# Narrative-Driven Case Elicitation

Karl Branting[1,*], Sarah McLeod[2]

[1] The MITRE Corporation, McLean, VA, USA
[2] The MITRE Corporation, Seattle, WA, USA

**Abstract**
This paper proposes an approach for narrative-driven case elicitation that uses schemas induced from corpora of legal case facts to distinguish relevant from irrelevant utterances and to identify facts that could distinguish between competing hypotheses. This approach to narrative-driven case elicitation builds on recent research in narrative schema induction.

**Keywords**
case elicitation, narrative schema, law, computational linguistics, machine-learning, human-computer interface

## 1. Introduction

> Law is full of stories, whether these are stories that are told in the courtroom as lawyers try to weave together compelling and competing versions of an event, in the legislative histories that subtend a body of statutes, or in stories about the origins and acceptance of legal systems [1].

The facts of legal cases are more than mere collections of events. Instead, case facts are *narratives* that have settings, characters with goals and motives, and events triggered by the characters' actions. Just as not every set of facts is a story, not every story is legally significant. It is the role of a legal counsel to identify the goals that the client hopes to achieve through a legal process and then to elicit a coherent narrative that is relevant to one or more possible legal remedies that could achieve those goals.

Lay (non-attorney) clients typically have little understanding of what facts are relevant to a possible legal remedy. Accordingly, attorneys must help clients express the facts that are relevant to possible legal remedies and steer clients away from facts irrelevant to those remedies. The legal remedy that appears most likely to achieve client's goals may change during the interview, requiring a reframing of the relevance of the facts that have been previously expressed and redirecting the elicitation toward the facts relevant to the new remedies.

This paper sets forth an approach to narrative-driven case fact elicitation and situates that approach within a broader architecture for induction and use of legal narratives schemas. The next section provides a background on the role of narrative understanding in providing legal assistance, and recent research in narrative schema induction is reviewed in Section 3. Section 4 presents an algorithm that uses case schemas for narrative-driven case elicitation. Section 5 describes an architecture that incorporates the narrative-driven case elicitation into a framework that includes narrative schema induction. Section 6 sets forth the text processing steps shared by both the schema induction and case elicitation components. Six new corpora for narrative schema induction and case elicitation are described in Section 7, and future steps are proposed in Section 8.

## 2. The Role of Narrative in Law

Law can be viewed as a framework of rules under which legal arguments often consist of alternative narratives that lead to opposite consequences. Empirical evidence has shown that jurors often decide cases based on which of two competing narratives imposes the highest degree of coherence on the evidence presented at trial [2] [3]. Unsurprisingly, the outcomes of trials often depend on the relative story-telling ability of attorneys and witnesses [4] [5].

In view of the importance of coherent legally relevant narratives to success in litigation, narrative elicitation is widely recognized as a vital legal skill. One study of client interviews revealed the importance of permitting a client ample time to speak, during which the attorney acknowledges the information received and expresses interest but otherwise "refrains from interrupting" except to direct the narrative away from "precarious" legal grounds and toward "possible remedies" "until the client has gone on for long enough to establish the problem at hand" [6]. In general, attorneys try to elicit "the causal and temporal connections that contribute to giving the events contextual meanings … with the aim of defining

*Workshop on Artificial Intelligence for Access to Justice (AI4AJ 2023), June 19, 2023, Braga, Portugal.*
*Corresponding author.
lbranting@mitre.org (K. Branting); smcleod@mitre.org (S. McLeod)
0000-0002-9362-495X (K. Branting)

'Who has done what, how, when, why and where?'" [7]. Clients' narratives are often "redefined to be a legally relevant narrative" by counsel, a process that can sometimes diminish emotionally salient background information [8]. An "account of the situation the client faces in light of the law" is sometimes term a *case theory* [9]. Case theories "unite possible client narratives with possible legal theories" [9].

In the context of protective order interviews, it has been shown that the interviewer "acts to reshape, if not repair, the narratives of domestic violence victims, so that they conform to the requirements of an affidavit that must be submitted to a judge if a protective order is to be issued" [10]. This is necessary because protective order applicants on their own often "tell their accounts of violence in a manner that differers quite markedly" from the "formulaic ... structure and ... thematic content" required for affidavits [10].

In summary, the achievement of a client's goals depends on the degree to which the clients story can be presented, whether in a written petition or to a judge or jury, in a compelling, coherent, and legally relevant way. This depends in turn on the ability of the client's (human or automated) counsel to elicit case facts in a manner that has these characteristics.

## 3. Narrative Schema Induction

A story has been described as "a character-based and descriptive telling of a character's efforts, over time, to overcome obstacles and achieve a goal" [5]. Early computational studies of narrative were motivated by efforts in cultural anthropology to formalize commonalities among folk stories [11]. Story grammars were an effort to characterize narratives in a rigorous way [12]. However, the brittleness of story grammars eventually led to their almost complete abandonment [13] in favor of "scripts," stereotypical sequences of events that create expectations and fill in missing details to assist in story understanding [14].

Unlike story grammars, scripts ordinarily lack the hierarchical and recursive structure needed to account for some important properties of stories (e.g., in a picaresque novel or a protective order application there can be a variable number of episodes, and episodes can have subepisodes). Moreover, both story grammars and scripts initially were entirely manually constructed, which made them unscalable.

However, a series of research advances have made it increasingly feasible to induce "narrative schemas" (i.e., "scripts" in Schank/Abelson terminology) from examples. The seminal work by Chambers and Jurafsky [15] defined "narrative chains" (later termed "narrative schemas" [16]) as "partially ordered set[s] of narrative events that share a common actor." The most common model for narrative schemas are Markov chain models that assign a probability to an event based on cooccurring (e.g., prior) events. Such models can be used to detect distinguish expected (high probability) events from unexpected (low probability) events and to predict the most likely events at a given point in an event sequence.

The performance of narrative schemas is typically measured by a narrative cloze test, i.e., accuracy in predicting the next or a missing event [16]. Improving narrative cloze performances has been obtained by using stricter constraints on multi-argument consistency [17], topic-specific training sets [18], and alternative language models, e.g., Hidden-Markov [19], Log-Bilinear [20], and Association Rule models [21].

The continuing progress in narrative schema induction techniques suggests that this approach will be an increasingly effective computational story model notwithstanding the limitation that narrative schemas are, as mentioned above, only a partial representation of the elements of story. This is because the aspect of performance that narrative schema induction seeks to optimize, the cloze test, can play a central role in the model of legal case elicitation, as described in the next section.

## 4. Using Schemas for Case Elicitation

Section 2 described how effective case elicitation requires helping a client articulate the events giving rise to a legal claim in a manner consistent with known legally meaningful narratives, e.g., in the linear fashion required for protective order affidavit. Section 3 described how narrative schemas can be induced from training sets of stories and then used to distinguish expected, unexpected, and missing events. This section describes an approach to narrative-driven case elicitation that uses the predictive capability of narrative schemas derived from prior case facts to guide interactions with a client.

A high-level view of a process for identifying the information that is most relevant to a litigant's legal goals, in the sense of being the fact that would best discriminate among legally relevant narratives, is set forth in Algorithm 1. The key requirement for the narrative elicitation process is a set of <schema, goal> pairs, (SGs), where each schema is a model capable of evaluating the relative likelihood of a given event sequence, and each goal is a (possibly negated) legal remedy, e.g., a protective order, a child custody order, a finding of employment discrimination, etc. Each schema is induced from a corpus of related case facts for a given area of law using the techniques described in Section 3.[1]

[1] A companion paper to this work describes techniques developing such schemas from representative legal case facts [22].

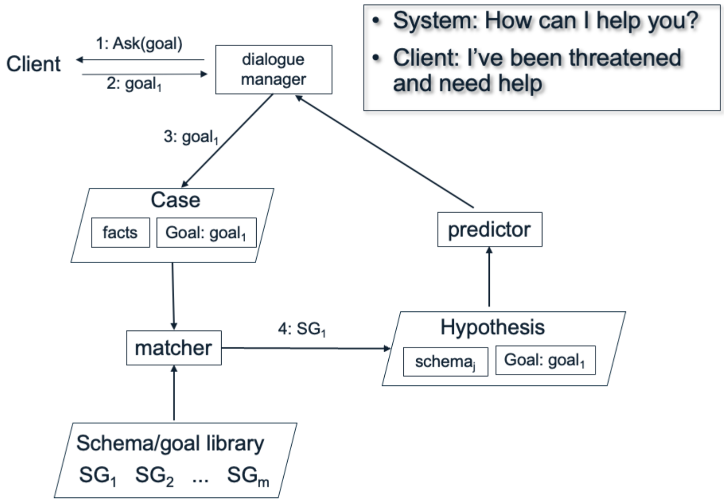


**Figure 1:** First two steps of an elicitation session.

In Algorithm 1, *sgs* is a library of SGs, the *goal* is the objective that the client hopes to achieve through a legal process, the *hypothesis* is the best matching SG, and the *facts* consist of the events and relations elicited thus far from the client.

**Algorithm 1** Legal narrative elicitation

sgs ← {SG*}
hypothesis ← ∅
facts← ∅
goal ← ask(goal)
hypothesis ← BestMatch(facts,goal,sgs)
**while** match(facts,goal,hypothesis)< threshold **do**
  newFact ← ask(getMissing(facts,hypothesis)
  add(facts,newFact)
  hypothesis ← BestMatch(facts,goal,sgs)
**end while**
**return** facts, goal, hypothesis

Algorithm 1 is only a high-level depiction of the actual complex process of case elicitation. The identification of the goal or goals of a client can itself be a complex process requiring mixed-initiative dialogue techniques beyond the scope of this paper. However, but there is an extensive body of research on goal-directed dialog (e.g., [23] [24]), and recent work has addressed the specific task of categorizing a legal aid clients' problems [25].

The function "match(facts,goal,SG)" measures the degree of match between the client's goal and current facts and the SG. An appropriate baseline function for "match" is the probability of the facts under the schema, or 0.0 if the goals don't match (i.e., the schema is irrelevant if the goals don't match, regardless of how well the facts match). The *hypothesis* is the currently best matching SG ∈ sgs. The function "getMissing(facts,hypothesis)" returns the missing fact or event that, if added to the facts, would most increase the probability of those facts. Performance on this task is equivalent to performance on the narrative cloze test [16], so as advances in narrative schema induction improve performance on cloze, the "getMissing" function should improve as well. The *hypothesis* is updated after each new fact has been elicited, reflecting the way that an attorney's assessment of a case may change as more facts are learned.

The process of eliciting additional facts and refining the hypothesis continues until there is no more progress or the match between the hypothesis and the facts and goal exceeds a success threshold. Only at this point are the legal rules for the remedy applied to the case facts. This corresponds to an attorney eliciting the full story from a client before turning to an assessment of the viability of a claim based on that story.

Figure 1 illustrates the initial steps of an elicitation under process shown in Algorithm 1. The dialogue manager starts by asking the client's goal (Step 1). The client replies by stating there was a threat and that the client wishes for help against the threat (Step 2). The *dialogue manager* uses this goal to instantiate a new case with a goal (Step 3) and the start of an event sequence. The *matcher* searches for the SG that best matches the goal and initial event sequence of the new case (i.e., the SG whose remedy matches the goal and whose schema maximizes the probability of the case's fact sequence) and makes that SG the current hypothesis (Step 4).

Figure 2 shows the *predictor* using the schema of the current hypothesis to predict the most probable missing

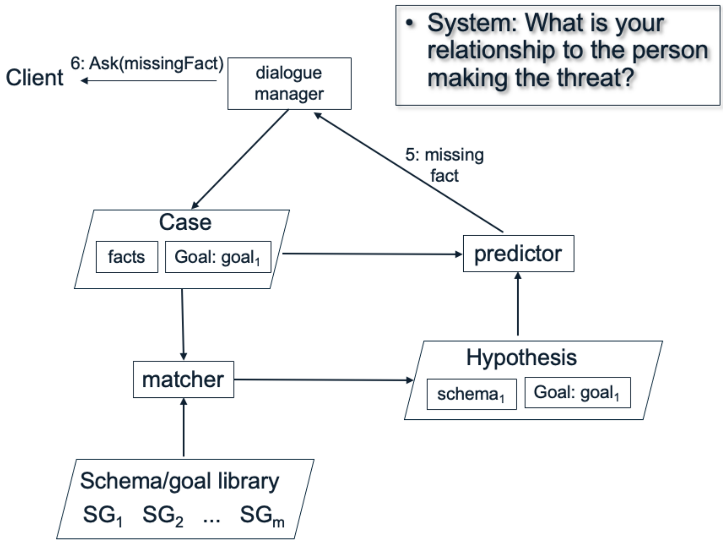


**Figure 2:** Steps 5 and 6 of an elicitation session.

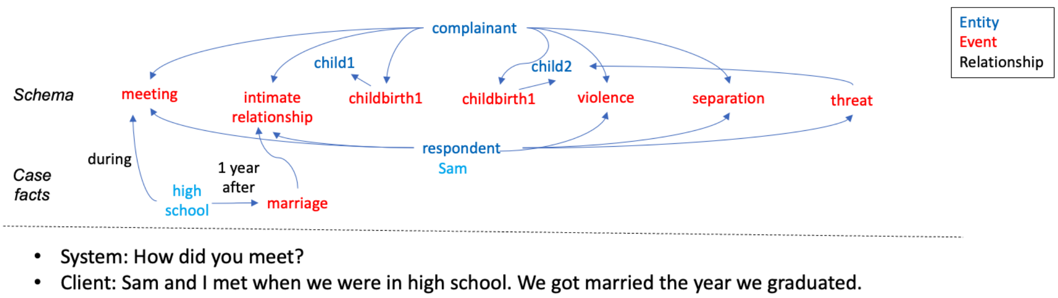


**Figure 3:** An early stage of a protective order case elicitation.

fact in the current case (Step 5) and the dialogue manager converting the fact into an appropriate discourse action (Step 6).

A more detailed (although still simplified) visualization of the initial stages in Algorithm 1 is shown in Figure 3. The sequence of events starting with "meeting," "intimate relationship," etc., represents the most probable event sequence under a typical protective order schema. The "case facts" represent a relational representation of a client's answer to the question, "How did you meet?"

Any practical implementation of Algorithm 1 must ensure a common representation for the schemas and case facts; otherwise, the matching and prediction steps would not be possible. The next section describes an architecture to achieve this common representation.

# 5. RIM: An Architecture for Schema Induction and Use

A system for narrative schema-based fact elicitation depends on two coordinated capabilities: acquiring narrative schemas from examples; and using those schemas to guide interactions with litigants. Figure 4 sets forth an architecture for providing these two capabilities. This architecture is termed *RIM*, short for "Relevant, Irrelevant, and Missing," since the key functionality of the system is identifying these three categories of events.

The left side of Figure 4 details an off-line mechanism for inducing schemas from narrative corpora. The right side of Figure 4 depicts the real time component, including the process described in the previous section of using these schemas to distinguish relevant from irrelevant utterances and to identify facts that could distinguish among legal schemas if confirmed or disconfirmed. The

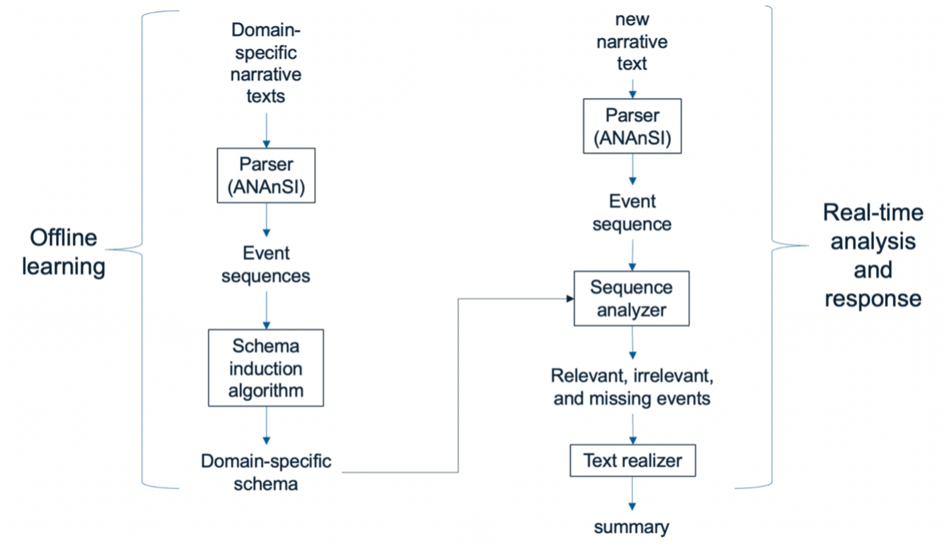


**Figure 4:** The RIM architecture for narrative case elicitation.

*Text Realizer* generates questions to determine whether missing events can be confirmed or disconfirmed. Additional events elicited in this manner can distinguish among partially matching narratives or refine the match to the most similar narrative.

The real-time processing depicted on the right side of Figure 4 depends on the existence of a narrative schema for each area of law for which facts are to be elicited. The process of induction of schemata from event sequences, depicted on the left side of Figure 4, is detailed in [22]. However, both the offline and real-time aspects of depend on conversion of raw text into event sequences, as shown as the second and third steps on both sides of Figure 4.

# 6. Text to Event Sequence Conversion

Offline narrative schema induction and online mixed-initiative case elicitation depend on a shared representation for event sequences.

## 6.1. Parsing

The first step in converting text to event sequences is to parse each sentence into individual events and, for each event, identify the entities that fill the semantic roles of that event. The next step is analyzing the relationships among events by resolving coreferences and determining the discourse relations among the events. Many alternative approaches could be used to perform these two steps; we use ANAnSI (Advanced Narrative Analytics System Infrastructure) [26], a system that integrates the output of the Stanford Core NLP [27] constituency parser and cTakes [28] into a temporal, causal, and intentional graph represented in Neo4j [29].

Figure 5 shows a portion of the graph for the sentence "During my employment, Respondent placed me on a leave of absence and required me to pass a medical exam in order to return to work" showing temporal and intentional links between pairs of events.

## 6.2. Graph Linearization

The resulting graph representation for a collection of one or more sentences is then linearized into an event sequence with arguments and semantic roles standardized in the manner proposed in [17] to three alternatives: agent, patient, and other complement. For example, in the event sequence shown in Figure 6, the pronouns "I" and "me" are normalized to "I".

## 6.3. Lemma Normalization

As discussed below in Section 7, corpora of legal narratives are, in general, many orders of magnitude smaller than the corpora used in previous narrative schema elicitation research, such as the Gigaword corpus. Such small corpora produce sparse transition matrices with little predictive value, e.g., most event pairs in a new (or held out) event sequence will have never been seen before, meaning that there is no frequency data on which to base cloze predictions.

Several normalizations were therefore applied applied

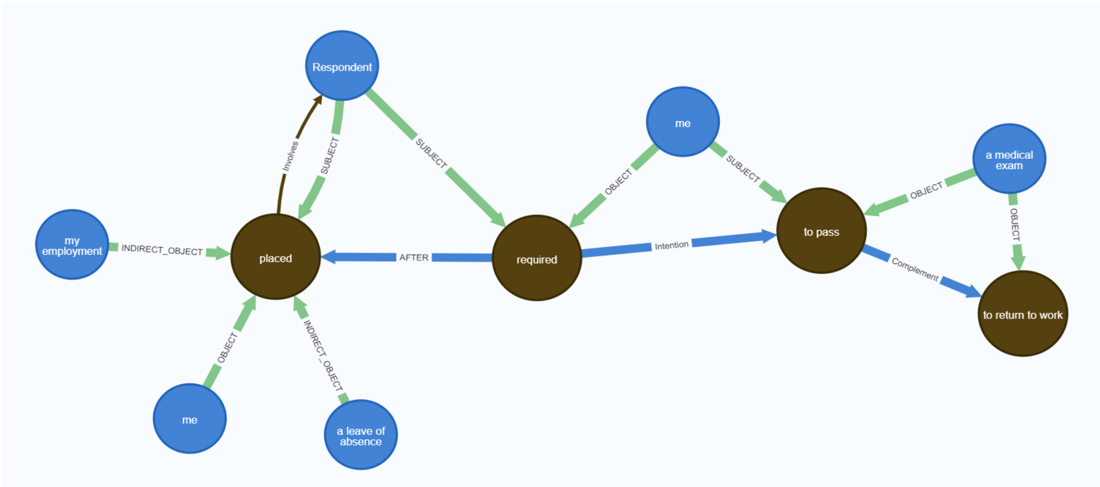


**Figure 5:** A narrative fragment showing the temporal, causal, and intentional graph relationships extracted by ANAnSI.

```
{1134: [1429::hire(Respondent, I, -),
        1427::be engineer(I, -, -),
        1416::employment(I, -, -),
        1424::place(Respondent, I, ['my employment', 'a leave of absence']),
        1423::require(Respondent, I, -),
        1418::pass(I, a medical exam, -),
        1421::return(-, -, ['work']),
        1408::file(I, a Charge of Discrimination, -),
        1405::present(Respondent, I, ['a severance agreement']),
        1413::unlawful(I, -, -),
        1414::be unlawful(included terms, -, -),
        1403::discharge(-, I, -),
        1397::believe(I, -, -),
        1401::discriminate(-, I, ['retaliation', 'engaging', 'violation of Title Vll', 'the Civil Rights Act of 1964']),
        1399::engage(-, -, -),
        1400::amend(-, -, -),
        1392::believe(I, -, -),
        1388::discriminate(-, I, ['I', 'violation of the Americans', 'Disabilities Act of 1990']),
        1394::amend(-, -, -)],
```

**Figure 6:** A linearization of ANAnSI's temporal, causal, and intentional graph.

to reduce vocabulary size to improve matching. The most important and general of these was lemma normalization, which consists of clustering events in semantic embedding space[2] and replacing each event with the most central member of the cluster in which it occurs. For example, Figure 7 show the results of complete-linkage hierarchical clustering of events in our EEOC corpus (described below) with a minimum cosine threshold of 0.75. For example, both 'harass' and 'threaten' are replaced by 'intimidate,' and 'ask,' 'hear,' 'know,' and 'let' are all replaced by 'tell.'

A second normalization that was particularly useful in the employment domains was to replace each occurrence of a form of "to be" that has as an argument the name of an occupation with the event "be OCCUPATION."[3]

Lemma normalization shrinks the vocabulary size of the narrative, increasing transition matrix density and therefore reducing the likelihood that event cooccurrences will never have been observed in the training corpus. This reduction in vocabulary size comes at the cost of reducing the specificity of the event representation.

[2]We used the spaCy large English model, https://spacy.io/models/en.

[3]We used the list of 1,156 occupations, from "accountant" to "zoologist" set forth in https://github.com/johnlsheridan/occupations/blob/master/occupations.csv.

```
{'accuse': ['accuse', 'blame'],
 'advise': ['advise', 'inform'],
 'allow': ['allow', 'permit', 'require'],
 'believe': ['admit', 'believe', 'say', 'thing'],
 'convince': ['convince', 'prove'],
 'decide': ['decide', 'decision'],
 'employment': ['employment', 'wages'],
 'explain': ['explain', 'question'],
 'fear': ['fear', 'scared'],
 'have': ['be', 'have'],
 'help': ['help', 'learn'],
 'insist': ['deny', 'insist', 'refuse'],
 'intimidate': ['harass', 'intimidate', 'threaten'],
 'keep': ['keep', 'leave'],
 'look': ['feel', 'look'],
 'make': ['give', 'make', 'put', 'take'],
 'meet': ['meet', 'meeting'],
 'regard': ['regard', 'respect'],
 'send': ['email', 'send'],
 'shift': ['change', 'shift'],
 'speak': ['speak', 'talk'],
 'suspend': ['revoke', 'suspend'],
 'tell': ['ask', 'hear', 'know', 'let', 'tell'],
 'terminate': ['terminate', 'termination'],
 'time': ['start', 'time'],
 'trouble': ['avoid', 'trouble'],
 'want': ['come', 'do', 'get', 'go', 'want']}
```

**Figure 7:** Lemma normalization by clustering event types in semantic embedding space.

# 7. Corpora

A key challenge for narrative schema-based case elicitations is the difficulty of obtaining significant numbers of narrative texts representative of text produced by litigants. In general, such text contains sensitive personal information that precludes sharing in the form of public corpora. Documents filed in legal or administrative bodies are typically public, so statements of facts in petitions, complaints, and other filings can be a source of narrative texts. However, counsel for litigants often draft the statements in facts of court filings, so the text of such statements seldom contains language used by litigants themselves except in the case of self-represented (*pro se*) litigants, i.e., those who have no attorney to draft their statements in fact. The ideal corpus would consist of statements of fact in pro se litigants' filings, but such filings are difficult to obtain in bulk.

In this research, we obtained one small corpus of texts by pro se litigants together with five other data sets intended to reflect various characteristics of fact statements:

1. EEOC complaints. The complaints were transcribed from handwritten texts in the field titled "The facts supporting the plaintiff's claim of discrimination," in thirty employment discrimination complaints filed in the Northern District of Illinois in 2016. These texts are representative of litigant-generated narrative texts.
2. Multi-LexSum Summaries of Civil Cases. These three hundred sixty four summaries of civil rights lawsuits were created for training and evaluating legal case summarization [30]. The Multi-LexSum text were included to typify procedural histories, a type of narrative required for appeals that court personnel have identified as being challenging for pro se appellants.
3. WIPO cases. The "background facts" of 6,000 decisions by World Intellectual Property Organization in domain name disputes. These fact statements were drafted by the panel deciding the case and are therefore not representative of pro se text. However, the similarity among these fact statements suggests that they could be a benchmark for narrative induction.
4. Board of Veteran Affairs decisions. The "Introduction" section of 1,680 Board of Veterans Appeals (BVA) cases. As with the WIPO cases, these texts are drafted by the judge writing the opinion and are therefore not representative of pro se text but potentially useful as a benchmark for narrative induction.
5. SPOT-HO online housing questions. Two hundred sixty three questions posed to the Suffolk University Law School's Legal Innovation and Technology (LIT) Lab issue spotting service [31].
6. SPOT-WO online employment questions. Two hundred ninety five employment questions posed to the SPOT site.

The size, type, and authors of each of the corpora are summarized in Table 1.

# 8. Summary and Discussion

This paper has proposed an approach to narrative-driven case elicitation that builds on recent research in narrative schema induction. This approach uses schemas induced from corpora of legal case facts to distinguish relevant from irrelevant client utterances and to identify facts that could distinguish among competing hypotheses if confirmed or disconfirmed.

Only the offline portion of the RIM model has been implemented in this project, as described in [22]. A working prototype of the narrative elicitation portion of RIM, which is the focus of this paper, would require assembling a library of scheme-goal pairs for a given area of law from a suitable corpus, such as the BVA or WIPO corpora described above. The matching and prediction functions described in Algorithm 1 are basic capabilities of the schemas described in [22], and as mentioned

| Corpus | Size | Text Type | Author Type |
|---|---|---|---|
| EEOC | 30 | complaints | pro se litigant |
| SPOT-WO | 295 | legal advice requests | lay public |
| SPOT-HO | 263 | legal advice requests | lay public |
| Multi-Lexum | 364 | procedural history | federal judge |
| BVA | 1,680 | background facts | administrative law judge |
| WIPO | 6,000 | background facts | administrative law judge |

**Table 1**
The size, type, and authors of each corpus of narrative texts.

above there is an extensive literature on the goal-directed mixed-initiative dialogue techniques needed for the *dialogue manager*. Thus, there are no significant technical obstacles to implementing a prototype.

The narrative-driven case elicitation approach described here is quite unlike the dominant techniques for automated legal assistance, which are overwhelming organized around form-filling [32] or backchaining through logical representations of legal rules [33] (see generally [34]). The narrative-driven approach is intended to change the focus of client interviewing from the structure of the target legal artifacts (completed petitions and legal rules) to the life experiences that give rise to legal claims and to human interactions between client and attorney. While the narrative-driven model proposed here is a drastic simplification of the actual interview process between clients and attorneys, we hope that it will be a first step toward more faithful model and is feasible to implement with current narrative schema induction and dialogue technology.

## Acknowledgments

The authors express their gratitude to Charlotte Alexander for permitting us to experiment with her collection of EEOC complaints and to Charles Horowitz and Karine Megerdoomian for their assistance with ANAnSI. The MITRE Corporation is a not-for-profit company, chartered in the public interest. This document is approved for Public Release; Distribution Unlimited. Case Number 23-1184. 

## References

[1] A. Tait, L. Norris, Narrative and the origins of law, Law and Humanities 5 (2011) 11–22.
[2] N. Pennington, R. Hastie, A cognitive theory of juror decision making: The story model, Cardozo L. Rev. 13 (1991) 519.
[3] N. Pennington, R.Hastie, Explaining the evidence: Tests of the story model for juror decision making, Journal of Personality and Social Psychology 62 (1992) 189–206. https://doi.org/10.1037/0022-3514.62.2.18.
[4] C. D. Phillips, Reconstructing reality in the courtroom: Justice and judgment in american culture, American Political Science Review 77 (1983) 275–276. doi:10.2307/1956108.
[5] K. D. Chestek, Competing stories: A case study of the role of narrative reasoning in judicial decisions, Legal Comm. & Rhetoric: JAWLD 9 (2012) 99.
[6] S. Kashimura, Hearing clients' talk as lawyers' work: the case of public legal consultation conference, in: M. L. Baudouin Dupret, T. Berard (Eds.), Law at Work: Studies in Legal Ethnomethods, Oxford University Press, Oxford, 2015, p. 87–113.
[7] F. D. Donato, Fact-finding in contexts: framing clients' agentivity within judicial and administrative procedures, file:///Users/lbranting/Downloads/DIDONATO_position_paper_Rotterdam_22.9.15.pdf, last accessed February 5, 2023, 2015.
[8] M. Livingston, A Sad Story in Not a Legal Defense: Defining Legal Issues, Master's thesis, University of Colorado at Boulder Dept. of Communication, 2014.
[9] A. Shalleck, How clients and lawyers construct facts. the stories they tell each other and the stories that guide investigations into the world, in: Dignifying and Undignified Narratives in and of (the) Law, Proceedings of the IVR World Congress 2019 WS Law and Narrative, 2020, pp. 119–144.
[10] S. L. Trinch, S. Berk-Seligson, Narrating in protective order interviews: A source of interactional trouble, Language in Society 31 (2002) 383–418.
[11] A. J. Greimas, C. Porter, Elements of a narrative grammar, Diacritics 7 (1977) 23–40.
[12] D. E. Rumelhart, Notes on a schema for stories, in: D. G. BOBROW, A. COLLINS (Eds.), Representation and Understanding, Morgan Kaufmann, San Diego, 1975, pp. 211–236.
[13] S. ANDERSEN, B. SLATOR, Requiem for a theory: the ' story grammar' story, Journal of Experimental & Theoretical Artificial Intelligence 2 (1990) 253–275.

[14] R. C. Schank, R. Abelson, Scripts, Plans, Goals, and Understanding, Hillsdale, NJ: Earlbaum Assoc, 1977.
[15] N. Chambers, D. Jurafsky, Unsupervised learning of narrative event chains, in: Proceedings of ACL-2007, Association for Computational Linguistics, Columbus, Ohio, 2008, pp. 789–797.
[16] N. Chambers, D. Jurafsky, Unsupervised learning of narrative schemas and their participants, in: Proceedings ACL-2009, Association for Computational Linguistics, USA, 2009, p. 602–610.
[17] K. Pichotta, R. Mooney, Statistical script learning with multi-argument events, in: Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, 2014, pp. 220–229.
[18] E. Rahimtoroghi, E. Hernandez, M. A. Walker, Learning fine-grained knowledge about contingent relations between everyday events, arXiv preprint arXiv:1708.09450 (2017).
[19] J. W. Orr, P. Tadepalli, J. R. Doppa, X. Z. Fern, T. G. Dietterich, Learning scripts as hidden markov models, CoRR abs/1809.03680 (2018).
[20] R. Rudinger, P. Rastogi, F. Ferraro, B. Van Durme, Script induction as language modeling, in: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, 2015, pp. 1681–1686.
[21] A. Belyy, B. Van Durme, Script induction as association rule mining, in: Proceedings of the First Joint Workshop on Narrative Understanding, Storylines, and Events, Association for Computational Linguistics, Online, 2020, pp. 55–62.
[22] L. K. Branting, S. McLeod, K. Megerdoomian, B. Park, Induction of narrative models for legal case elicitation, in: Proceedings of the Sixth Workshop on Automated Semantic Analysis in Law (ASAIL), 2023.
[23] B. Peng, M. Galley, P. He, C. Brockett, L. Liden, E. Nouri, Z. Yu, B. Dolan, J. Gao, Godel: Large-scale pre-training for goal-directed dialog, 2022. arXiv:2206.11309.
[24] J. Lester, K. Branting, B. Mott, Conversational agents, in: M. Singh (Ed.), Practical Handbook of Internet Computing, CRC Press, 2004.
[25] M. Mistica, J. H. Lau, B. Merrifield, K. Fazio, T. Baldwin, Semi-automatic triage of requests for free legal assistance, in: Proceedings of the Natural Legal Language Processing Workshop 2021, 2021, pp. 217–227.
[26] K. Megerdoomian, K. Branting, C. Horowitz, A. Marsh, S. Petersen, E. Scott, Automated narrative extraction from administrative records., in: AIAS@ ICAIL, 2019, pp. 38–48.
[27] C. D. Manning, M. Surdeanu, J. Bauer, J. R. Finkel, S. Bethard, D. McClosky, The stanford corenlp natural language processing toolkit., in: ACL (System Demonstrations), The Association for Computer Linguistics, 2014, pp. 55–60.
[28] G. K. Savova, J. J. Masanz, P. V. Ogren, J. Zheng, S. Sohn, K. C. Kipper-Schuler, C. G. Chute, Mayo clinical text analysis and knowledge extraction system (ctakes), JAMIA 17 (2010) 507–513.
[29] Neo4j, Neo4j -graph data platform, 2023. http://neo4j.com/, last accessed 7 March 2023.
[30] Z. Shen, K. Lo, L. J. Yu, N. Dahlberg, M. Schlanger, D. Downey, Multi-lexsum: Real-world summaries of civil rights lawsuits at multiple granularities, ArXiv abs/2206.10883 (2022).
[31] SPOT, The legal innovation & technology lab's spot api, https://spot.suffolklitlab.org, last accessed March 9, 2023, 2023.
[32] J. Frank, A2j author, legal aid organizations, and courts: Bridging the civil justice gap using document assembly, W. New Eng. L. Rev. 39 (2017) 251.
[33] M. J. Sergot, F. Sadri, R. A. Kowalski, F. Kriwaczek, P. Hammond, H. T. Cory, The British Nationality Act as a logic program, Commun. ACM 29 (1986) 370–386. URL: http://doi.acm.org/10.1145/5689.5920. doi:10.1145/5689.5920.
[34] M. Lauritsen, Q. Steenhuis, Substantive legal software quality: A gathering storm?, in: Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law, ICAIL '19, ACM, New York, NY, USA, 2019, pp. 52–62.