Cognitive Assistance in Government and Public Sector Applications
AAAI Technical Report FS-17-02

Cognitive Assistance for Administrative Adjudication

**L. Karl Branting, Alexander Yeh, Brandy Weiss,
Elizabeth Merkhofer, Bradford Brown**
The MITRE Corporation
7515 Colshire Dr
McLean, VA, 22102, USA
{lbranting,asy,bweiss,emerkhofer,bcbrown}@mitre.org

Abstract

The majority of adjudications in the United States are administrative, so improving the efficiency, accuracy, and consistency of administrative processes could significantly benefit agencies and citizens alike. We propose an approach to cognitive assistance based on identification of salient case features using predictive models induced from previous administrative decisions. We hypothesize that highlighting salient facts will make citizens' and decision makers' interactions with administrative processes more informed, efficient, and accurate. This paper describes three data sets for exploring this hypothesis: motion-rulings, Board of Veterans Appeals (BVA) decisions; and World Intellectual Property Organization (WIPO) domain name dispute decisions. Three different approaches for prediction in these domains were tested. Each approach was capable of predicting outcomes, with the simpler WIPO cases appearing to be much more predictable than BVA or motion-ruling cases. We explore several approaches to using predictive models to identify salient phrases in the predictive texts (i.e., motion or contentions and factual background), and propose a design for displaying this information to decision makers.

Introduction

The majority of adjudications in the United States are administrative, typified by routine licensing, permitting, immigration, and benefits decisions. Notwithstanding the simplicity and uniformity that typically characterizes such cases, citizens often struggle to understand what facts are relevant to their desired outcome and whether their claim is weak or strong. Moreover, the high volume of these administrative cases can lead to inconsistencies, case backlogs, and high resource loads for agency adjudicators.

This paper proposes an approach to cognitive assistance for decision makers and citizens based on identification of salient case features using predictive models induced from previous administrative decisions. More specifically, this approach consists identifying the aspects of the facts of a case that are most relevant to its outcome, retrieving prior cases that share the most relevant similarities to the facts of a given case, and predicting the outcome of the decision. We explore how assistance of this type can be provided by predictive models induced from previous administrative decisions.

Our research agenda consists of three parts. First, we have demonstrated the feasibility of predicting routine administrative decisions from collections of prior decisions. Second, we are exploring techniques for identifying the most relevant portions of case facts based on the predictive model. Finally, we plan to evaluate the degree to which decision-making accuracy and speed can be improved by a tool that highlights the text that appears to be most relevant under the predictive model and that retrieves prior cases based on the similarities between the most relevant text of those cases.

Data Sets

Our ultimate objective is to improve the efficiency, consistency, and transparency of the agencies, such as for veterans benefits, Social Security disability, immigration status, and Medicare appeals, that suffer from long backlogs owing to their inability to handle growing case loads with the available resources. However, privacy and sensitivity issues in these agencies make it desirable to perform initial research on representative but less sensitive data sets. In this section we describe three such data sets used in our initial research.

- **Motion Rulings**
  Our first data set consists of 6,866 motion/order pairs drawn from the docket of a United States federal district court.[1] Motions may be granted, denied, or granted in part and denied in part, and a single order may rule on multiple motions, potentially granting some and denying others. To obviate these procedural complexities, our initial data set is restricted to orders that either rule on a single motion or that have rulings of the same type for multiple motions, i.e., all granted or all denied. Each training instance consists of the text of the motion, which may contain OCR errors (the original filings were in PDF format), together with a classification as either "granted" or "denied."
- **Board of Veterans Appeals Decisions**
  Adjudicative bodies vary in the extent to which case facts and decisions are published. Many adjudicative bodies publish only decisions but not the factual record on which each decision is based. Many agencies, such as the United States Veterans Benefits Administration, publish only appellate decisions, not the original decisions. The

[1] Document filings in US federal courts are "semi-public" in that they are publicly accessible through PACER (https://www.pacer.gov/login.html), but per-page charges and primitive indexing impede wholesale document mining.

absence of published case records can create a cart-and-horse problem in which agencies are unwilling to share sensitive data for an unproven decision-support tool, but the decision-support tool can't be demonstrated because there is no access to the data on which it must be trained.

A method of finessing this problem exploits the convention that decisions generally contain statements of the fact of the case. Decisions with clear sections can be segmented, with the statement of facts treated as a summary of the actual case record, and the decision treated as the classification of those facts in terms of legal outcome. This "bootstrapping" approach was used to demonstrate the feasibility of predicting decisions of the European Court of Human Rights in (Aletras et al. 2016) (Ale16). Of course, decision drafters routinely exclude facts that are irrelevant to the decision and often tailor statements of relevant facts to fit the intended conclusions. As a result, bootstrapping is merely a proxy for the actual task of predicting decisions from raw case facts. However, demonstrating that decisions can be predicted from statements of fact, even if those statements are filtered, is an essential first step in demonstrating the feasibility of prediction in more realistic settings.

Board of Veterans Appeals (BVA) cases[2] have clear sections: Issues, Introduction, Findings, Conclusions, and Reasons. The Issues and Introduction sections contain only facts and contentions, and the decision on each issue is set forth in the Conclusions section. BVA cases often involve multiple issues, but issues are consistently numbered in Issues, Findings, and Conclusions sections. We therefore split each published BVA opinion with n issues into n instances, one for each issue, in which the facts consist of an issue and the entire Introduction, and the classification is extracted (using regular expressions) from the numbered paragraph of Conclusion that corresponds to the Issue (i.e., that has the same numbering). The possible decisions on each issue are (1) the requirements for benefits have been met, (2) the requirements have not been met, (3) the case must be remanded for additional hearings, and (4) the case must be reopened. Conversion of all published BVA cases in this fashion yields 3,844 4-class instances or 1605 2-class (met or unmet) instances.

Unfortunately, the Findings section of BVA cases sometimes contain conclusions about facts not discussed in the Issues and Introduction section, so these sections are an incomplete proxy for the actual case record. This incompleteness makes it impossible in principle to predict the outcome of all BVA cases from just the Issues and Introduction.

- **WIPO Domain Name Dispute Decisions**
  The World Intellectual Property Organization (WIPO) publishes decisions resolving complaints brought against the holder of a domain name that "is identical or confusingly similar" to a trademark belonging to the complainant.[3] WIPO cases have only two possible outcomes: the domain name is transferred to the complainant or it is not. WIPO cases are clearly segmented into seven sections: Parties, Domain Name, History, Background, Contentions, Findings, and Decision. The facts of each instance consist of the concatenation of the first 5 sections, and the classification is "transferred" or "not transferred." The WIPO data set consists of 5587 instances with a roughly 10-to-1 class skew in favor of "transferred."

[2] https://www.index.va.gov/search/va/bva_search.jsp.

[3] http://www.wipo.int/amc/en/domains/decisionsx/index-gtld.html.

Prediction

The first step in confirming the hypothesis that predictive models induced from previous administrative decisions can improve subsequent decision-making processes is to demonstrate that decision outcomes can be predicted. We experimented with 3 predictive techniques: hierarchical attention networks; support vector machines (SVM); and maximum entropy classification.

Hierarchical Attention Networks

In our first approach, we extended the hierarchical neural network model presented in Yang et al. (Yan16) to predict the outcome of legal cases from free-text sections of their case records. The original model takes as input a sequence of sentences. A sentence representation is built for each sentence with a bidirectional gated recurrent unit (GRU) layer over word embeddings. An attention mechanism determines the weight of each time-step's contribution to a sentence vector. Then, a second GRU layer operates over the sentence vectors, an attention mechanism is applied, and the weighted sentence representations are summed to form a hidden document representation. In prior work the document representation was used to predict the ratings of Yelp and movie reviews.

The hierarchical model was extended to account for the deeper structure of legal case documents. With the intuition that human decisions are informed by some combination the text in each section, we altered the model architecture for each dataset. The WIPO cases take as input three sections: history, background and contentions. We feed each section separately into Yang et al.'s document model, sharing weights. The resulting section representations are combined to create the case representation. The architecture used for BVA cases, shown in Figure 1, considers two sections: the issue and the introduction. The issue is nearly always only one sentence, so was treated as a single sequence of words. The introduction may be tens of sentences long and is passed through the hierarchical architecture described in the paper. The case representation is a learned transformation of the issue and introduction sections.

In our experiments, a fully-connected layer appeared to better combine sections' hidden representations than a recurrent layer. We therefore used a hidden layer size of 50 for the WIPO cases and 64 for the BVA cases. We pre-trained word embeddings using the word2vec algorithm of (Mil13). For the WIPO cases, we pre-train on only the WIPO dataset; for the BVA cases, we use a separate dataset of approximately 50,000 appeals. We apply 30% dropout to delay overfitting

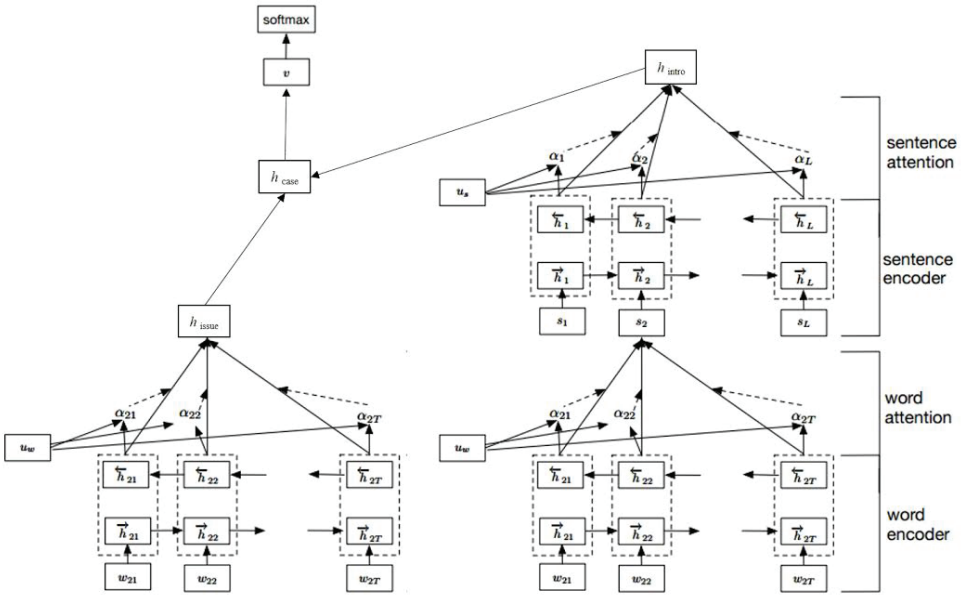

Figure 1: Hierarchical neural model architecture for BVA cases. h_case is a learned function of h_issue, built from the words in the issue section, and h_intro, built from a hierarchical combination of the words-in-sentences and sentences in the case's introduction section.

these small datasets and use the Adam optimizer. Our models are trained on 80% of data, developed with an additional 10%, and the remaining 10% is reserved for testing.

The BVA model achieved a mean F1 of .738 and overall accuracy of 74.7%. The architecture reached a mean F1 of .944 on the WIPO cases, with an F1 of .64 for the 10-times-less frequent negative class. That model has 94.4% accuracy.

Support Vector Machine

The second approach to decision prediction was Support Vector Machine (SVM) learning. For the WIPO and BVA data sets, text was converted into n-gram frequency vectors for n=1–4, with only those n-grams retained that occur at least 8 times. The result was converted into sparse arff format,[4] loaded into WEKA (Hal09), and evaluated in 10-fold cross-validation using WEKA's implementation of Platt's algorithm for sequential minimal optimization(Pla99; Kee01). Because of memory issues, the WEKA SVM was run against only a subset of the entire WIPO data set consisting of 649 instances from each category.

In 10-fold cross validation the SVM approach achieved a mean F1 of 0.731 on the BVA data set, with an overall accuracy of 73%. A mean F1 of 0.950 was achieved on the WIPO data set, yielding an overall accuracy of 90.5%.

Maximum Entropy Classification

The third approach to decision prediction that we explored was Maximum Entropy (Maxent) classification (Ber96) (often termed *logistic regression*). We used the jCarafe[5] implementation of Maxent, which adds regularization to mitigate overfitting, to build a model to predict whether a motion will be granted. Our features consisted of the party filing the motion, the judge ruling on the motion, the sub-type of motion, and the sequences of 1 to 4 tokens (alphanumeric character sequences having non-alphanumeric characters on both the left and right sequence borders) that occur in the text of the motion.

We observed that the motions contain many tokens that appear only in one motion and seem to be the result of OCR errors (as noted above, the documents were filed in PDF format, and some were created by scanning images to PDF). To remove these artifacts, any token that only appeared in only one motion in the collection was removed.

There are many different sub-types of motions, e.g., for extension of time to file, for summary judgment, etc. We found that better accuracy was obtained by training separate models for motion subtypes rather than training a single model for all subtypes. Accordingly, we split the motions into the following 3 large classes of sub-types and build a separate prediction model for each class:

- Extension-type motions, such as a motion to extend a filing due date, which tend to have higher grant rates than motions in general
- Motions of the letter sub-type, which tend to have a slightly lower grant rate than motions in general
- Motions not included in either of the 2 classes above

We used 10-fold cross validation to build and test separate models for each of the 3 large classes above and then combined the results. The combined results were an accuracy of 75%, and the recall, precision and balanced F-score for "granted" were 54%, 66% and 59% respectively.

[4] http://www.cs.waikato.ac.nz/ml/weka/arff.html

[5] https://github.com/wellner/jcarafe

	max-ent	SVM	HAN
motion-rulings	0.742	0.757	
BVA		0.731	0.738
WIPO		0.950	0.944

Table 1: Frequency-weighted mean F1 for predictive algorithms applied to three decision data sets. Note that the SVM result on the WIPO data set is on a balanced subset, rather than the entire WIPO collection, whereas the hierarchical attention network was applied to the entire skewed set.

The predictive results for the three experiments are summarized in Table below:

Decision Support

The results of the prediction experiments indicate that routine adjudications and orders are predictable to a degree that depends on the complexity of the underlying domain from models trained from text representing the facts of the case (in the WIPO and BVA data sets) or the motion text (for the order-prediction data set). Since this approach does not perform argumentation mining and has no explicit model of the applicable legal issues and rules, there is a limit to the predictive accuracy that this approach can achieve except in highly routine and predictable domains, such as WIPO decisions. However, our objective isn't replacement of human discretion, but rather support for human decision making. Our hypothesis is that predictive models can assist human decision makers by identifying the portions of the predictive text, e.g., statements of case facts or motion texts, that are most predictive of the outcome. We hypothesize that a decision maker may benefit from having the predictive text identified even when the decision disagrees with the models prediction. This hypothesis is based on the observation that one of the challenges of decision making is sifting through irrelevant portions of the case record to locate the most important facts.

We distinguish two uses of predictive text:

- Highlighting the parts of a document most relevant outcome, e.g., granting or denying a motion, or accepting or rejecting a claim for benefits, so that the decision maker can quickly identify the facts determinative of the outcome.
- Highlighting the parts of one document most relevant to assessing the similarity or difference between the cases. The Common-Law doctrine of *stare decisis*, under which a decision in one case is binding on subsequent similar cases, is generally inapplicable to administrative adjudications, even in countries with Common Law legal systems. Nevertheless, we hypothesize that enabling decision makers to compare the current case to the most similar prior cases could make decision making faster and more consistent.

We therefore turn to the issue of how predictive texts can be identified. In the context of algorithms for prediction based on text, identification of the most predictive text is a special case of the more general problem of feature selection (Guy03).

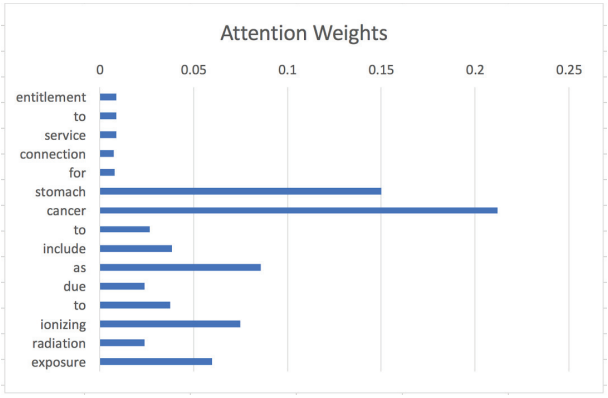

Figure 2: Attention weights for words in a representative sentence in the introduction of a BVA case. These weights reflect the relative importance of the words in determining the label assigned by the hierarchical attention network to the case.

Salient Fact Detection

Neural Network Attention Attention mechanisms for neural networks allow the network to learn to weights as part of its representation. Bahdanau et al. (2014) (Bah14) introduced neural attention for natural language processing, learning a soft alignment for machine translation such that certain input words contribute most to output words. In the context of text classification, the attention mechanism determines relative contributions of words in the input sequence to the prediction. This hierarchical model has attention over the words in each sentences and over the sentences that make up each section. The attention operates on output from a bidirectional recurrent layer, meaning that each time-step folds in context from surrounding words or sentences but is most responsive to the word or sentence at that time-step.

Extracting the attention weights enables analysis of the model's predictions. Figure 2 illustrates the word-level weights of a representative sentence from a BVA case, showing that words related to the veteran's medical condition had a disproportionate weight in the neural network's prediction.

The use of attention weights to identify the texts most salient to a decision maker is illustrated Figure 3, which shows an excerpt from a BVA decision in which the highest-attention sentence are highlighted in color. The sentence in blue received 74% of the attention weight, and the next most important, shown in yellow, received 9% of the attention weight. The sentence in blue is, in fact, highly relevant in that it recharacterizes the issue from being moot into something that can be granted.

While the accuracy of the hierarchical attention network is similar to the other learning models, it has the significant advantage that attention weights are specific to token instances rather than, as in the case of linear model weights, global. A phrase that is insignificant in one context can be very significant in a different context; this distinction can be identified

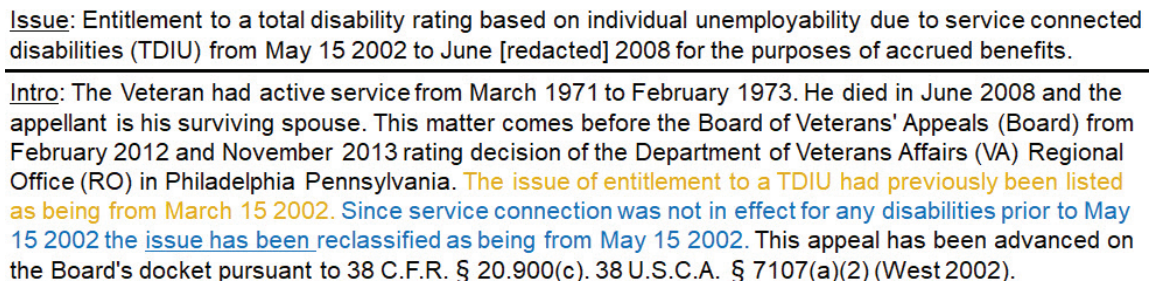

Figure 3: A portion of a BVA case. The sentence with the highest proportion of attention weight, 74%, is shown in blue, and the sentence with the next highest weight, 9%, is shown in yellow.

by hierarchical attention networks but not by models that produce global weights.

Linear Model Weights An alternative approach to using a predictive model to identify the most salient case facts makes use of the feature weights learned during model training. In linear models, such as maximum entropy and SVM with linear kernels, feature weights are indicative of relevance of features to the model's predictions (Mla04), which in our application consists of predicted case decisions. These feature weights will have a similarity to point-wise mutual information (PMI) above, with features increasing the chances of a positive prediction tending to have positive weights and features decreasing the chances tending to have negative weights. One difference between PMI and the feature weights for logistic regression with regularization is that a regression feature weight differs when a feature either only occurs infrequently (weight magnitude is diminished) or is correlated with other features in the model (weight is adjusted for the effects of correlated features on the model).

For example, when maximum entropy/logistic regression (with regularization) is applied to the domain of motions for an extension of time, phrases having a relatively large positive feature weight for "granted = true" include "to dismiss" (feature weight of 0.319), "dismiss" (0.310) and "with the consent" (0.285), whereas phrases with relatively large negative feature weights include "a stipulation" (-0.385) and "stipulation" (-0.735).

User Interface for Decision Support

The data derived from the predictive models is intended to be displayed to administrative claimants and decision makers using a Graphical User Interface (GUI). The complexity of the displayed content along with a need to prevent user error necessitates a usable interface. We are therefore exploring interface designs to present this information in a manner that best facilitates its use, that is, to present data in a manner that improves a decider's speed and accuracy (Nie93). The preliminary design concept shown in Figure 4 contains several features that we hypothesize will assist users with deciding on cases efficiently and accurately. One feature of this design concept allows a user to view the most relevant cases in multiple ways (i.e., multi-case comparison, high-level comparison, in depth comparison). Providing multiple formats for case comparison allows a user the flexibility to decide how in-depth they would like to view the current case and previous case information. Another design feature provides the user with convenient access to relevant information (e.g., the rules) during the review process. This design concept leverages the pattern of open/close panels, which allow the user the ability to customize their view as they go through the evaluation process. In order to support efficient comparison, this design also provides a highlighting feature that is intended to allow a user to compare the similarities between current and previous cases. A future evaluation of this design concept will provide the information needed to iterate on the design patterns and features. The overall goal is to provide a satisfactory user experience while also assisting the decision maker to make quick and accurate assessments of cases.

We plan on conducting an initial experimental evaluation to assess the overall ability of the combined predictive model and user interface to facilitate improved speed and accuracy in decision making. We hypothesize that the speed and accuracy of decision making can be improved by highlighting the phrases in cases with facts having the greatest weight under a predictive model and by retrieving prior cases with the strongest similarity to the current case in terms of the highest weight phrases. The initial evaluation will be performed using non-lawyers as subjects and WIPO case outcome detection as the predictive task, since WIPO cases have relatively simple and predictable facts and issues. After initial evaluation, we plan to conduct future evaluations that validate this concept using lawyers and other end-users of this type of decision support system.

Summary and Future Work

This paper proposes an approach to cognitive assistance that uses predictive models induced from previous administrative decisions to highlight relevant portions of the case record, for retrieval of relevant prior cases, and for prediction. Three data sets were developed: motion-rulings, BVA issue decisions, and WIPO domain name dispute decisions. The ability to predict outcomes in these three domains was demonstrated using three different approaches for prediction: maximum entropy over token n-grams; SVM over token n-grams; and a hierarchical attention network applied to the full text. This initial evaluation did not establish the superiority of one approach over another, but rather indicates that the outcome

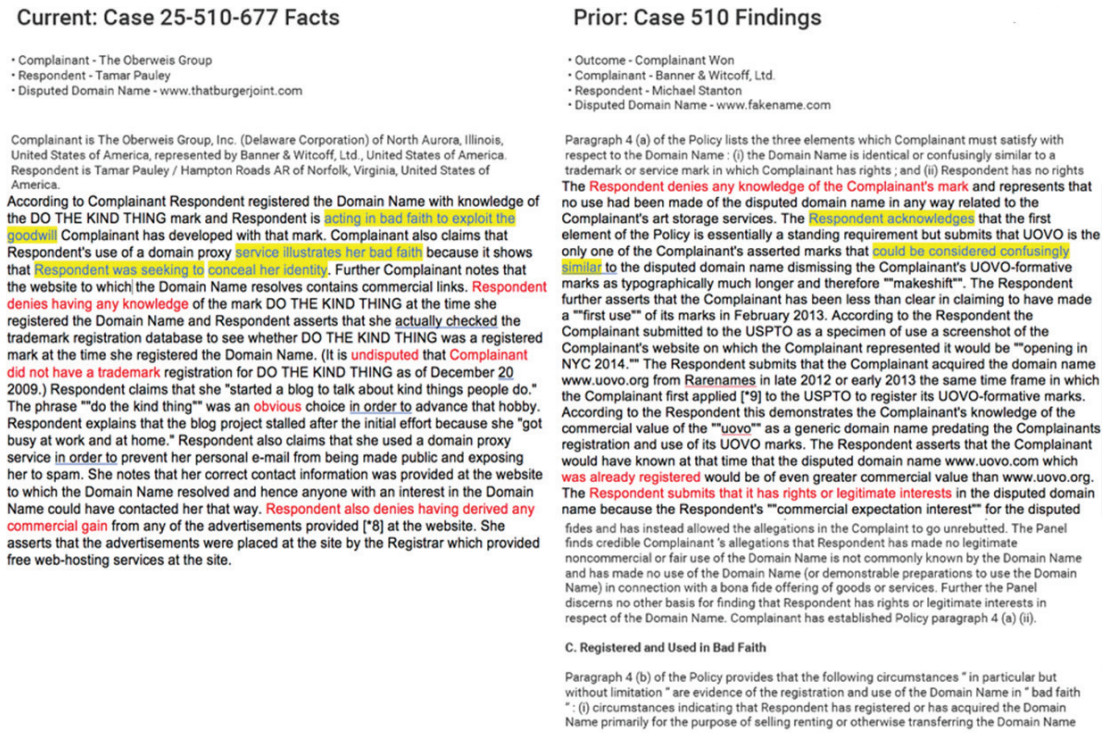

Figure 4: A prototype decision-support interface design illustrating how the most salient case facts can by highlighted to assist with analysis of the case record and comparison between cases. Phrases highlighted in yellow are associated with rulings in favor of complainants, whereas phrases highlighted in red are associated with rulings in favor of respondents.

of routine decisions is predictable using multiple alternative models from the text of the motion or contentions and factual background alone and that predictive accuracy varies depending on the domain and the nature of the predictive texts.

We are currently exploring techniques for integrating network attention weights or feature weights from other predictive models into our decision-support tool. We have designed and plan to execute an empirical evaluation that will measure the degree to which salient text highlighting impresses speed and accuracy of decision making by non-lawyers and by lawyers not familiar with WIPO rules.

The ultimate objective of this work is to improve the efficiency, accuracy, and consistency of administrative decision making, the form of adjudication that has the greatest impact on most citizens, by integrating automated decision models into the human decision process. This work represents an initial step towards this objective.

Acknowledgments

The MITRE Corporation is a not-for-profit company, chartered in the public interest, that operates multiple federally funded research and development centers. This document is approved for Public Release; Distribution Unlimited. Case Number 17-1719.

References

Aletras, N., Tsarapatsanis, D., Preotiuc-Pietro, D., Lampos, V.: Predicting judicial decisions of the European Court of Human Rights: a natural language processing perspective. PeerJ CompSci (October 24 2016), https://peerj.com/articles/cs-93/

Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473 (2014)

Berger, A.L., Pietra, V.J.D., Pietra, S.A.D.: A maximum entropy approach to natural language processing. Computational Linguistics 22(1), 39–71 (Mar 1996), http://dl.acm.org/citation.cfm?id=234285.234289

Branting, L.K.: Data-centric and logic-based models for automated legal problem solving. Artificial Intelligence and Law 25(1), 5:27 (2017), http://dx.doi.org/10.1007/s10506-017-9193-x

Gallager, R.G.: Information Theory and Reliable Communication. John Wiley & Sons, Inc., New York, NY, USA (1968)

Guyon, I., Elissee, A.: An introduction to variable and feature selection. J. Mach. Learn. Res. 3, 1157–1182 (Mar 2003), http://dl.acm.org/citation.cfm?id=944919.944968

Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The weka data mining software: An update. SIGKDD Explorations 11(1) (2009)

Keerthi, S., Shevade, S., Bhattacharyya, C., Murthy, K.: Im-

provements to platt's smo algorithm for svm classifier design. Neural Computation 13(3), 637–649 (2001)

Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. CoRR abs/1301.3781 (2013), http://arxiv.org/abs/1301.3781

Mladenic, D., Brank, J., Grobelnik, M., Milic-Frayling, N.: Feature selection using linear classifier weights: Interaction with classification models. In: Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 234–241. SIGIR '04, ACM, New York, NY, USA (2004), http://doi.acm.org/10.1145/1008992.1009034Keerthi2001

Nielsen, J.: Usability Engineering. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1993)

Peterson, M., Waterman, D.: Rule-based models of legal expertise. In: Walters, C. (ed.) Computing Power and Legal Reasoning, pp. 627–659. West Publishing Company, Minneapolis, Minnesota (1985)

Platt, J.C.: Advances in kernel methods. chap. Fast Training of Support Vector Machines Using Sequential Minimal Optimization, pp. 185–208. MIT Press, Cam- bridge, MA, USA (1999), http://dl.acm.org/citation.cfm?id=299094.299105

Sidhu, D.: Moneyball sentencing. Boston College Law Review 56(2), 672–731 (2015)

Yang, Z., Yang, D., Dyer, C., He, X., Smola, A., Hovy, E.: Hierarchical attention networks for document classification. In: Proceedings of NAACL-HLT. pp. 1480– 1489 (2016)