques allowed it to make better predictions of the principles and past cases that are likely to be relevant in the analysis of new cases than it made without that information. Both methods employed the same case representation and compared cases in the first stage using content vectors. The critical difference was that Non-Op SIROCCO did not use the core set of operationalization techniques: no code and case instantiations, no code and case groupings, and no reuse of operationalizations.

The experiment provides strong evidence of the epistemological contribution of the Board's core operationalizations and of the appropriateness of SIROCCO's representation for capturing this contribution.

This result shows that the Board's explanations of its decisions over time in the foundational cases flesh out the definitions of the code principles by example. Intuitively, the instantiations and groupings of codes and cases are part of the dialectic between principles and cases. As the Board makes and explains citations, it conceptually links critical facts, code principles, and past cases. In the process, it provides applicability and relevance conditions for the code principles and past cases, information which is valuable for retrieving and applying codes and cases to new problems.[72]

The experiment confirms that SIROCCO represents and can reuse this operationalization information when analyzing new cases. The program captures the Board's contribution in extensionally defining the abstract codes and utilizes that information to improve retrieval. In either type of matching, exact or inexact, the effect is substantial. In the exact-match, the ablated operationalization information accounted for 38% of SIROCCO's mean F-measure. In the inexact-match, it accounted for 33%.[73] The effect may be greater for the exact-match because exact-match is a more demanding task. The program must pick the precise code or case cited in the Board's analysis. The less constrained inexact-match increases the chances for matching by "lucky guesses." It follows that the core operationalization techniques make a greater contribution in the exact-match test. In a future line of investigation, additional ablation experiments could allow

---

72. We also suspect that the Board's decisions flesh out temporal event-ordering conditions under which a code provision may reasonably apply. In a subsequent experiment we plan to investigate this contribution empirically.

73. Each percentage is computed as follows. Let s be the mean F-measure for SIROCCO and n be the mean F-measure for Non-Op SIROCCO. The contribution of the ablated operationalization information c is $(s - n)/s$. Using the data from Figure 15, for the exact-match, s is .21, n is .13, and c is 38%. For the inexact match, s is .46, n is .31, and c is 33%.

one to determine the relative importance of each kind of operationalization information in the core set.

The capture of the operationalization information results from the case-enterers' recording information about the Board's citations as in Figures 5 and 6. SIROCCO does not involve the case enterers or program designers in defining special rules for deducing how the abstract principles apply to realistic scenarios. This would be an alternative representation approach, but such an effort would be impossibly difficult and would result, in any event, in a body of non-authoritative rules.[74] Instead, the representation of operationalization information in SIROCCO is a natural result of the case entry process.

A skeptic may argue that the ablation experiment's result is neither surprising nor significant. Because SIROCCO cannot understand the texts of either code principles or past cases expressed in natural language, the cases must be represented in a much less expressive artificial language that can be manipulated by machine. Given these impoverished case representations, the skeptic may assert, if one compares the program's results with and without taking into account some additional relevant information, the so-called operationalization information, the program's performance is bound to improve.

According to this argument, the result does not necessarily imply anything about what happens when the ethics code provisions, cases, and concepts are represented and understood in natural language text. People do not use SIROCCO's impoverished case representation. For them, the terms have meaning. Anyone who understands the texts of the code principles, past cases, and new case facts is likely to perceive which codes or past cases are relevant, with or without operationalization information.

The argument is plausible, but there are a number of reasons why the ablation experiment is meaningful nonetheless. The fact that the program does significantly better with the operationalization information suggests that this is the relevant information for fleshing out the meanings of the code provisions and past cases. When the Board explains its citation to a code principle or to a past case, it explains which facts in the specific fact situation it deems critical with regard to their application to a questioned fact. At the same time, it makes links among scenarios where it is reasonable to consider that the code may apply and provides judgments as to which of these scenarios are analogous and which are distinguishable.

As a result, even where codes and cases are represented in natural language terms, one would expect operationalization information to improve performance. The operationalization information is both relevant and specific. The meanings of the terms of the code principles alone are not likely to imply information as specific as in the operationalizations. The code provisions are very abstract; there are few clues about the sorts of scenarios to which they may apply.

As a practical matter, many people do not know or understand the texts of either the ethics codes or the past cases. While some people may be particularly good at applying the codes, it is not clear whether they are drawing from the

74. JONSEN & TOULMIN, *supra* note 11, at 6–10.

meanings of the code provisions' terms or whether experience may have trained them about the kinds of scenarios associated with that language. If the latter, then the experiment provides an example of some experiential information that is effective in fleshing out the code's abstractions and a possible mechanism for applying that information. In any event, there is no way to conduct such an experiment with natural language as the representation. Computer programs do not understand it, and human beings cannot turn certain information on and off as computers can.

B. Significance for Legal Reasoning

Given the common law significance of precedent in interpreting legal rules and principles, the results of the experiment should come as no surprise to attorneys or judges. Although the terminology of "operationalization" may be unfamiliar in law, it is likely that judges operationalize abstract legal rules and precedents in much the same way as the Board does. Instantiation and other operationalization techniques are likely to flesh out the meanings of abstract *legal* standards and precedents much as they do ethical principles and past cases.

The NSPE ethics cases and legal cases have much in common. Both involve concrete factual scenarios, present normative issues to be decided by authoritative decision makers who must rationalize their decisions in terms of abstract normative standards, including ethical principles and open-textured legal rules. In writing opinions, ethical or legal, the Board and judges both explicitly cite the normative principles and past cases. They explain the relevance of the citations in the current factual context. They develop analogies and distinctions between a problem situation and cited cases. They resolve competing standards in a given factual context, pose hypothetical factual variations, and group together conceptually related standards and cases.

The experimental results, then, are significant to the study of legal reasoning, not because they identify a new phenomenon, but because they provide empirical confirmation of a phenomenon long recognized. As far as is known, the dialectic between principles and cases in either ethics or law has never been captured in the sense of having been represented, applied, and demonstrated empirically. Here an AI model has enabled exactly that.

Despite the similarities between the NSPE cases and legal cases, there are some differences. As noted previously, past ethics cases may be persuasive, but they have no *stare decisis* effect. The Board's opinions are purely educational; they do not have a binding legal effect. Also, strictly speaking, decisions in engineering ethics cases are not as constrained to binary conclusions (for example, plaintiff winning or losing), as is often the case in law. While the Board tends to frame the moral questions presented in terms of whether the actions were ethical or not, it often suggests ways in which the protagonist might have avoided the ethical dilemma altogether or ways to correct an unjust action after the fact. Presumably, the Board deems it important for engineering practitioners to recognize the conflicting values in moral dilemmas and to learn to apply "creative

middle way" solutions to those situations (that is, resolutions that at least partially reconcile the conflicting values).[75]

As a practical matter, the legal and ethical domains also differ in terms of access to case data. Full-text retrieval services and a vast array of substantive conceptual indexes put legal practitioners in touch with an enormous body of on-line cases. The engineering ethics domain, by contrast, offers professionals much less sophisticated resources with far fewer case examples and opinions and far less conceptual indexing.

SIROCCO focuses on how on-line case retrieval for engineering ethics can be improved. Specifically, computationally representing the operationalizing links between abstract principles and factual case narratives improves retrieving relevant codes and cases. Because the same operationalization techniques apply in law, conceivably, the same computational techniques could be applied in turn to improve legal information retrieval.

C. Comparison to Other Work

Other AI and law models have represented legal claims and statutory legal concepts extensionally in terms of cases and examples. The HYPO program provided a model of comparing legal cases in terms of dimensions, stereotypical patterns of fact that tended to strengthen or weaken a side's argument concerning a legal claim.[76] CABARET represented tax concepts involving the home office deduction using cases and HYPO-style dimensions.[77] GREBE represented statutory terms in worker's compensation law using portions of judges' explanations why the terms were satisfied or not in particular cases.[78]

Still other AI and law models have focused on the phenomenon that legal opinions, via their citations, create a conceptually linked network and have modeled such networks for purposes of information retrieval. In the FLEXICON program, citations to cases and statutes were included in case document profiles along with legal concepts and factual terms. All were treated as additional terms for purposes of applying TF/IDF weighting.[79] SCALIR, another legal information system, linked legal opinions to other documents that it cited or that shared terms in common.[80] Upon entering a query using terms in the network, the documents connected to the input terms in the network were activated and the activation spread throughout the network to other linked terms and documents. Each

---

75. HARRIS ET AL., *supra* note 2, at 64–72.

76. *See generally* KEVIN D. ASHLEY, MODELING LEGAL ARGUMENT: REASONING WITH CASES AND HYPOTHETICALS (1990).

77. Edwina L. Rissland & David B. Skalak, *CABARET: Rule Interpretation in a Hybrid Architecture*, 34 INT'L J. MAN-MACHINE STUD. 839 (1991).

78. *See* L. KARL BRANTING, REASONING WITH RULES AND PRECEDENTS: A COMPUTATIONAL MODEL OF LEGAL ANALYSIS (2000).

79. Daphne Gelbart & J.C. Smith, *FLEXICON: An Evaluation of a Statistical Ranking Model Adapted to Intelligent Legal Text Management*, PROC. FOURTH INT'L CONF. ON ARTIFICIAL INTELLIGENCE & L. 142, 144 (1993).

80. Daniel E. Rose & Richard K. Belew, *A Connectionist and Symbolic Hybrid for Improving Legal Research*, 35 INT'L J. MAN-MACHINE STUD. 1, 11 (1991).

document's degree of activation was a function of its terms' TF/IDF weights. The documents were ranked for retrieval according to their degree of activation. Given a problem scenario, another program, BankXX searched a network of interrelated precedents, legal theories (that is, tests expressed in terms of dimensions) and concepts, opportunistically gathering the best cases and arguments according to a set of argument evaluation factors.[81]

SIROCCO is the first such program to model how decision makers' explanations of their decisions in ethical cases operationalize abstract ethical principles and the first to demonstrate empirically an epistemological contribution of the dialectic between principles and cases. It represents the operationalization information computationally and uses it to improve information retrieval of relevant ethical principles and cases. In addition, unlike SIROCCO, none of the above programs represent problems and cases as narratives. SIROCCO's representation enables linkages from abstract concepts like principles into specific events and their chronological relationships, which in turn can be mapped structurally into new problems. Thus, although the other programs represent case examples and citation links, they cannot represent the same kind of narrative information at the appropriate grain size for structural matching to new scenarios.

Although GREBE does not represent cases as narratives, it does represent judges' explanations of why a case's facts are or are not instances of open-textured terms in the relevant statutes and supports structural mapping. Its case-based and rule-based model of causal and evidential relations in the domain of worker's compensation cases is more finely grained than SIROCCO's and enables it to generate more detailed case analyses and arguments. SIROCCO's use of code instantiations, its two-stage retrieval algorithm, and its use of the A* search algorithm in the second stage were all inspired by GREBE's explanation-based exemplars and criterial facts.[82]

SIROCCO has three main advantages over GREBE. First, as noted above, unlike SIROCCO, GREBE does not represent cases as temporally ordered, narrative descriptions of events. SIROCCO's representation language provides a total of 190 actions and events compared to from 70 to 90 in GREBE. SIROCCO's representation includes formally defined temporal relations among facts, and a well-defined algorithm for matching temporal relations. GREBE does not account systematically for temporal relations, and does not take temporal ordering into account in structural mapping. Since temporal ordering appears to be important in determining moral obligations (for example, one cannot have an obligation to disclose information before one learns of the information), and is often important in determining legal obligations, this is an important difference.

Second, SIROCCO provides practical means to encourage consistency in describing cases, and is designed to tolerate a much greater level of inconsistency in descriptions of scenarios. GREBE's attempts at structural mapping will fail unless a particular fact is represented the same way in whatever cases it appears.

---

81. Edwina L. Rissland et al., *BankXX: Supporting Legal Arguments Through Heuristic Retrieval,* 4 ARTIFICIAL INTELLIGENCE & L. 1, 19 (1996).

82. *See* BRANTING, *supra* note 78.

SIROCCO's structural mapping can tolerate greater variety in the way the fact is expressed. SIROCCO's case acquisition web site provides a standardized language for representing cases with guidelines and examples. Its techniques for matching facts abstractly reduce the need for perfect consistency. SIROCCO can match actions and events at various levels of abstraction, something that GREBE could not do. Whereas the program designer represented GREBE's 41 cases,[83] 12 case enterers represented SIROCCO's 242 cases, none of whom had been involved in developing the program. SIROCCO's cases appear to span a wider range of factual scenarios and normative issues than GREBE's.

Third, while GREBE captures something like instantiations for a set of open-textured statutory terms, the work on SIROCCO identifies and implements more techniques for operationalizing abstract principles and past cases. GREBE does not address abstract principles like the ethics code provisions, nor does it have equivalents to SIROCCO's operationalization techniques for grouping principles or cases, resolving conflicts between principles, instantiating cases, elaborating principles, hypothesizing facts, revising principles, and reusing operationalizations. While an ingenious evaluation showed that GREBE's arguments were comparable to those of law students, the GREBE experiment did not isolate or provide evidence for the epistemological contribution of operationalization information in fleshing out the meanings of principles and cases.

D. For Normative Information Retrieval and Instruction

The ubiquitous World Wide Web affords opportunities for fielding on-line resources of ethical and legal information. The question is whether ordinary information retrieval, even though delivered via the Web, is sufficient to help practitioners make better decisions or learn skills of ethical analysis.

The NSPE makes its professional code of ethics and BER cases available online via the Web as a pedagogical aid. Nevertheless, the ethics code is so comprehensive and abstract that engineers and students may find it of limited utility in practical decision making. Publishing advisory opinions is intended to flesh out the meanings and conditions for applying the abstract principles but at the cost of even more material for engineers to search and read.

The work described here points to the design of an intelligent aid for retrieving abstract normative standards and relevant past cases in fields like professional ethics and law. Conceivably, intelligent access on-line to *relevant* standards and examples, given the facts of a problem situation, may help guide decision making and aid pedagogy.

The results of comparing SIROCCO to MG and Extended-MG show that the former performs better than its most likely competitor, a full-text retrieval method not unlike LEXIS or WESTLAW. The fact that SIROCCO performs significantly better than a term vector approach is clear evidence of the value of its case representation.

---

83. *See id.* at 66–67.

The fact that SIROCCO outperformed Extended-MG, which, in turn, outperformed MG, is also significant. Since Extended-MG makes use of what might be termed a weak kind of operationalization information (that is, it prefers to cite code provisions cited more frequently in the top X selected cases), its improvement over MG also supports the hypothesis.

From a practical viewpoint, however, SIROCCO's advantage in performance over MG and Extended-MG comes at a significant cost. Constructing the latter programs' inverted index is a simple and comparatively quick mechanical process. Filling out SIROCCO's representations for the fact chronology and aspects of the Board's rationalization is much more time-consuming and requires human interpretation.

For a practical information retrieval tool, a more convenient mode of case entry is required. Automatically or semi-automatically representing case facts, however, remains a long-term research goal. SIROCCO's representation may provide a conceptual framework for experimenting with natural language extraction techniques that may someday fill out the representations automatically.[84] We hope that SIROCCO's use of a limited language focusing on the important verbs in a domain may help in this regard.

In addition, currently we are comparing SIROCCO to a different kind of full-text retrieval approach using Latent Semantic Analysis (LSA). LSA is another computational technique for representing textual cases as points in a vector space, one in which case texts may be represented automatically as with term vectors, but whose dimensions may capture interactions among concepts better than term vectors.[85]

Another research goal is to automatically represent the Board's explanations of its citations. It is not difficult to automatically identify places where the Board has cited code principles or past cases. SIROCCO uses text manipulation software to accomplish this. Relating those citations to important factual features in the case's fact situation, however, presents the same problems as automatically representing facts. It might be an easier task if the Board and judges cooperated with efforts to "mark-up" the content of their opinions (that is, applying machine-readable tags to parts of the text to indicate their content such as "facts consistent with conclusion" or "facts contra conclusion"). Considerable activity is under way to support substantive mark-up of legal documents, including cases.[86] The work described here indicates the complexity of the kind of substantive mark-ups that would be required to support instantiating principles and other operationalizations. SIROCCO's representation captures conceptual information such as the narrative of events in a case, their temporal relations, and the dialectic between

---

84. Stefanie Brüninghaus & Kevin D. Ashley, *Improving the Representation of Legal Case Texts with Information Extraction Methods*, PROC. SEVENTH INT'L CONF. ON ARTIFICIAL INTELLIGENCE & L. 42 (2001); Jodie J. Daniels & Edwina L. Rissland, *Finding Legally Relevant Passages in Case Opinions*, PROC. SIXTH INT'L CONF. ON ARTIFICIAL INTELLIGENCE & L. 39 (1997).

85. *See* Scott Deerwester et al., *Indexing By Latent Semantic Analysis*, 41 J. AM. SOC'Y FOR INFO. SCI. 391 (1990).

86. *See, e.g.*, LegalXML, About LegalXML, *at* http://www.legalxml.org/about (last visited Feb. 14, 2004).

principles and cases. As the experiment shows, the representation of this information makes a significant difference in the quality of the codes and cases retrieved. Developing a standardized set of annotations in a legal mark-up language to capture this information would be a subtle and difficult exercise.

In the shorter term, we plan to incorporate SIROCCO into a tutoring environment for practical ethics. The prototype PETE program leads students through the Harris process of analyzing ethics problems.[87] An extended PETE will lead small groups of students (who may be at different places or even participating at different times) in analyzing and discussing ethics cases via the Web. SIROCCO will serve as an on-line resource for participants to search for relevant code principles and past cases. From a pedagogical viewpoint, it may be beneficial for students to manually represent problem scenarios as narratives of temporally ordered events in a fact chronology. It would induce them to consider the facts of a case more carefully. As an aid to understanding cases, generating fact chronologies and citation information is not unlike first-year law students' case-briefing exercises. SIROCCO's explanations of its outputs could also be pedagogically useful. Full-text retrieval schemes alone cannot generate such explanations.

An important goal of an ethics tutoring system is to teach students to compare problem situations to cases in a process of line-drawing[88] and to consider the normative consequences of hypothetical variations of facts (for example, as in operationalization technique 2 in Figure 7). As noted above, the HYPO program compares legal cases in terms of dimensions, stereotypical patterns of fact that tended to strengthen or weaken a side's argument.[89] HYPO applied its model in posing hypothetical variations of trade secret problems to illustrate their consequences for the legal analysis. Reasoning with dimensions or factors appears to be useful in modeling practical ethical reasoning as well.[90] Dimensions would be helpful, for instance, in constructing a more detailed computational model of the Board's posing of hypothetical fact variations and in modeling line-drawing. In future work, the author hopes to combine HYPO's dimensional model with SIROCCO's narrative one.

Such a combination would address a practical need. In professional fields like engineering, teachers may need to support students in comparing not only ethics cases, but legal ones. In order to understand—or model—the ethical analysis of

---

87. Ilya M. Goldin et al., *Introducing PETE: Computer Support for Teaching Ethics,* PROC. EIGHTH INT'L CONF. ON ARTIFICIAL INTELLIGENCE & L. 94 (2001).

88. HARRIS ET AL., *supra* note 2, at 59–64.

89. *See* ASHLEY, *supra* note 76, at 36–38; Kevin D. Ashley, *Designing Electronic Casebooks that Talk Back: The CATO Program,* 40 JURIMETRICS J. 275 (2000) (showing that improvements of the CATO model have been used to engage law students in simulated courtroom arguments); Vincent Aleven, Teaching Case-Based Argumentation Through a Model and Examples (1997) (unpublished Ph.D. dissertation, University of Pittsburgh) (showing that improvements of the model have been used to teach law students basic skills of arguing with cases) (on file with Learning Research and Development Center), *available at* http://www-2.cs.cmu.edu/~aleven/dissertation.html (last visited Feb. 14, 2004).

90. Carson Strong, *Justification in Ethics, in* MORAL THEORY AND MORAL JUDGMENTS IN MEDICAL ETHICS 193–211 (Baruch A. Brody ed., 1988).

these scenarios, students may need to understand the legal analysis and vice versa. The interplay of ethical and legal considerations is evident in Case 97-5. In not disclosing the risks to public safety, the engineer is said to rely on the fact that he entered into a confidentiality agreement with his municipal employer. From a legal viewpoint, the agreement might be unenforceable as against public policy. Under the relevant standard,[91] this would depend on, among other things, the parties' justified expectations and the likelihood that a refusal to enforce the agreement would further the public policy. Presumably, provisions of the NSPE codes like II-1-c are evidence of the public policy and condition the justified expectations of the parties to an engineering consultant contract.[92] The magnitude and likelihood of risk to public safety affect the legal assessment of the likelihood that the policy will be furthered just as they affect the ethical assessment of the need for disclosure. The intertwining of ethical and legal issues in case 97-5 appears to be very common in the NSPE BER cases and is discussed at length in the Harris text.[93]

Developing SIROCCO's practical contributions to normative information retrieval and instruction will require much additional work. In the meantime, its techniques for representing citation information and explanation have enabled it to make a more theoretical contribution. It demonstrates the dialectic between principles and cases and the practical effect of operationalization information in fleshing out the meaning of abstract normative principles as they are applied, cited, and explained.

91. *See, e.g.*, RESTATEMENT (SECOND) OF CONTRACTS §178 (1979) (grounds for when a term is unenforceable on grounds of public policy).

92. Criton A. Constantinides, *Professional Ethics Codes in Court: Redefining the Social Contract Between the Public and the Professions*, 25 GA. L. REV. 1327, 1366 (1991).

93. HARRIS ET AL., *supra* note 2, at 163–66, 169.

FIGURES

Figure 1. Three Sample NSPE Code Provisions

II-1-a. Engineers shall at all times recognize that their primary obligation is to protect the safety, health, property, and welfare of the public. If their professional judgment is overruled under circumstances where the safety, health, property or welfare of the public are endangered, they shall notify their employer or client and such other authority as may be appropriate.

II-1-c. Engineers shall not reveal facts, data, or information obtained in a professional capacity without the prior consent of the client or employer except as authorized or required by law or this Code.

III-4. Engineers shall not disclose confidential information concerning the business affairs or technical processes of any present or former client or employer without his consent.

Figure 2. Facts of Case 97-5[94]

Engineer A, a principal in ABC Engineering, an environmental engineering firm, submits qualifications and a proposal to a local municipality to be considered as the consultant for the research and analysis of a former dump site which is being considered for reclamation as a wetland. The dump has been closed for many years after being used for several decades for commercial waste disposal, possibly without any regulation or control. In a meeting with Engineer A, the municipality indicates the possibility that there could be hazardous and toxic wastes encountered in the dump. Upon being awarded the contract, Engineer A is informed by the city that, as part of the contract, a confidentiality clause must be signed which precludes Engineer A from disclosing any results or information concerning the project without the city's written permission. Engineer A signs the contract and the clause.

Preliminary research by Engineer A confirms that the dumpsite is not closed according to the hazardous and solid waste regulations of the state. Tests of the surface soils on the site are inconclusive but reveal a possibility that very high contaminant levels of hazardous and toxic waste could, over time, become exposed at the surface, due to erosion of the cover, and even washed into a river that flows immediately adjacent to the site. The city is considering plans to build a children's park, recreation and picnic area, bike/jogging trail, and parkway near the reclaimed areas, and the river is used for drinking water intake for cities on the other side of the river and downstream. Upon receiving the initial data, the city terminates the contract, saying that the development will be moved to another site, citing the political ramifications of revealing the findings and the economics of having to clean up the property as its reasons for not continuing. Engineer A responds that the city has a responsibility to the public to proceed to remediation, even if the development is moved elsewhere, but the city refuses and reminds Engineer A of its confidentiality clause and the legal consequences of going public with the confidential information. Engineer A decides not to inform the appropriate authorities.

94. Engineering Ethics Transcription Exercise, Case 97-5, *at* http://www.pitt.edu/~bmclaren/ethics/caseframes/97-5.html (last visited Feb. 25, 2004).

Figure 3. Case 97-5 BER Analysis[95]

Questions and Conclusions:

1. Was it ethical for Engineer A not to inform the appropriate regulatory agencies of the engineer's findings and the potential dangers to the public health and the environment? No.

2. Did Engineer A behave ethically in signing the confidentiality clause restricting him from revealing information concerning dangers to the public health and the environment, after being informed by the city that there was a possibility that the site could contain hazardous and toxic wastes? No.

Discussion:

The responsibility of engineers for the protection of the public health and safety is generally considered the most fundamental ethical principal [sic] related to the practice of engineering. . . . However, the view . . . is not universally shared within and outside of the engineering profession. Among the reasons cited by dissenters is the fact that engineers are generally employees or are retained by clients and that their most basic ethical obligation is to their employer or their client and not to the public.

The Board has considered several cases involving the protection of the public health and safety and also the duties of engineers in connection with hazardous waste material.[96] For example, BER Case No. 92-6 . . . . In considering whether it was ethical for Engineer B merely to inform the client of the presence of the drums and suggest that they be removed, and whether Engineer B had an ethical obligation to take further action, the Board noted that the extent to which an engineer has an obligation to hold paramount the public health and welfare in the performance of professional duties (See Code I.1) overlaps the duty of engineers not to disclose confidential information concerning the business affairs, etc. of clients (See Code III.4)[97]. . . . The Board noted that Engineer B's responsibility under the facts was to bring the matter of the drums possibly containing hazardous material to the attention of the client with a recommendation that the material be analyzed. To do less would be unethical. If analysis demonstrates that the material is indeed hazardous, the client would have the obligation of disposing of the material in accordance with applicable federal, state, and local laws.[98]

In an earlier case, BER Case No. 89-7, an engineer was retained to investigate the structural integrity of a 60-year-old, occupied apartment building, which his client was planning to sell. Under the terms of the agreement with the client, the structural report written by the engineer was to remain confidential. In addition, the client made it clear to the engineer that the building was being sold "as is," and the client was not planning to take any remedial action to repair or renovate any system within the building. . . . [D]uring the course of providing services, the client confided in the engineer that the building contained deficiencies in the electrical and mechanical systems, which violated applicable codes and standards. While the engineer was not an electrical or mechanical engineer, he did realize that those deficiencies could cause injury to the occupants of the building and so informed the client. In his report, the engineer made a brief mention of his conversation with the client concerning the deficiencies; however, in view of the terms of the

---

95. *Id.*

96. An example of Operationalization Technique 8, Grouping cases. *See supra* Part II.C and *infra* fig. 7.

97. An example of Operationalization Technique 5, Grouping code principles. *See supra* Part II.C and *infra* fig. 7.

98. An example of Operationalization Technique 9, Reuse of an application; also an example of Operationalization Technique 2, Hypothesizing facts that affect how the code applies. *See supra* Part II.C and *infra* fig. 7.

agreement, the engineer did not report the safety violations to any third parties. In determining that it was unethical for the engineer not to report the safety violations to appropriate public authorities, the Board, citing cases decided earlier, noted that the engineer "did not force the issue, but instead went along without dissent or comment." . . . The Board concluded that the engineer had an obligation to go further, particularly because the NSPE Code uses the term "paramount" to describe the engineer's obligation to protect the public safety, health, and welfare.

In BER Case No. 90-5, the Board reaffirmed the basic principle articulated in BER Case No. 89-7. . . .

The case presently before the Board is similar to each of these earlier cases to some degree, and most probably closest to the situation faced by Engineer A in Case No. 89-7. Although the facts are somewhat different because Case No. 89-7 involved a building containing obvious fire code violations, which had an immediate impact on the building's residents, the Board is convinced that the reasoning in Case No. 89-6 is applicable to this case.[99] Despite a written agreement not to disclose confidential information, Engineer A is bound by the NSPE Code of Ethics and has a paramount duty in matters involving the public health and safety to notify the employer or client, and such other authority as may be appropriate, where the engineer's professional judgment is overruled.

Under the facts, there is ample reason for Engineer A to conclude that a serious public health danger could occur if the project is permitted to proceed as scheduled without a remediation of the hazardous material on the site. . . . Engineer A cannot remain a party to a "conspiracy of silence" against the public health and safety, but instead must identify the appropriate regulatory officials and come forth to explain his professional findings and recommendations.

With regard to Engineer A's actions in signing a confidentiality agreement, while such agreements are relatively common and are usually consistent with Code III.4, the Board is deeply troubled with the fact that Engineer A agreed to sign the agreement knowing that there was a possibility that his professional services would encounter hazardous and toxic material.[100] While the Board does not believe that a confidentiality agreement per se is inappropriate, a confidentiality agreement that "ties the hands" of an engineer to report dangers to the public health and safety is a clear violation of the NSPE Code of Ethics.[101]

---

99. An example of Operationalization Technique 6, Instantiate case as precedent. *See supra* Part II.C and *infra* fig. 7.

100. An example of Operationalization Technique 1, Instantiating code principles. *See supra* Part II.C and *infra* fig. 7.

101. An example of Operationalization Technique 4, Resolution of conflicting code principles. *See supra* Part II.C and *infra* fig. 7.

Figure 4. Fact Chronology for Case 97-5

Key: The Fact Chronology represents the facts of Case 97-5, Figure 2, in a manner that SIROCCO can process. It includes ten facts and their temporal relations. It also indicates that facts 5 and 10 are the basis of the ethics questions put to the BER. English translations are in brackets.

No.	Fact	Time Qualifier
1.	(**owns-the-company** Engineer-A ABC-Engineering) [Engineer A owns ABC Engineering Company.]	pre-existing-fact
2.	(**submits-a-proposal-to-for** Engineer-A Municipality-M Dumpsite-Reclamation) [Engineer A submits a proposal to Municipality M concerning a dumpsite reclamation.]	(After-the-start-of (1))
3.	(**informs-that** Municipality-M Engineer-A ((**may-be-hazardous-to-safety** Dumpsite-Reclamation))) [Municipality M informs Engineer A that the dump-site may be hazardous to safety.]	(After-the-conclusion-of (2))
4.	(**hires-the-services-of-for** Municipality-M Engineer-A ((**provides-engineering-advice-to-regarding** Engineer-A Municipality-M Dumpsite-Reclamation))) [Municipality M hires Engineer A to provide engineering advice regarding the dumpsite reclamation.]	(After-the-conclusion-of (3))
5.	(**signs-the-agreement-with** Engineer-A ((**does-not-disseminate-to** Engineer-A ((**may-be-hazardous-to-safety** Dumpsite-Reclamation)) Public)) Municipality-M) [Engineer A signs an agreement with Municipality M that A will not disseminate information to the public that the dumpsite reclamation may be hazardous to safety.]	(After-the-conclusion-of (4)) (*Questioned-Fact 2*)
6.	(**discovers-that** Engineer-A ((**is-a-safety-hazard** Dumpsite-Reclamation)) [Engineer A discovers that the dumpsite reclamation is a safety hazard.]	(After-the-conclusion-of (5))
7.	(**terminates-the-services-of** Municipality-M Engineer-A) [Municipality M terminates the services of Engineer A.]	(Ends (4))

No.	Fact	Time Qualifier
8.	(**asks-for** Engineer-A Municipality-M ((**should-be-informed-about-the-hazard-or-potential-hazard** Public)) [Engineer A asks Municipality M to inform the public about the safety hazard.]	(After-the-conclusion-of (7))
9.	(**disagrees-with-regarding** Municipality-M Engineer-A ((**should-be-informed-about-the-hazard-or-potential-hazard** Public)) [Municipality M disagrees with Engineer A about informing the public about the safety hazard.]	(After-the-conclusion-of (8))
10.	(**does-not-inform-that** Engineer-A Public ((**is-a-safety-hazard** Dumpsite-Reclamation)) [Engineer A does not inform the public that the dumpsite reclamation is a safety hazard.]	(after-the-conclusion-of (9)) (*Questioned-Fact 1*)

Figure 5. Representing Citation Information re Code Principles

Key: This table summarizes information in the BER analysis of Case 97-5, Figure 3, concerning code principles the Board cites. It indicates whether the Board said that the principle supports its conclusion, whether the principle was violated, the particular facts most critical for that conclusion, and the textual basis. The information (except the textual basis) is represented in a manner that SIROCCO can process.

Code	Support Conclusion?	Code Violated?	Facts	Textual Basis
I-1 Hold paramount safety, health and welfare of public.	Supports	Violated	6, **10**	Engineer is involved in a professional situation in which a safety, health or welfare issue is at stake. Engineer's action does not hold paramount the safety, health, and welfare of the public.
II-1 Engineers shall hold paramount safety, health and welfare of public.	Supports	Violated	6, **10**	Same as above.
II-1-a If engineers' judgment is overruled under circumstances that endanger life or property, they shall notify their employer or client and such other authority as may be appropriate.	Supports	Violated	6, 8, 9, **10**	Engineer's judgment is overruled in a particular professional circumstance. Overruling the Engineer's judgment may lead to the endangerment of the safety, health, property or welfare of the public. Engineer does not notify their employer, client, or other appropriate authority.
II-1-c Engineers shall not reveal facts, data or information without prior consent of client or employer except as authorized or required by law or this Code.	Conflicts	Not-violated	4, 3, 6, **10**	Engineer has a client. Engineer obtains confidential facts, data, or information through work for the client. Engineer does not reveal confidential facts, data, or information to unauthorized parties.

Code	Support Conclusion?	Code Violated?	Facts	Textual Basis
III-4 Engineers shall not disclose, without consent, confidential information concerning business affairs or technical processes of any present or former client or employer, or public body on which they serve.	Supports[a]	Not-violated	3, 4, 6, **10**	Engineer obtains confidential information concerning the business affairs or technical processes of a present client. Engineer does not disclose the confidential information.
II-1-b Engineers shall approve only those engineering documents which are in conformity with applicable standards.	Supports	Violated	**3, 5**	Engineer is in a situation in which he must approve or disapprove an engineering document or documents. Engineer approves a document that is not safe for public health, property, or welfare.

a. Although SIROCCO's representation for Case 97-5 indicates that III-4 supports the BER's conclusion, the case enterer appears to have made a mistake. III-4 conflicts with that conclusion, just as II-1-c does.

Figure 6. Representing Significance of Past Cases

Key: This table summarizes information in the BER analysis of Case 97-5, Figure 3, concerning past cases the Board cites. It indicates whether the Board says the case is analogous to or distinguishable from the problem, the particular facts most critical for that conclusion, and the textual justification. The information (except the textual justification) is represented in a manner that SIROCCO can process.

Case	Analogous or Distinguishable	Facts	Textual Justification
89-6	Analogous-precedent	5, 6, **10**	There is a conflict of interest between an engineer's professional responsibilities and his duty to the public. The engineer's duty to the public is deemed to override his professional responsibilities.
89-7	Analogous-precedent	**5**, 6, **10**	An engineer found evidence of a safety hazard but did not inform the proper authorities. In both the cited and instant cases, there was an agreement to keep the engineer's findings confidential.
90-5	Analogous-precedent	**5**, 6, **10**	An engineer found evidence of a safety hazard but did not inform the proper authorities. The engineer did not disclose his findings due to a confidentiality agreement.
92-6	Distinguishing-precedent	**5**, 6, **10**	An engineer found evidence of a safety hazard but did not inform the proper authorities. Distinguished because in the cited case there was no agreement of confidentiality regarding the engineer's findings.

Figure 7. Operationalization Techniques

1. Instantiate code principles by linking them to clusters of questioned and critical facts.
2. Hypothesize facts that affect how code principles apply.
3. Revise a code principle over time in light of cases.
4. Resolve conflicting code principles in specific cases.
5. Group code principles in specific cases.
6. Instantiate cases as precedents by linking them to clusters of questioned and critical facts, and by analogizing or distinguishing them.
7. Apply, define or elaborate an issue or principle as raised in past case.
8. Group cases.
9. Reuse specific applications of any of the above techniques from previous analyses.

Figure 8. SIROCCO's Analysis of Case 97-5, Question 1

SIROCCO has the following suggestions for evaluating "97-5-1: Signing a Confidentiality Agreement - Duty to Disclose Danger to the Public Health [Ques. #1]"

Possibly Relevant Codes:

III-4: Do not Disclose Confidential Information Without Consent
II-1-C: Do not Reveal Confidential Information Without Consent
I-4: Act as a Faithful Agent or Trustee
III-2-B: Do not Complete or Sign Documents that are not Safe for Public
II-1-A: Primary Obligation is to Protect Public (Notify Authority if Judgment is Overruled)
I-1: Safety, Health, and Welfare of Public is Paramount
II-1-E: Report Alleged Code Violations

Possibly Relevant Pre-Dated Cases:

93-3-1: Appropriate Notification and Review of Another Engineer's Work [Ques. #1]
76-4-1: Public Welfare - Knowledge of Information Damaging to Client's Interest
96-8-1: Peer Review - Confidentiality Agreements
89-7-1: Duty To Report Safety Violations

Additional Suggestions:

1. There may be an important distinction between the possibly relevant case 93-3-1 and the current case, 97-5-1, with respect to III-4 ("Do not Disclose Confidential Information Without Consent") and II-1-C ("Do not Reveal Confidential Information Without Consent"). In case 93-3-1 the board concluded that Engineer B was unethical. The current case is similar in many respects to 93-3-1, but the cases differ on these two facts: 93-3-1: "ENGINEER-B INFORMS-THAT ENGINEER-A ((HIRES-THE-SERVICES-OF-FOR FRANCHISER ENGINEER-B ((REVIEWS-AND-ANALYZES ENGINEER-B DESIGN-WORK))))" 97-5-1: "ENGINEER-A DOES-NOT-INFORM-THAT PUBLIC ((IS-A-SAFETY-HAZARD DUMPSITE-RECLAMATION))"

2. There may be an important distinction between the possibly relevant case 96-8-1 and the current case, 97-5-1, with respect to II-1-E ("Report Alleged Code Violations"). In case 96-8-1 the board concluded that Engineer A was ethical. The current case is similar in many respects to 96-8-1, but the cases differ on these two facts: 96-8-1: "ENGINEER-A INFORMS-THAT AUTHORITIES ((FAILS-STANDARDS-AND-MAY-BE-HAZARDOUS-TO-SAFETY DESIGN-WORK))" 97-5-1: "ENGINEER-A DOES-NOT-INFORM-THAT PUBLIC ((IS-A-SAFETY-HAZARD DUMPSITE-RECLAMATION))"

3. There may be an important distinction between the possibly relevant case 93-3-1 and the current case, 97-5-1, with respect to I-4 ("Act as a Faithful Agent or Trustee"). In case 93-3-1 the board concluded that Engineer B was unethical. The current case is similar in many respects to 93-3-1, but the cases differ on these two facts: 93-3-1: "ENGINEER-B INFORMS-THAT ENGINEER-A ((HIRES-THE-SERVICES-OF-FOR FRANCHISER ENGINEER-B ((REVIEWS-AND-ANALYZES ENGINEER-B DESIGN-WORK))))" 97-5-1: "ENGINEER-A DOES-NOT-INFORM-THAT PUBLIC ((IS-A-SAFETY-HAZARD DUMPSITE-RECLAMATION))"

4. The codes I-1 ("Safety, Health, and Welfare of Public is Paramount") and II-1-A ("Primary Obligation is to Protect Public (Notify Authority if Judgment is Overruled)") may override codes III-4 ("Do not Disclose Confidential Information Without Consent") and I-4 ("Act as a Faithful Agent or Trustee") in this case. See case 76-4-1 for an example of this type of code conflict and resolution.

5. The code II-1-E ("Report Alleged Code Violations") may override code III-4 ("Do not Disclose Confidential Information Without Consent") in this case. See case 96-8-1 for an example of this type of code conflict and resolution.

6. The cases 82-2-1, 85-4-1, and 87-2-1 were cited by 89-7-1 to highlight or elaborate a general principle or common scenario. Since 89-7-1 has been suggested as possibly relevant to the present case, its cited cases may also be relevant. Check whether the general scenario of the cited cases is relevant to the present case: "Engineer has a client" "Engineer obtains confidential facts, data, or information through work for the client."

7. The case 67-10-1 was cited by 76-4-1 to highlight or elaborate a general principle or common scenario. Since 76-4-1 has been suggested as possibly relevant to the present case, its cited case may also be relevant. Check whether the general scenario of the cited case is relevant to the present case: "Engineer is involved in a professional situation in which the public welfare is at stake."

Figure 9. Comparison of Citations by NSPE Board vs. SIROCCO

Citations by NSPE Board	Citations by SIROCCO
Code I-1	**Code I-1**
No match	Code I-4 (Arguably relevant to Case 97-5)
Code II-1	No match (Equivalent to Code I-1)
Code II-1-a	**Code II-1-a**
Code II-1-b	No match (Equivalent to Code III-2-b)
Code II-1-c	**Code II-1-c**
No match	Code II-1-e (Arguably relevant to Case 97-5)
No match	Code III-2-b (Equivalent to Code II-1-b)
Code III-4	**Code III-4**
No match	Code III-4-a (Related to Codes II-1-c, III-4)
No match	Case 76-4-1 (Arguably relevant to Case 97-5)
Case 89-6	No match (Apparent typo in Board opinion)
Case 89-7-1	**Case 89-7-1**
Case 90-5	No match
Case 92-6	No match
No match	Case 93-3-1 (Arguably relevant to Case 97-5 but distinguishable)
No match	Case 96-8-1 (Arguably relevant to Case 97-5)

Figure 10. Normalized TF/IDF Weights for Text of Fact Description of Case 97-5

citi 0.329	children 0.087	remedi 0.068
dump 0.321	trail 0.087	...
sit 0.237	*contamin 0.087*	analysi 0.050
hazard 0.215	ramif 0.087	commerci 0.049
confidenti 0.199	intak 0.087	*disclos 0.049*
waste 0.196	inconclus 0.087	appropri 0.049
river 0.192	preclud 0.087	becom 0.049
clause 0.181	bik 0.087	environment 0.048
toxic 0.174	wash 0.087	due 0.047
surfac 0.160	decad 0.087	data 0.047
reveal 0.132	wetland 0.087	cover 0.046
inform 0.121	drink 0.087	control 0.046
mov 0.121	contract 0.085	decid 0.046
clos 0.112	regul 0.083	*public* 0.045
research 0.112	...	test 0.045
being 0.105	sign 0.073	develop 0.045
possibl 0.103	...	refus 0.044
consid 0.097	encount 0.071	award 0.044
used 0.089	solid 0.071	accord 0.042
municip 0.087	...	...
could 0.087	flow 0.071	consult 0.016
expos 0.087	soil 0.068	project 0.016
picnic 0.087		stat 0.012
ero 0.087		firm 0.011
reclaim 0.087		engin 0.002

Figure 11. SIROCCO's Architecture[102]

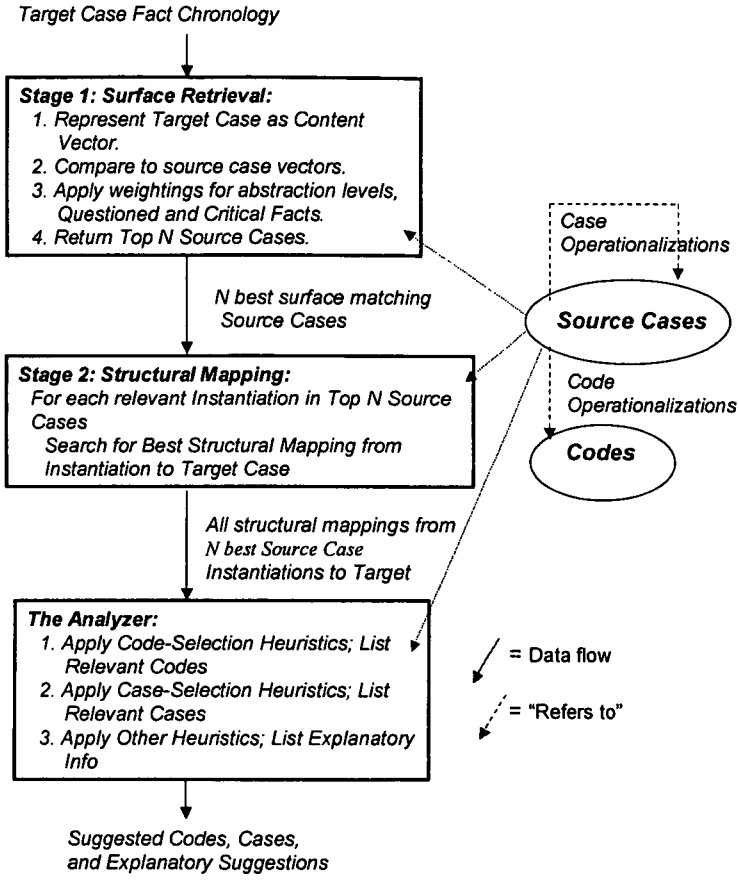

102. McLaren, *Ph.D. dissertation, supra* note 1, at 64 fig. 3-1.

Figure 12. Content Vector for Case 97-5

owns-the-company (1); submits-a-proposal-to-for (1); informs-that (1); may-be-hazardous-to-safety (2); hires-the-services-of-for (1);	provides-engineering-advice-to-regarding (1); signs-the-agreement-with (1); does-not-disseminate-to (1); discovers-that (1); is-a-safety-hazard (2);	terminates-the-services-of (1); asks-for (1); should-be-informed-about-the-hazard-or-potential-hazard (2); disagrees-with-regarding (1); does-not-inform-that (1)

Figure 13. SIROCCO's Detailed Explanation of Some Recommendations

Fact Situation 97-5: Signing a Confidentiality Agreement—Duty to Disclose Danger to the Public Health (Case 97-5-1 is Question # 1 of Fact Situation 97-5)

Possibly Relevant Codes:

II-1-A: Primary Obligation is to Protect Public (Notify Authority if Judgment is Overruled).

Reasons:

- Cited by 3 of the 6 best surface matching cases.
- 66.7% match and ques. fact match (Source and Target: FACT-PRIMITIVE) in case 76-4-1.
- > 50.0% match in multiple cases: 88-6-1 76-4-1
- Grouped with code I-1 in case 76-4-1; Good match to ques. facts.

Mapping:

*To Case 76-4-1 (Corresponding steps of Case 97-5-1 indented; * = questioned fact)*

ENGINEER-DOE DISCOVERS-THAT ((FAILS-STANDARDS-AND-MAY-BE-HAZARDOUS-TO-SAFETY DISCHARGE))
>ENGINEER-A DISCOVERS-THAT ((IS-A-SAFETY-HAZARD DUMPSITE-RECLAMATION))
*ENGINEER-DOE DOES-NOT-INFORM-THAT CONTROL-AUTHORITY ((FAILS-STANDARDS-AND-MAY-BE-HAZARDOUS-TO-SAFETY DISCHARGE))
*>ENGINEER-A DOES-NOT-INFORM-THAT PUBLIC ((IS-A-SAFETY-HAZARD DUMPSITE-RECLAMATION))

Possibly Relevant Pre-Dated Cases:

96-8-1: Peer Review - Confidentiality Agreements

Reasons:

- 60.0% match and ques. fact match (Source and Target: FACT-PRIMITIVE) in code III-4.
- 100.0% match to 2 critical facts in code II-1-E.
- 100.0% match and ques. fact match (Source and Target: FACT-PRIMITIVE) in code II-1-E.
- 100.0% match to 2 critical facts (citation to 76-4-1).
- 100.0% match and ques. fact match (Source and Target: FACT-PRIMITIVE) (citation to 76-4-1).

Mappings:

*To Code Inst. III-4 (Corresponding steps of Case 97-5-1 indented; * = questioned fact)*

ENGINEER-A DISCOVERS-THAT ((FAILS-STANDARDS-AND-MAY-BE-HAZARDOUS-TO-SAFETY DESIGN-WORK))

>ENGINEER-A DISCOVERS-THAT ((IS-A-SAFETY-HAZARD DUMPSITE-RECLAMATION))

PEER-REVIEW-PROGRAM INSTRUCTS-TO ENGINEER-A ((DOES-NOT-INFORM-THAT ENGINEER-A ANYONE ((REVIEWS-AND-ANALYZES ENGINEER-A DESIGN-WORK))))

>ENGINEER-A ASKS-FOR MUNICIPALITY-M ((SHOULD-BE-INFORMED-ABOUT-THE-HAZARD-OR-POTENTIAL-HAZARD PUBLIC))

*ENGINEER-A INFORMS-THAT AUTHORITIES ((FAILS-STANDARDS-AND-MAY-BE-HAZARDOUS-TO-SAFETY DESIGN-WORK))

*>ENGINEER-A DOES-NOT-INFORM-THAT PUBLIC ((IS-A-SAFETY-HAZARD DUMPSITE-RECLAMATION))

*To Code Inst. II-1-E (Corresponding steps of Case 97-5-1 indented; * = questioned fact)*

ENGINEER-A DISCOVERS-THAT ((FAILS-STANDARDS-AND-MAY-BE-HAZARDOUS-TO-SAFETY DESIGN-WORK))

>ENGINEER-A DISCOVERS-THAT ((IS-A-SAFETY-HAZARD DUMPSITE-RECLAMATION))

*ENGINEER-A INFORMS-THAT AUTHORITIES ((FAILS-STANDARDS-AND-MAY-BE-HAZARDOUS-TO-SAFETY DESIGN-WORK))

*>ENGINEER-A DOES-NOT-INFORM-THAT PUBLIC ((IS-A-SAFETY-HAZARD DUMPSITE-RECLAMATION))

89-7-1: Duty To Report Safety Violations

Reason:

- 65.0% match and ques. fact match (Source and Target: FACT-PRIMITIVE) (citation of 97-13-2).

Figure 14. Information Retrieval Measures for Citation Overlap

Key: The right oval represents the NSPE BER's citations for a case.The left oval represents the citations returned for that case by Method X (e.g., SIROCCO, MG, Non-Op SIROCCO, etc.) The intersection contains the BER citations that Method X succeeds in returning. These are True Positives (TP). False Positives (FP) are citations Method X returns that are not among the BER's citations. False Negatives (FN) are BER citations that Method X fails to return. True Negatives (TN) would be citations that the BER did not make and that Method X did not return. Recall (R) is the percentage of the relevant citations that Method X retrieved. Precision (P) is the percentage of the retrieved citations that were relevant. The F-measure is a heuristic combination of precision and recall; it quantifies the degree of citation overlap between Method X and the Board. Beta (β) represents the relative importance of precision versus recall. β was set at 1.0, assigning P and R equal importance.

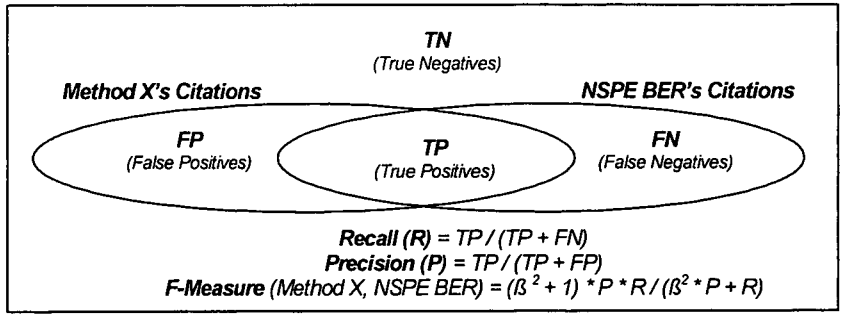

Figure 15. Mean F-Measures for Six Methods

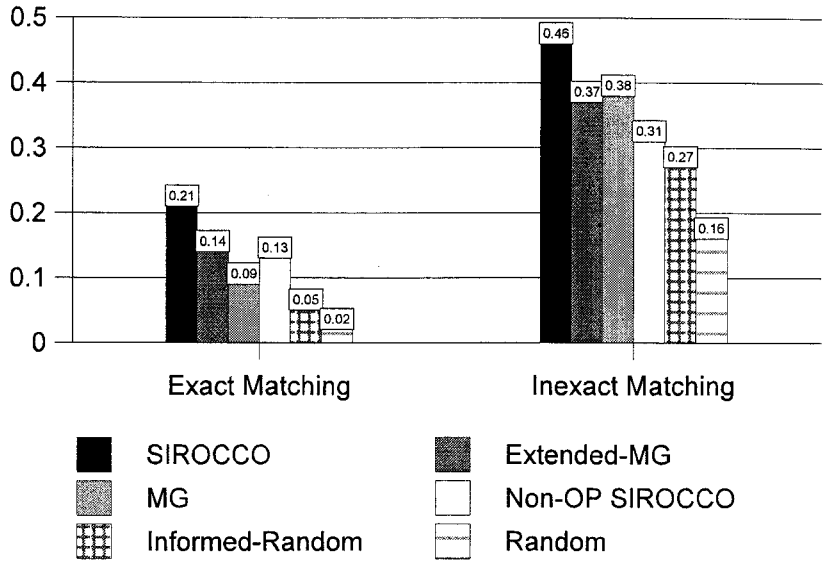

Figure 16. Mean F-Measures of SIROCCO vs. Extended MG after Adjustment for Extra Citations