

## **A Framework for Self-Explaining Legal Documents**

L. Karl Branting  
*Department of Computer Science*  
*University of Wyoming*  
*Box 3682*  
*Laramie, WY 82071*  
*karl@index.uwyo.edu*

James C. Lester  
*Department of Computer Science*  
*North Carolina State University*  
*Box 8206*  
*Raleigh, NC 27606*  
*lester@adm.csc.ncsu.edu*

### **Abstract**

Legal document drafting is an essential professional skill for attorneys and judges. To maintain stylistic and substantive consistency and decrease drafting time, new documents are often created by modifying previous documents. This paper proposes a framework for document reuse based on an explicit representation of the illocutionary and rhetorical structure underlying documents. Explicit representation of this structure facilitates (1) interpretation of previous documents by enabling them to “explain themselves,” (2) construction of documents by enabling document drafters to issue goal-based specifications and rapidly retrieve documents with similar intentional structure, and (3) maintenance of multi-generation documents. The applicability of this framework to a representative class of judicial orders—jurisdictional show-cause orders—is demonstrated.

### **1 Introduction**

Legal problem solving subsumes a number of distinct tasks, including analysis of the legal consequences of actual or hypothetical sequences of actions, argumentation, advising clients, planning transactions, and drafting legal documents. Legal document drafting is an essential professional skill for attorneys and judges. In the U.S., a significant portion of attorneys’ workload consists of drafting documents intended to precisely stipulate legal relationships, such as wills, contracts, and leases, and persuasive documents arising from litigation, such as pleadings, motions, and briefs. Judges routinely draft performative documents, such as orders and decisions.

Complex legal documents are typically created by modifying previous documents. Adaptation and reuse of previous documents is an almost universal practice in U.S. law firms. Document reuse is beneficial because it reduces drafting time and promotes stylistic and substantive consistency. However, document reuse requires access to the original intentions underlying the document, which may not be readily apparent from the document’s surface text. For example, when both parties to a contract agree that modifications should be made,

the assumptions behind the original contract must be reconstructed to determine the precise textual changes required.

To address the problem of recovering the intentions underlying documents to facilitate their reuse, we propose a *self-explaining documents* framework. We say that a document containing a given discourse is “self-explaining” if it contains an explicit representation of the illocutionary and rhetorical<sup>1</sup> structure underlying the discourse. Because self-explaining documents record intentional knowledge, they offer significant potential for the interpretation and maintenance of complex, multi-generation documents. In particular, they can explain why a particular clause was included, suggest how an existing document should be modified to apply to slightly different circumstances, and present arguments for the pros and cons of alternative clauses.

This paper presents a framework for the use of justification structures for document reuse. Section 2 describes a representative class of legal documents—appellate jurisdictional show cause orders—and illustrates how the self-explaining document approach would assist drafting in this application. Section 3 proposes a dual justification structure that combines illocutionary and rhetorical structures to represent document intent. Section 4 illustrates the envisioned behavior of the system when applied to the domain of show-cause order drafting. Section 5 describes future implementation plans for this work, and Section 6 outlines related work.

## **2 Document Reuse in Judicial Drafting**

In the Anglo-American legal system, the most prominent judicial documents are appellate decisions. Appellate opinions typically contain a summary of the facts of the case, identification of the issues of law raised in arguments by counsel for each of the parties, pronouncement of the legal propositions supported by the controlling authorities, and declaration of a decision that resolves the issues by applying the legal propositions to the facts of the case. The complexity and individuality of appellate opinions makes automated assistance for such documents unfeasible. However, courts produce a number of other more routine documents having considerable stylistic and substantive consistency, including various types of orders issued in response to motions or *sua sponte*. Jurisdictional show-cause orders are typical of such orders.

Jurisdictional show-cause orders are generally issued during jurisdictional screening, a process of determining whether the requirements for an appeal have been satisfied. Jurisdictional screening is typically performed at the earliest possible stage of an appeal to permit cases with jurisdictional defects to be recognized as soon as possible. This minimizes unnecessary consumption of limited judicial resources.

In this paper, discussion of appellate jurisdictional screening will focus on the Colorado Court of Appeals, where one of the authors, Karl Branting, worked for several years as a staff attorney. The Colorado Court of Appeals receives on average over 100 new cases per month. Screening these appeals is too complex for clerical personnel, but must instead be performed by a staff attorney. The staff attorney examines the case file to determine whether the subject-matter, finality, and timeliness requirements for appellate jurisdiction have been met. If there

---

<sup>1</sup> An illocutionary operator is a speech act such as informing, requesting, warning, or promising. A rhetorical operator is a discourse or coherence relation, such as exemplification, generalization, sequence, or elaboration. See Allen [1987] for a more detailed discussion.

*A Framework for Self-Explaining Legal Documents*

appears to be a jurisdictional defect, the staff attorney drafts a show-cause order that sets forth the apparent defect and orders the appellant to rebut the defect within a fixed time period or face dismissal of the appeal.

Colorado Court of Appeals No. 90CA0274	Order Tr. Ct. No. 79DR221
In re the Marriage of	
SUSAN W. KIRKPATRICK	Appellant,
and	
JOHN B. KNEZOVICH	Appellee.
To: Susan W. Kirkpatrick, appellant, and her attorney, Susan M. Lach	
From the notice of appeal filed by appellant and the register of actions submitted by the clerk of the district court, it appears that this appeal is from an order denying a motion to change venue or for an order declining jurisdiction as an inconvenient forum under Section 14-13-108, C.R.S. It further appears that this is not a final judgment because it does not end "the particular action in which it is entered, leaving nothing further for the court pronouncing it to do in order to completely determine the rights of the parties involved in the proceeding." <i>Harding Glass Co. v. Jones</i> , 640 P.2d 1123 (Colo. 1982); <i>D.H. v. People</i> , 192 Colo. 542, 561 P.2d 337 (1965).	
IT IS THEREFORE ORDERED that the appellant shall show cause, if any she has, in writing on or before April 17, 1990, why this appeal should not be partially dismissed without prejudice for failure to file a final appealable order.	
BY THE COURT	
Dated: April 3, 1990	
Copies to: Counsel of Record	

**Figure 1: A typical show-cause order.**

Figure 1 shows a typical show cause order. This order identifies an apparent defect—lack of finality—and orders the appellant to show cause with 14 days why the appeal should not therefore be dismissed. Show-cause orders typify legal documents that are produced in relatively high volume (several hundred per year), are complex enough to require drafting by an attorney, but have a sufficient degree of stylistic and substantive consistency to facilitate reuse.

*Jurix '96: Branting and Lester*

Currently, staff attorneys at the Colorado Court of Appeals draft show cause orders manually. One approach is for the staff attorney to reuse only those “boilerplate” text blocks that are common to all show-cause orders (e.g., “From the notice of appeal filed by appellant...”). Alternatively, a staff attorney may refer to a collection of previous show-cause orders. The attorney can search this collection for a previous order that involved defects similar to those in the current case. The reusable language from the previous case may then be transcribed into the current case, or the staff attorney may photocopy the previous order, cross out the portions inapplicable to the new case, and write in portions specific to the new case.

Manual drafting of show-cause orders has several clear disadvantages. First, the process is very time-consuming and laborious. The first approach, which reuses only the most general boilerplate language, entails repeated replication of drafting effort and creates a high likelihood of inconsistent language. The second approach, reuse of similar orders, depends on the staff attorney’s ability to find and appropriately modify previous orders. This in turn depends on the attorney’s ability to understand the relevant similarities and differences between the goals that must be achieved by the current show-cause order and the intentions underlying previous orders.

The difficulty of finding and adapting previous documents is exacerbated by the frequent personnel changes. For example, at the Colorado Court of Appeals few staff attorneys are willing to do jurisdictional screening for more than six months, and many do screening for as little as three months. As a result, jurisdictional screening is typically performed by attorneys familiar with only the most recent show-cause orders. Staff attorneys may therefore have difficulty reconstructing the intent behind particular clauses, impeding interpretation and modification.

One approach to providing automated assistance in the drafting of show-cause orders would be to design a set of templates for various show-cause orders, e.g., Word Perfect macros. However, the wide variety of possible jurisdictional defects and the even wider variety of factual situations that can give rise to jurisdictional defects make it unlikely that a adequate collection of macros could be devised. Moreover, even if a sufficient set of macros could be devised, these macros would present the user with an overwhelming number of choices, making selection of the most appropriate macro unlikely.

We argue that the goals of accuracy, efficiency, and stylistic consistency are best served by an approach to drafting routine legal documents, such as show-cause orders, based on retrieval and modification of previous documents. The task of drafting by document reuse can be summarized as follows:

**Given:**

- A set of goals to be accomplished by the document to be drafted.
- A library of existing documents.

**Do:**

- Retrieval. Find the existing document(s) (or combination of document components) that best satisfy the current goals.
- Analysis. Display the goals achieved by the retrieved document or document component
- Comparison. Display the differences, if any, between, the goals achieved by the retrieved document(s) and the current goals
- Adaptation. Remove the portions of text whose only purpose is to satisfy goals that aren’t present in the current situation (excision), and add text to satisfy any of the current goals not satisfied by the retrieved text (augmentation).

### *A Framework for Self-Explaining Legal Documents*

A system for legal drafting by document reuse must therefore enable users to pose at a minimum three classes of queries:

- Goal Exemplification: What are some examples of archival documents or document segments that achieve a given goal or set of goals? Goal exemplification is necessary for the retrieval stage of drafting.
- Goal Identification: Why was a given segment included in a document? Goal identification is necessary for the analysis stage of drafting.
- Goal Comparison: To what segments of archival documents is a given document segment most similar? What are the intentional similarities and differences between two documents (at the most specific abstraction level at which they differ)? Goal comparison is necessary for the comparison and adaptation stages of drafting.

In addition, the system must provide the functionalities required for systematically substituting case-specific data from the current case for case-specific data in the previous case, e.g., names of parties, dates, etc., and compile the result into a new document.

### **3 Representational Requirements for Self-Explaining Legal Documents**

Legal documents can serve a variety of illocutionary goals, including eliciting information, persuading, memorializing events such as reciprocal communications, or accomplishing performative goals, such as creating or revoking legal relationships. For example, the goals of show-cause orders include to make-findings about relevant facts such as the apparent judgment and filing dates, and make-conclusions concerning the existence of a jurisdictional defect, and inform the appellant of the apparent defect, time limit for response, and sanction if no adequate response is filed. Self-explaining legal documents therefore require a vocabulary of illocutionary goals sufficiently rich to permit accurate document retrieval, analysis, and adaptation.

However, the illocutionary goal structure of a document is not per se sufficient to completely determine the document's surface text. In general, the illocutionary goal structure leaves unspecified rhetorical features such as (1) the order of the textual elements that satisfy various illocutionary subgoals, and (2) textual elements and stylistic constraints imposed by the particular genre of the text. Accordingly, self-explaining documents must include the documents' rhetorical structure as well as their illocutionary structure.

The minimum representational requirements for self-explaining documents therefore include the following:

- A taxonomy of illocutionary goals sufficiently expressive to permit retrieval of documents, comparison of documents, and explanation of document components. The necessary granularity of the leaf nodes of this taxonomy depends on the requirements of the particular document genre and the pragmatics of the user's application.
- A taxonomy of rhetorical goals.
- A representation of templates at a level of granularity corresponding to the leaf nodes of the illocutionary goals. The substitutable elements of text templates should be tagged with a data-type so that entire documents, or document components, can be viewed either as uninstantiated templates or as fully instantiated texts.
- A set of link annotations, e.g., annotations providing the legal authority under which a given legal goal is satisfied by a given set of performative subgoals, and annotations explaining why a given rhetorical goal is satisfied by a particular set of subgoals in a given document genre.

Our model of the intentional structure of jurisdictional show-cause orders includes five illocutionary operators: **establish**, **inform**, **find**, **rule**, and **order**. These operators are necessary because of the nature of performative judicial documents (i.e., judicial documents intended to create or alter legal relationships). Typically, such documents must make findings of relevant facts, rulings of the applicable law, and order some change in legal status. In addition, there may be intermediate reasoning steps that must be established, and the recipient of the document may need to be informed of facts which are neither findings, rulings, or orders, such as the applicable legal authority.

To illustrate, a simplified representation of the illocutionary structure of the show-cause order shown in Figure 1 is set forth in Figure 2. The root illocutionary goal is to establish (**est**) the prerequisites for dismissal. Achieving this goal requires the operators displayed as children of the root: set forth the preamble (**inform**); establish a jurisdictional defect; and order a response. The goal to establish a jurisdictional defect, in turn, is achieved by establishing that the order being appealed is non-final. This goal, in turn, is established by **finding** that the order is a motion to change venue, **ruling** that a motion to change venue is not a final order, and **informing** the appellant of the authority that supports the ruling. The subgoals of the **order** requirement are **informing** the appellant of the date a response is due, the sanction for failing to respond, and the defect. Each leaf illocutionary goal is connected to a text segment intended to achieve that goal.

A simplified representation of the rhetorical structure of the show-cause order is shown in Figure 3. Unlike the illocutionary structure, the rhetorical structure is closely connected to the surface text of the document. The top-level node of the rhetorical structure, show-cause order, indicates the genre of the document. The children of the root node consist of a set of elements sufficient for a document of that genre. Together, the annotated illocutionary and rhetorical goal structures constitute the *justification structure* of a document.

#### **4 Example: Use of Self-Explaining Show-Cause Orders**

This section illustrates through an example the manner in which the justification structure described in the previous section could be used to assist in drafting new show-cause orders. Suppose that a relatively inexperienced staff attorney, Jones, is screening an appeal in the case “*In re the Marriage of Herbert W. Smythe and Catherine Smythe*” and discovers an apparent jurisdictional defect. Under the Colorado Appellate Rule 4(a), an appeal to the Court of Appeals is timely only if a notice of appeal is filed within 45 days of (1) the date of entry of judgment, if the parties are present when the judgment is announced, or (2) the date of mailing of the notice, if notice of the judgment is transmitted by mail. However, if the appellant files a motion for post-trial relief the due date for the notice of appeal is extended until the motion is ruled upon or expires.

*A Framework for Self-Explaining Legal Documents*





Figure 2: A simplified representation of the illocutionary structure of the show-cause order in Kirkpatrick.

*Jurix '96: Branting and Lester*



**Figure 3: The rhetorical structure of the show-cause order in Kirkpatrick.**

Suppose that the Smythe file shows that notice of the judgment was mailed on September 22, 1995, but the notice of appeal was not filed until November 7, 1995 (more than 45 days later). Jones therefore wishes to draft a show-cause order having to do with an untimely notice of appeal. Representation of the justification structure of existing documents could assist Jones in the following steps:

**Initial Document Retrieval**

- User: Initiates the document planning session by selecting “Construct new document” from pull-down menus.
- Document Planner: Presents user with dialog form listing various types of orders and notices. Among his choices is Show-Cause Order.
- User: Selects Show-Cause Order.
- Document Planner: Presents user taxonomic hierarchies of illocutionary and rhetorical goals and substantive legal areas.

### *A Framework for Self-Explaining Legal Documents*

- User: Selects marriage from the taxonomy of substantive legal areas.
- Document Planner: Asks user if he would like to see a typical show cause order pertaining to marriage.
- User: Answers affirmatively.
- Document Planner: Retrieval system indexes into a partition of the case library containing show-cause documents pertaining to marriage. It retrieves the prototypical show-cause order for marriage, Kirkpatrick
- User: Studies Kirkpatrick, considers how it relates to his current task, and highlights a region of the document that seems different than the one he would like to have in the current: "It further appears that this is not a final judgment because it does not end the 'particular action in which it is entered, leaving nothing further for the court pronouncing it to do in order to completely determine the rights of the parties involved in the proceeding.' "
- Document Planner: Asks user if he would like to know why this section was included.
- User: Answers affirmatively.
- Document Planner: Uses the illocutionary structure to generate an explanation of the communicative goals behind the document. First it locates the substructure of the illocutionary structure that covers the region:

**find** (non-final-order-type(motion-to-change-venue))

It then ascends the supergoals to find

**est** (jurisdictional-defect(no-finality))

From this it generates the explanation "The text you have highlighted was included to establish that there was a lack of a final appealable order." It then poses the question, "Is this true in the case for which you are drafting an order?"

- User: The user selects "No."
- Document Planner: The planner inspects its operator library to determine other reasons for establishing jurisdictional defects. It then finds all of the operators that achieve the goal of establishing jurisdictional defects and asks the user about the applicability to the current case. After asking about two others to which the user responds negatively, it encounters

**est** (jurisdictional-defect(failure-to-file-timely-noa))

<

**inform** (defective-order),

**est**(untimely-NOA(failure-to-file-timely-NOA)).

and asks the user, "Is the jurisdictional defect in your case a failure to file a timely notice of appeal?"

- User: Responds affirmatively.
- Document Planner: Asks the user if he would like it to search for show-cause orders in which the jurisdictional defect is a failure to file a timely notice of appeal.
- User: Responds affirmatively.
- Document Planner: The planner again locates the sub-partition of the document library pertaining to documents whose illocutionary structure contains an instantiated Establish operator with the following pattern

**est** (jurisdictional-defect(failure-to-file-timely-NOA))

It locates Cohen, whose illocutionary structure has

**est**jurisdictional-defect(failure-to-file-timely-NOA))

*Jurix '96: Branting and Lester*

<  
**est**(untimely-NOA(failure-to-file-timely-NOA))

It then asks the user if he would like to use Cohen as a model.

- User: Responds affirmatively

**Interactive Document Compilation**

- Document Planner: The planner then begins to construct an illocutionary structure for the user's show-cause order by traversing the illocutionary structure of Cohen, performing a structure mapping, and querying the user for information. As it traverses the tree, it encounters

**est** (time-limit-begins(cohen))

<

**find**(parties-present-at-judgment(cohen))

and so it asks if both parties were present at the Smythe judgment.

- User: Responds in the negative.

**Interactive Document Adaptation and Constraint Propagation**

- Document Planner: As a result of the interaction, the planner has established that the Cohen case is not a precise analogue of the Smythe case. It can now provide the user with a good deal of assistance by exploiting its knowledge of the illocutionary structures of other documents. Its first step is to determine alternative means of achieving the goal of establishing when the time limit begins. It searches the operator library and discovers

**est**(commencement-of-time-limit)

**find**(mailed-noticed-of-judgment).

so it asks the user if, in the Smythe case, the judgment notice was mailed to the parties.

- User: Responds affirmatively.
- Document Planner: Next it asks if the user if he would like to see an excerpt of a show cause order dealing with a mailed judgment.
- User: Responds affirmatively.
- Document Planner: Indexes into the document library and examines the illocutionary structures. It is unable to locate a show cause order pertaining to marriage in which the judgment notice was mailed. It then searches other types of show cause orders and it finds a document with an illocutionary structure including the instantiated operator

**est** (time-limit-begins(canada))

<

**find**(mailed-noticed-of-judgment(canada))

It then climbs the illocutionary structure of Canada up two levels to obtain the context in which the decision was made. Here it finds the instantiated goal

**est**(duration-of-time-limit("45 days", "C.A.R. 4(a)"))

Next, it locates all text in the Canada document that was generated by this goal and all its progeny goals. It then spawns a document viewer, and displays this text to the user:

## *A Framework for Self-Explaining Legal Documents*

However, it appears that the trial court's order granting summary judgment was entered on February 9, 1987 and mailed to counsel of record on February 10, 1987, and that no C.R.C.P. motion for post-trial relief was filed.

It then asks the user if Smythe presents a similar situation.

- User: Responds affirmatively.
- Document Planner: It now structure-maps the Canada illocutionary structure for the "mailing" segment over to the Smythe case by interacting with the user. It first asks when the judgment was mailed.
- User: Enters "September 22, 1995".
- Document Planner: This triggers a constraint propagation that calculates a new value for the ?DUE-DATE, which is 45 days after the mailing date. The planner explains the calculation to the user, using the authority clause (?AUTHORITY = "C.A.R. 4(a)") to point the user to the justification. The planner then requests the user to confirm the due date.
- User: Confirms date.
- Document Planner: Encounters the goal in Canada  
**find** (no-post-trial-relief-motion(canada))  
and asks if in Smythe, there was a post trial relief motion.
- User: The user responds affirmatively.
- Document Planner: The planner then integrates the sub-structure of the illocutionary structure that it analogized from Canada into the super-structure that it is in the process of building from Cohen.

### **Completing Document Compilation**

The planner passes the newly constructed illocutionary structure for the Smythe case to the document compiler, which performs the following actions. It first checks if any internal constraints have been violated (which may happen as a result of adaptive interactions). In this case, no violations have occurred. It next passes the illocutionary structure to the document generator, which takes all **inform** operators and realizes them in text. Finally, it locates the rhetorical structure associated with documents of the given class (in this case, Show-Cause orders). The rhetorical structure is used to (1) order the text resulting from the **inform** operators and (2) format the document. The resulting document is then displayed in the document studio and saved in the archive.

## **5 Implementation Plans**

We are currently applying the self-explaining document approach to Colorado Court of Appeals show-cause orders. As argued above, show-cause orders appear to typify legal documents that are produced in relatively high volume (several hundred per year), are complex enough to require drafting by an attorney, but have sufficient stylistic and substantive consistency to facilitate reuse.

We have developed an initial implementation of illocutionary rules for a set of representative show-cause orders and are currently engaged in developing rhetorical rules for these documents. To supply a uniform representation for all of the knowledge structures, all illocutionary structures, rhetorical structures, and planning operators are to be implemented in a

### *Jurix '96: Branting and Lester*

unification-based constraint system [Elhadad, 1991]. Use of this single representation permits all adaptation procedures (document compilation and consistency checking) to be performed with a single mechanism: constraint propagation. Although this brings with it an initial development overhead, we expect to witness significant software engineering benefits in the form of rapid extensions to the domain theory. Once the system for our first class of regulatory documents is complete, we will run extensive empirical evaluations to measure the gains in efficiency and correctness.

## **6 Related Work**

Our approach to self-explaining documents draws on four different lines of research: discourse structure analysis; the theory of argumentation; explanation generation; and automated document drafting. The primary focus of research in discourse structure has been accounting for the coherence of expository or other communicative text through hierarchical structures of rhetorical and other discourse relations, e.g. Grosz et al. [1986] and Hobbs [1979]. The formalization of inter-sentential discourse relations is a key requirement for the development of self-explaining documents.

The most directly relevant portion of research in discourse structure is speech act theory. Initiated by J.L. Austin, who was primarily concerned with explicit performatives [Austin, 1962], speech act theory addresses the illocutionary content of discourse, that is, the goals that a speaker intends to accomplish through that discourse [Grice, 1975; Searle, 1969].

The theory of argumentation addresses texts intended to persuade, establish, or prove. For example, Toulmin [1958] analyzed argumentative texts in terms of the concepts of warrant, ground, conclusion, backing, and qualification. This model has been widely applied to the analysis [Marshall, 1989; Zeleznikow et al., 1995] and creation [Bench-Capon et al., 1995] of legal documents. Argument structure, like other forms of illocutionary goal structure but unlike rhetorical structure, does not directly address the "surface" form of texts. This line of research is particularly relevant to the analysis of the illocutionary structure of persuasive or dispositive documents such as legal briefs and judicial decisions [Branting, 1993a].

The explanation community has extensively studied the process of planning and realizing text given a set of discourse specifications. Over the past decade, their work research on discourse planning [McKeown, 1985; Paris, 1988; Hovey, 1990; Hovey, 1993; Cawsey, 1992; Suthers, 1993; Moore, 1995; Mittal, 1993; Lester et al., 1996] has produced a variety of techniques for determining the content and organization of many genres of text. Perhaps because of the necessity of coping with the myriad underlying rhetorical, illocutionary, and argument structures in discourse generation, this work has yielded a variety of mechanisms for determining the content and organization of multi-sentential text, a key capability of self-explaining documents.

Automated document drafting research is the fourth relevant research area. Two important areas of automated document drafting research are automated legal drafting and automated report generation. A large number of automated legal drafting systems have been developed in recent years, but most involve creation of text templates that are then instantiated to create particular documents [Lauritsen, 1992]. Some progress has been made in exploiting explicit representations of the relationship between generic documents and document instances and of constraints among document components [Daskalopulu et al., 1995]. However, there is a growing recognition in the Law and AI community that a declarative representation of the knowledge underlying the selection and configuration of

## *A Framework for Self-Explaining Legal Documents*

textual elements is essential for the development of tools that embody the expertise of legal drafting experts [Gordon, 1989; Lauritsen, 1993].

The automated report generation community has addressed another form of text production from an underlying domain structure: the derivation of technical documentation from program traces generated during software development or use [McKeown et al., 1995; Korelsky et al., 1993]

### **7 Summary**

Document drafting is a critical task for attorneys and judges. To cope with the complexities of constructing and maintaining complicated, multi-generation documents, we have proposed the self-explaining document framework. We are currently instantiating the framework in a system for drafting appellate jurisdictional show-cause orders. It is our hypothesis that, by equipping documents with a justification structure that includes both illocutionary and rhetorical components, self-explaining document systems can assist document designers with constructing, maintaining, and interpreting complex documents.

Our initial focus has been on legal documents in the common law system. However, there is nothing about the self-explaining document approach that is limited to the common-law system or, indeed, legal documents. This approach is applicable to any documents that accomplish well-defined illocutionary goals and are subject to consistent rhetorical constraints.

### **Acknowledgments**

This research is supported in part by grants from the National Center for Automated Information Research and by NSF Faculty Early Career Development Grant IRI-9502152.

### **References**

- J. Allen. *Natural Language Understanding*. Benjamin/Cummings, Menlo Park, CA, 1987.
- J. Austin. *How to do things with words*. Oxford U. Press, New York, 1962.
- T. Bench-Capon and G. Staniford. PLAID - proactive legal assistance. In *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pages 81--88, 1995.
- L. K. Branting. An issue-oriented approach to judicial document assembly. In *Proceedings of the Fourth International Conference on Artificial Intelligence and Law*, pages 228--235, Amsterdam, The Netherlands, June 15--18, 1993. ACM Press.
- L. K. Branting. A reduction-graph model of ratio decidendi. In *Proceedings of the Fourth International Conference on Artificial Intelligence and Law*, pages 40--49, Amsterdam, The Netherlands, June 15--18, 1993. ACM Press.
- A. Cawsey. *Explanation and Interaction: The Computer Generation of Explanatory Dialogues*. MIT Press, 1992.

*Jurix '96: Branting and Lester*

- A. Daskalopulu and M. Sergot. A constraint-driven system for contract assembly. In *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pages 62--70, 1995.
- M. Elhadad. FUF: The universal unifier user manual version 5.0. Technical Report CUCS-038-91, Department of Computer Science, Columbia University, 1991.
- T. Gordon. A theory construction approach to legal document assembly. In *Pre-Proceedings of the Third International Conference on Logic, Informatics, and Law*, pages 485--498, Florence, 1989.
- H. Grice. Logic and conversation. In P. Cole and J. Morgan, editors, *Syntax and Semantics 2: Speech Acts*, pages 41--58. Academic Press, New York, NY, 1975.
- B. Grosz and C. Sidner. Attention, intention, and the structure of discourse. *Computational Linguistics*, 12(3), 1986.
- J. Hobbs. Coherence and co-reference. *Cognitive Science*, 3(1):67--82, 1979.
- E. H. Hovy. Pragmatics and natural language generation. *Artificial Intelligence*, 43:153--197, 1990.
- E. H. Hovy. Automated discourse generation using discourse structure relations. *Artificial Intelligence*, 63:341--385, 1993.
- T. Korelsky, D. McCullough, and O. Rambow. Knowledge requirements for the automatic generation of project management reports. In *Proceedings of the Eighth Knowledge-Engineering Conference*. IEEE Computer Society Press, September 20--23 1993.
- M. Lauritsen. Technology report: Building legal practice systems with today's commercial authoring tools. *Law and Artificial Intelligence*, 1(1), 1992.
- M. Lauritsen. Knowing documents. In *Fourth International Conference on Artificial Intelligence and Law*, pages 185--191, Amsterdam, 1993. ACM Press.
- J. C. Lester and B. W. Porter. Scaling up explanation generation: Large-scale knowledge bases and empirical studies. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, Portland, Oregon, to appear 1996.
- C. Marshall. Representing the structure of a legal argument. In *Proceedings of the Second International Conference on Artificial Intelligence and Law*, pages 121--127, Vancouver, B.C., June 13-16 1989.
- K. R. McKeown. *Text Generation: Using Discourse Strategies and Focus Constraints to Generate Natural Language Text*. Cambridge University Press, 1985.

*A Framework for Self-Explaining Legal Documents*

V. Mittal. *Generating Natural Language Descriptions with Integrated Text and Examples*. PhD thesis, University of Southern California, September 1993.

J. D. Moore. *Participating in Explanatory Dialogues*. MIT Press, 1995.

K. McKeown, J. Robin, and K. Kukick. Generating concise natural language summaries. *Information Processing and Management*, 1995. Special Issue on Summarization.

C. Paris. Tailoring object descriptions to a user's level of expertise. *Computational Linguistics*, 14(3):64--78, September 1988.

J. Searle. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, Cambridge, 1969.

D. D. Suthers. *An Analysis of Explanation and Its Implications for the Design of Explanation Planners*. PhD thesis, University of Massachusetts, February 1993.

S. E. Toulmin. *The Uses of Argument*. Cambridge University Press, 1958.

J. Zeleznikow and A. Stranieri. The split-up system: Integrating neural networks and rule-based reasoning in the legal domain. *In Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pages 185--194, 1995.