

# ML Techniques for Labeling Legal Opinions

Christopher Hundt chundt@eecs.berkeley.edu

## Automatic topic labeling of legal opinions

Legal opinions—written accounts by judges of the reasoning used in deciding a case—are an essential part of the U.S. legal system (and that of many other countries). Systems to organize legal opinions are of tremendous value to lawyers.

One important problem in such system is to determine the topic of a legal opinion, i.e., what area of law it concerns. This project aims to investigate whether this can be done with machine-learning techniques.

To the right is a passage from a recent opinion from the United States Court of Appeals for the Ninth Circuit with some interesting features highlighted.

1. Councils first argue that the questions Roommate poses to prospective subscribers during the registration process violate the Fair Housing Act and the analogous California law. Councils allege that requiring subscribers to disclose their sex, family status and sexual orientation “indicates” an intent to discriminate against them, and thus runs afoul of both the FHA and state law.<sup>13</sup>

[4] Roommate created the questions and choice of answers, and designed its website registration process around them. Therefore, Roommate is undoubtedly the “information content provider” as to the questions and can claim no immunity

<sup>13</sup>While the Conference Report refers to this as “[o]ne of the specific purposes” of section 230, it seems to be the principal or perhaps the only purpose. The report doesn’t describe any other purposes, beyond supporting “the important federal policy of empowering parents to determine the content of communications their children receive through interactive computer services.” H.R. Rep. No. 104-458, at 194 (1996) (Conf. Rep.), as reprinted in 1996 U.S.C.A.N. 10, 207-08.


<sup>14</sup>The Fair Housing Act prohibits any “statement . . . with respect to the sale or rental of a dwelling that indicates . . . an intention to make [a] preference, limitation, or discrimination” on the basis of a protected category. 42 U.S.C. § 3605(c) (emphasis added). California law prohibits “any written or oral inquiry concerning the” protected status of a housing seeker. Cal. Gov. Code § 12955(b).

## Text features: bag of words

The basic features I used for learning from the text of the opinions are the so-called “bag-of-words features,” in which a document is represented as an unordered collection of the words (or short phrases) that appear in it. Examples of interesting words or phrases are highlighted in orange above.

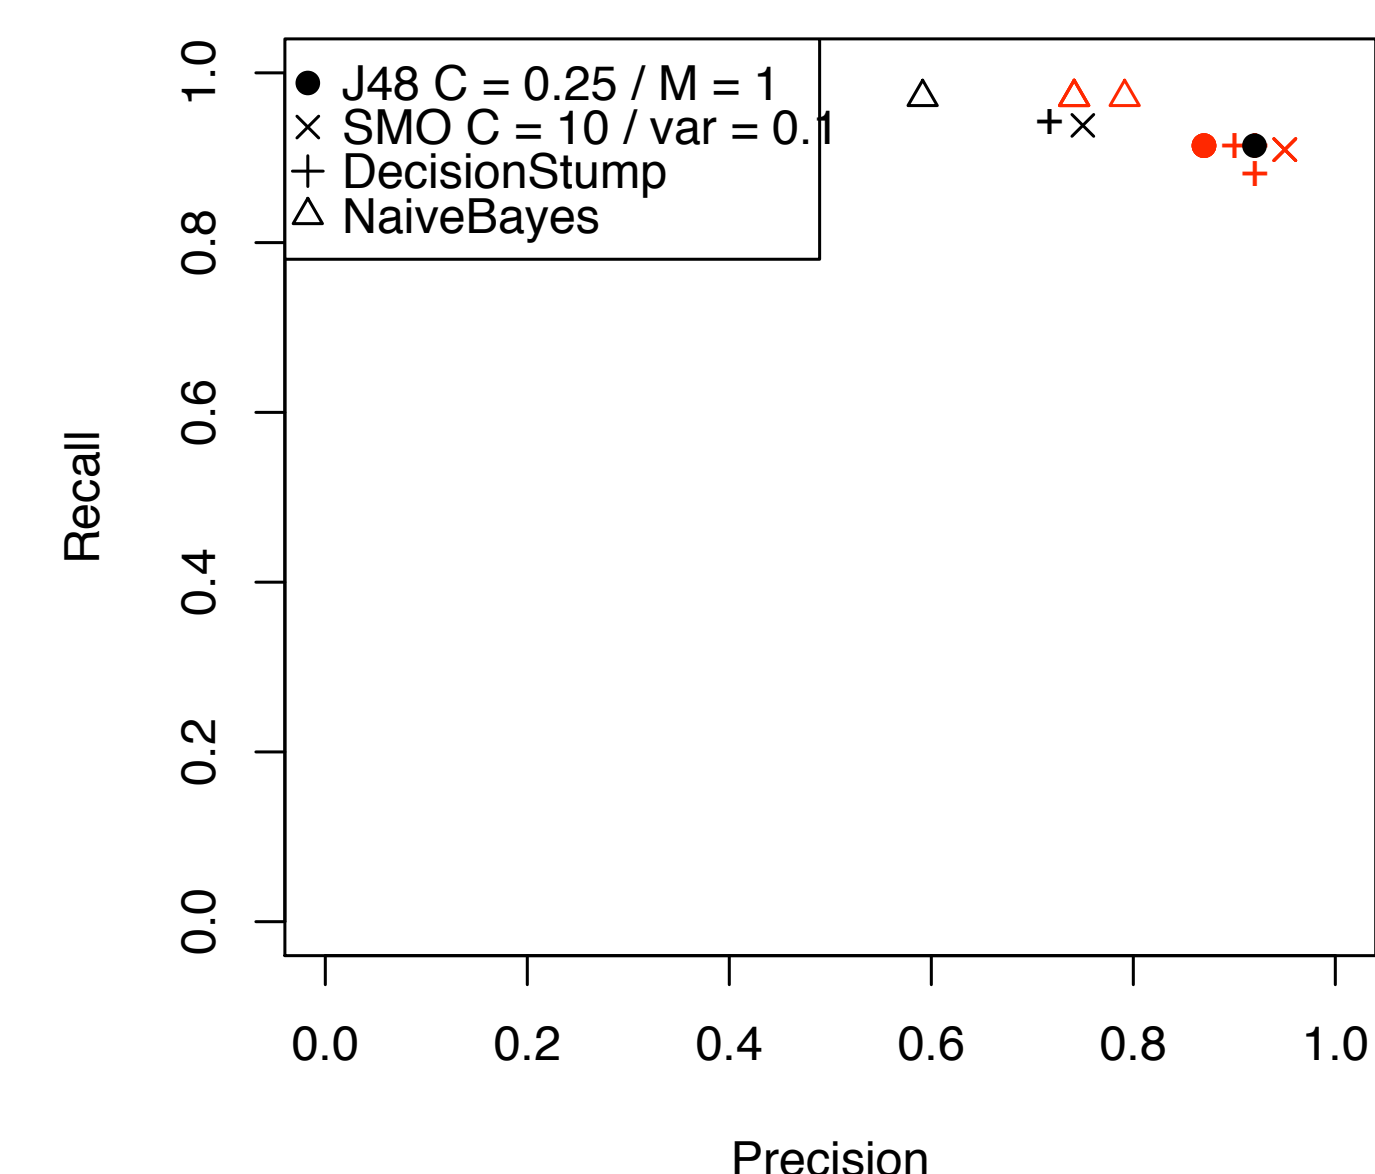
## Choosing classifier parameters

In using both SVMs (SMO algorithm) and decision trees (J48 algorithm) there are parameters to choose. Shown below are some of the effects of parameter choice.



## The effect of boosting

The metalearning algorithm AdaBoost is useful for strengthening weak classifiers. I tested it on various classifiers on this data with the bag-of-words features. Results are shown to the right. (The red marks are the boosted results.) Generally, AdaBoost did not help significantly when classifiers already got good results, but did help for classifiers that are truly weak on this data, like Decision Stump and Naïve Bayes.



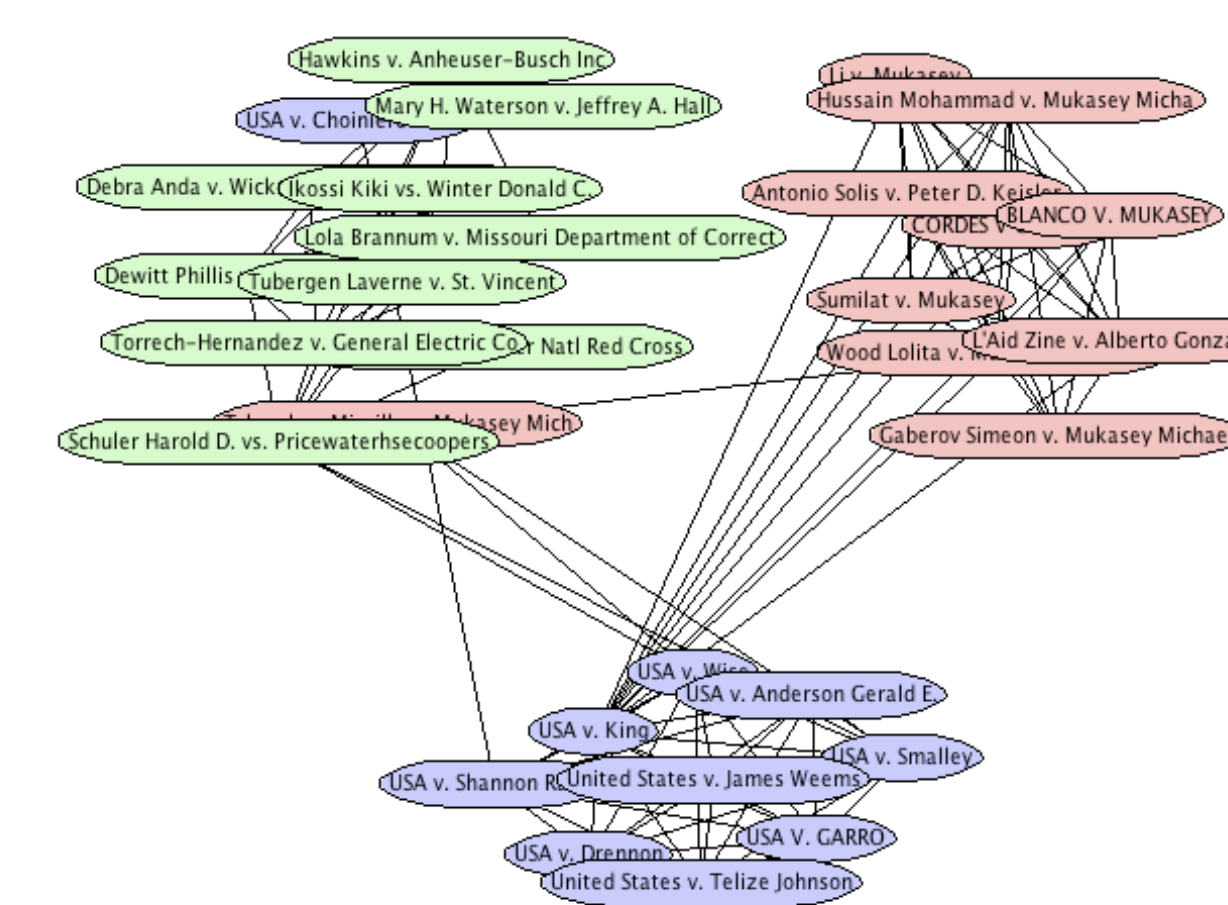
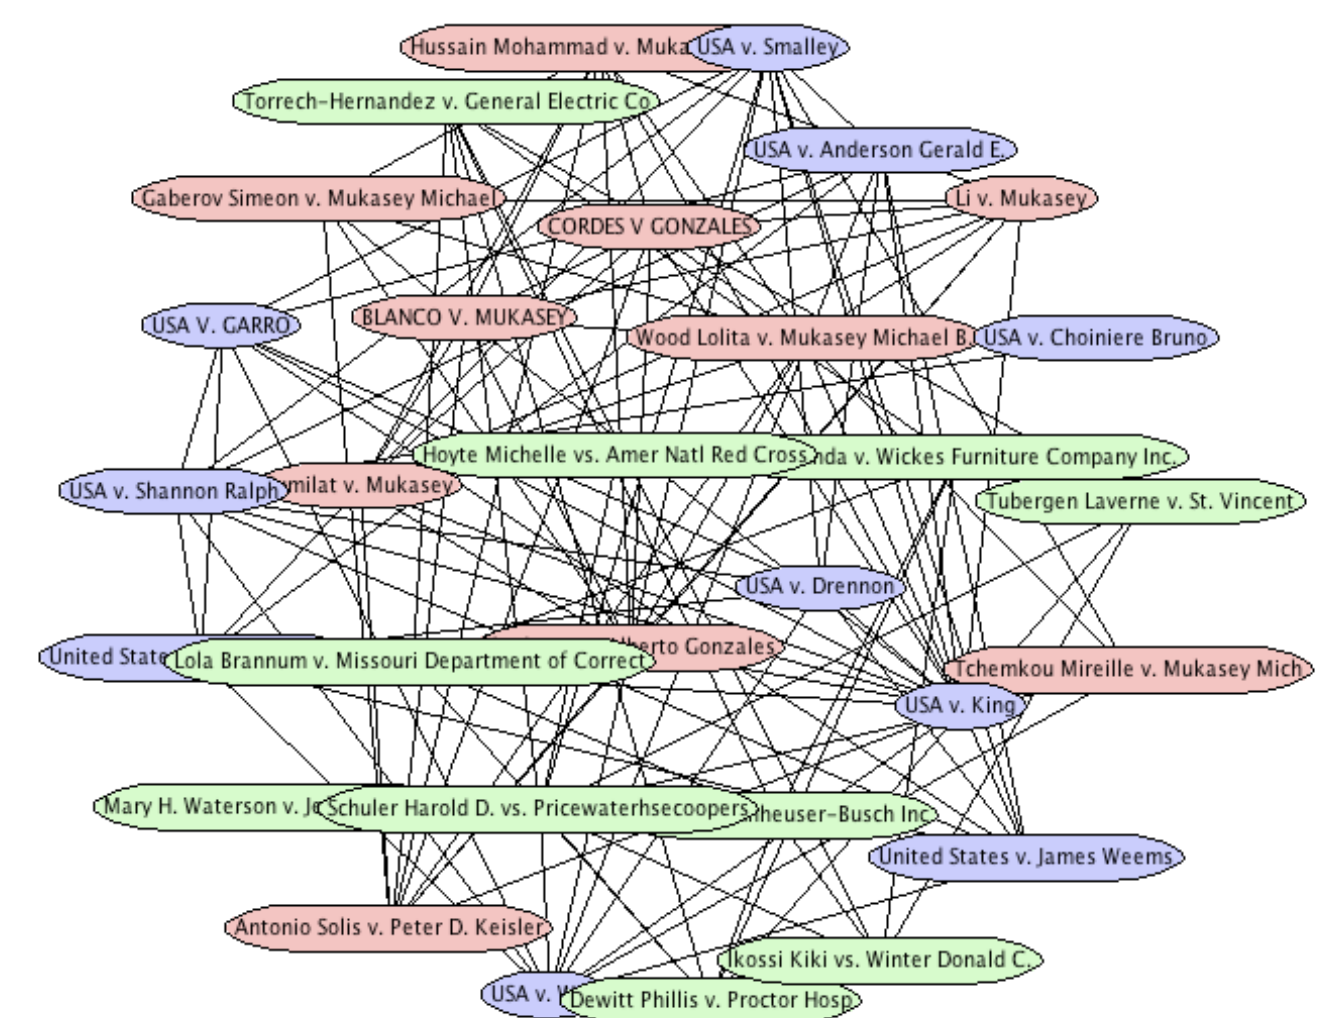
## Incorporating citation information

Legal opinions tend to cite many previous legal opinions, statutes, codes, regulations, and other authorities, and they tend to use a uniform citation format—some examples are highlighted in purple on the left. These citations may provide information regarding the topic of the opinion.

## Social networks and blockmodeling

One possibility is to use the citation information to establish “relationships” among opinions and use methods of learning from social networks to inform the classifier.

In order to do this I call two opinions “friends” if they both cite the same other opinion, statute, or other authority. On the right is a network of 30 labeled opinions. The colors indicate the labels: Immigration, Discrimination, and Sentencing.



Stochastic blockmodeling is a social-network model in which each node (opinion) in the network belongs to a single group and the probability of a relationship between two nodes depends only on the which groups they belong to.

Using only the relationship information, one can do inference to find the likely block structure in the network under the model assumptions. On the left the likely block structure is shown for the same collection of opinions as above (after inference by Gibbs sampling). The groups correspond very closely to the labels, even though labels were not inputs to the model.

## Experimental results

Generally, blockmodeling had significant classification power but, in the presence of all the text features, did not result in a major improvement. To create these charts, the J48 learner was used on subsets of most informative text features both with and without additional blockmodeling information. Blockmodeling helps significantly when there is little textual information (zero or one text features from the top ten and most subsets of features 11–20), but not much when the text features are very informative.

