# Ideological Drifts in the U.S. Constitution: Detecting Areas of Contention with Models of Semantic Change

**Abdul Z. Abdulrahim**
University of Oxford
Oxford, UK
abdul.abdulrahim@cs.ox.ac.uk

## Abstract

Recent years have seen a marked decline in public support for major US political institutions. This falling support could threaten the Supreme Court's legitimacy as a study on Congressional law-making found that Congress is more than twice as likely to overturn a Supreme Court decision when public support for the Court is at its lowest compared to its highest level. Using a corpus of federal court opinions and congressional speeches, we develop a method for tracking areas of ideological differences and how these have drifted closer or further from each other over the years. Building on existing models of tracking semantic change, we compare multiple approaches and evaluate the stability, coherence (intrinsic valuation) and predictive power (extrinsic valuation) of each model to determine the most appropriate method for our analysis and similar work in the area of legal and political informatics.

## 1 Introduction

In the United States of America (U.S.) — much like Europe — we have seen declining public support for major political institutions, and a general disengagement with the processes or outcomes of the three branches of government under the constitution. According to Pew's Public Trust in Government survey, "public trust in the government remains near historic lows" — with only 14% of Americans stating that they can trust the government to do what is right "most of the time". To better understand ideological drifts in the branches of government, we use models of semantic change to track the evolution of issues around reproductive rights and the potential bias inherent in the language used in discussing it. We use abortion as a case study because of the historical legal context and shifts in the Courts and Congress on the issue.

When using models of semantic change, a common problem faced in diachronic or temporal analysis of distributed text representation is the comparison of word vectors across different time periods. While it is relatively uncomplicated to train separate word representation models using time-specific corpora, it is more complicated to compare word vectors across different models. To overcome this, various suggestions have been made, but we focus on the primary approaches that have been adopted (incremental updates [Kim et al., 2014] and diachronic alignment [Hamilton et al., 2016]). Kim et al. [2014] proposed incrementally updated embedding models; that is, training a model on the year $y_i$, and then the model for the year $y_{i+1}$ is initialised with the word vectors from $y_i$. As such, all the models are inherently related to each other, allowing direct comparison of cosine similarities between the same word in different time periods. However, a drawback with using the incremental updates approach is that it is not as sensitive to the cultural shifts in the corpora because the word vectors are trained on a widening window that dulls the effects of new usage statistics [Dubossarsky et al., 2019]. Hamilton et al. [2016], in contrast, use orthogonal Procrustes to align the learned low-dimensional embeddings from each time period. By defining $W^{(t)} \in \mathbb{R}^{d\times|\nu|}$ as the matrix of the

Preprint. Under review.

word representation learned at time period $t$, they align across time-periods while preserving cosine similarities. This approach overcomes the issue with detecting local shifts noted with incremental updates, but exhibits some incoherence and instability with a small corpus [Dubossarsky et al., 2016, Kutuzov et al., 2018, Wendlandt et al., 2018].

## 2 Experimental Setup

To compare the ideological stance of the U.S. Congress and Supreme Court, we use two datasets that provide an insight into the opinions and views of the members of each branch of government. For the U.S. Congress, we use the text of speeches made by the members of Congress. For the U.S. Supreme Court, we use the opinions for cases tried in the appeals courts and Supreme Court. We expanded the scope of opinions used to include the appeals courts as the U.S. Supreme Court does not provide sufficient data to generate stable word representations, — particularly when using the diachronic alignment method. These are the closest datasets to the constitutional process of lawmaking in the U.S.

In preparing the data for our word representation algorithm, the following data wrangling and preprocessing steps were taken: (i) we lower-cased all tokens in the corpus before extracting word-context pairs; (ii) for pair extraction we chose a window size of 5 for all models; (iii) corpus tokens were skipped as word or context if they did not have a minimum frequency of 200 in the full corpus used or contained non-alphanumeric characters; and (iv) we regularised the text of the opinions to capture case or legislation citations. Some key characteristics of the corpora we use following preprocessing are provided in table 1.

Table 1: **Descriptive Statistics.** Key characteristics of our corpora, including the number of cases and speeches, unique words and the average number of words per case opinions or speeches.

| **Characteristics** | **U.S. Fed. Court Opinions** | **U.S. Congress Speeches** |
|---|---:|---:|
| No. of opinions/speeches used: | 1,077,605 | 16,031,032 |
| No. of judges/speakers: | ∼600 | ∼6,000 |
| Time period: | 01/01/1955 - 31/12/2014 | 03/01/1955 - 03/01/2015 |
| Size of corpus (in bytes): | 10.7GB | 14.3GB |
| Average length (in words): | 1,590 | 150 |
| No. of unique words: | 4,351,232 | 3,267,619 |
| No. of unique words (>100): | 131,632 | 94,431 |
| No. of unique words (>200): | 81,596 | 65,878 |

Conceptually, we formulate the task of discovering semantic shifts as follows. Given a time sorted corpus: $(corpus_1, corpus_2, ..., corpus_n)$, we locate our target word and its meanings in the different time periods. We chose the *word2vec* algorithms based on previous work by Antoniak and Mimno [2018]. To tune our algorithm, we performed two main evaluations (intrinsic and extrinsic) on samples of our corpora, comparing the performance across different hyperparameters (window size and minimum word frequency). We first evaluate the stability and coherence of the local changes reflected in the nearest neighbours; and second, we assess each model's predictive power. In tuning our hyperparameters, we ended up with the following: MIN = 200 (minimum word frequency), WIN = 5 (symmetric window cut-off), DIM = 300 (vector dimensionality), CDS = 0:75 (context distribution smoothing), K = 5 (number of negative samples) and EP = 1 (number of training epochs).

## 3 Results

To contextualise our evaluation of the method and the significance of the results we obtain, we use the assessment from an expert interview to categorise each word in table 3 and 4 as: (i) **a medically descriptive word**, i.e., it relates to common medical terminology on the topic; (ii) **a legally descriptive word**, i.e., it relates to case, legislation or opinion terminology; and (iii) **a potentially biased word**, i.e., it is not a legal or medical term and thus was chosen by the user as a descriptor.

We see from our results for the incremental updates and diachronic alignment in table 3 that in the decade up to 1965, the Courts had negative connotations with abortion but changed thereafter. Using

Table 2: **Nearest Neighbours Table Key.** Description of keys used to classify words in the nearest neighbours by type of terminology. These were based on the insights derived from expert interview.

| Category | Colour Code | Description |
|---|---|---|
| Medical | Yellow | Any medical or scientific terminology that is used to describe a person, condition or procedure e.g. "physician", "pregnant", "saline", etc. |
| Legal | Green | Any legal terminology relating to case names, statute or phrases used in opinions e.g. "roe v. wade", "medicaid", "partial-birth" etc. |
| User-Chosen | Orange | Other words that are not classed as medical or legal chosen by the user and as such could be biased e.g. "sodomy", "rape", "pro-life", etc. |

these terms associated with abortion, we are able to detect the key issues of that time period, though it is arguable whether there is any strong bias (positive or negative) in the Court's interpretation of the word. In contrast, our results from the corpus of speeches produced by the U.S. Congress show a clear view of the collective ideology held by the members. These are individuals that — of late — have been praised for some of their controversial opinions on the issues of reproductive rights. While it has been phrased as a red versus blue issue, the debate on abortion has seen some Democratic candidates campaign for legislation that restricts the rights granted to women.

## 4 Discussion

To formally measure semantic displacement, we compute the cosine distance (Cosine Distance $=$ $1-$Cosine Similarity) between a word's representation for different time-periods, i.e. $dist(w_t, w_{t+\delta})$, as a measure of semantic change. This is shown in figure 1, and is used to quantify rates of semantic change for our target word ("abortion") by looking at the displacement between consecutive time-points. We notice less displacement in the $SGNS_{IN}$ model and a higher yearly displacement in the court opinions than the congressional speeches. This perhaps confirms the view that although the interpretation of abortion has not changed significantly in recent years, more shifts have been made by the courts.

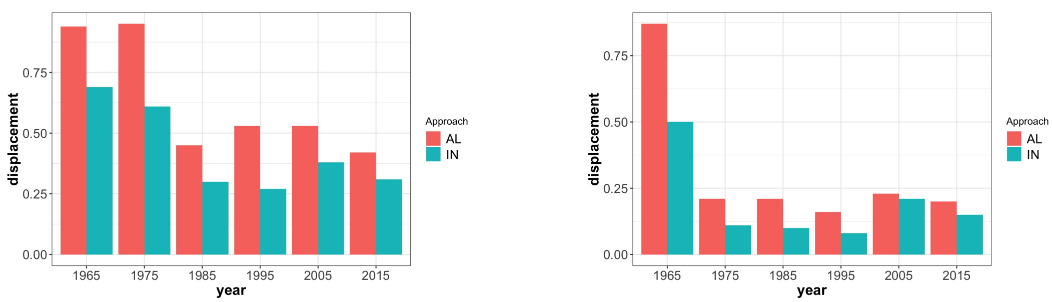


(a) Court Opinions

(b) Congress Speeches

Figure 1: **Semantic Displacement of Target Word.** This shows the semantic displacement from one time period to the next for our target word.

Overall, the results show that within the issue of reproductive rights, it is harder to detect how the ideology of a judge might feed into the words used in relation to the term "abortion". The same cannot be said for members of Congress who blatantly use or infer negative connotations with the term and in turn reveal the prevailing ideology over the time period we analyse. We also contend that this bias — overt or not — can determine how a case is decided to avoid conflict with Congress, or legislation is implemented to override a ruling by the Supreme Court. Our approaches using incremental updates and diachronic alignment proved insightful at detecting micro-cultural shifts in the corpora. We see good performance in the detection of relevant cases during a time period and also observe that late-term abortion is likely to be a debating point in the upcoming 2020 U.S. presidential elections. This might lead to the revival of the Pain-Capable Unborn Child Protection Act, which intends to — in most cases — make it unlawful to perform an abortion if the estimated post-fertilisation age of a fetus is 20 weeks or more.

Table 3: **Local Changes in U.S. Federal Court Opinions.** The top 10 nearest neighbours to the target word "abortion" ranked by cosine similarity for each model.

INCREMENTAL UPDATES

| 1965 | 1975 | 1985 | 1995 | 2005 | 2015 |
|---|---|---|---|---|---|
| *sodomy* | *pregnant* | *parent* | *fetus* | *clinic* | *clinic* |
| *chemically* | *unborn* | *roe v. wade* | *unborn* | *unborn* | *unborn* |
| *pharmaceutically* | *griswold v. conn.* | *school-age* | *clinic* | *fetus* | *abort* |
| *polyurethane* | *parent* | *doe v. bolton* | *medically* | *partial-birth* | *fetus* |
| *prolix* | *therapeutic* | *immature* | *saline* | *i.u.d.* | *i.u.d.* |
| *carnal* | *medicaid* | *anonymity* | *medicaid* | *pregnant* | *pregnant* |
| *trafficker* | *sterilization* | *sterilization* | *pregnant* | *woman* | *partial-birth* |
| *smallest* | *doe v. bolton* | *fetus* | *woman* | *reproductive* | *late-term* |
| *transmits* | *physician* | *unborn* | *hyde amendment* | *physician* | *woman* |
| *steal* | *sodomy* | *life-saving* | *demonstration* | *abort* | *physician* |

DIACHRONIC ALIGNMENT

| 1965 | 1975 | 1985 | 1995 | 2005 | 2015 |
|---|---|---|---|---|---|
| *prolix* | *pregnant* | *unemancipated* | *clinic* | *sterilize* | *women* |
| *polyurethane* | *physician* | *parent* | *medicaid-eligible* | *i.u.d.* | *clinic* |
| *sodomy* | *trimester* | *immature* | *fetus* | *contraception* | *reproductive* |
| *chemically* | *saline* | *school-age* | *roe v. wade* | *clinic* | *mcallen* |
| *trafficker* | *parent* | *anonymity* | *parent* | *partial-birth* | *partial-birth* |
| *carnal* | *medical* | *unborn* | *reproductive* | *f.g.m.* | *a.s.c's* |
| *conspiracy* | *roe v. wade* | *fetus* | *unborn* | *unborn* | *parent* |
| *steal* | *unborn* | *life-saving* | *protester* | *pregnancy* | *fetus* |
| *mayhem* | *doe v. bolton* | *well-informed* | *pregnancy* | *fetus* | *roe v. wade* |
| *purloin* | *munday* | *hyde amendment* | *hyde amendment* | *reproductive* | *i.u.d.* |

Table 4: **Local Changes in U.S. Congress Speeches.** The top 10 nearest neighbours to the target word "abortion" ranked by cosine similarity for each model.

INCREMENTAL UPDATES

| 1965 | 1975 | 1985 | 1995 | 2005 | 2015 |
|---|---|---|---|---|---|
| *sterilization* | *anti-abortion* | *anti-abortion* | *anti-abortion* | *partial-birth* | *partial-birth* |
| *unborn* | *sterilization* | *contraception* | *partial-birth* | *pro-life* | *pro-life* |
| *contraception* | *contraception* | *pregnancy* | *pregnancy* | *pregnancy* | *contraception* |
| *therapeutic* | *unborn* | *reproductive* | *infanticide* | *infanticide* | *pregnancy* |
| *fetus* | *fetus* | *pro-life* | *third-trimester* | *rape* | *infanticide* |
| *pregnant* | *rape* | *infanticide* | *rape* | *trimester* | *abortionist* |
| *pro-life* | *sodomy* | *parent* | *abortionist* | *reproductive* | *trimester* |
| *anti-abortion* | *infanticide* | *clinic* | *clinic* | *woman* | *woman* |
| *back-alley* | *pro-life* | *trimester* | *reproductive* | *contraception* | *late-term* |
| *pro-abortion* | *abortionist* | *ectopic* | *woman* | *abortionist* | *self-induced* |

DIACHRONIC ALIGNMENT

| 1965 | 1975 | 1985 | 1995 | 2005 | 2015 |
|---|---|---|---|---|---|
| *sherri* | *sterilization* | *anti-abortion* | *abortion-related* | *partial-birth* | *late-term* |
| *harrassment* | *unborn* | *non-therapeutic* | *pregnancy* | *late-term* | *partial-birth* |
| *preemptively* | *contraceptive* | *unborn* | *anti-abortion* | *third-trimester* | *sterilization* |
| *disingenuous* | *non-therapeutic* | *pregnancy* | *reproductive* | *pregnancy* | *abortionist* |
| *soviet-style* | *legalized* | *pro-abortion* | *pro-life* | *pro-life* | *unborn* |
| *beheading* | *fetus* | *abortion-related* | *clinics* | *abortionist* | *infanticide* |
| *doublespeak* | *pregnancy* | *infanticide* | *roe v. wade* | *gruesome* | *pro-life* |
| *theatrics* | *abortion-related* | *abortionists* | *infanticide* | *procedure* | *gosnell* |
| *soviet-backed* | *fetal* | *rape* | *gruesome* | *infanticide* | *elective* |
| *annan* | *alcoholic* | *hyde amendment* | *abortionist* | *roe v. wade* | *dismembering* |

## References


M. Antoniak and D. Mimno. Evaluating the stability of embedding-based word similarities. *Transactions of the Association of Computational Linguistics*, 6:107–119, 2018.

H. Dubossarsky, D. Weinshall, and E. Grossman. Verbs change more than nouns: a bottom-up computational approach to semantic change. *Lingue e linguaggio*, 15(1):7–28, 2016.

H. Dubossarsky, S. Hengchen, N. Tahmasebi, and D. Schlechtweg. Time-out: Temporal referencing for robust modeling of lexical semantic change. *arXiv preprint arXiv:1906.01688*, 2019.

W. L. Hamilton, J. Leskovec, and D. Jurafsky. Diachronic word embeddings reveal statistical laws of semantic change. *arXiv preprint arXiv:1605.09096*, 2016.

Y. Kim, Y.-I. Chiu, K. Hanaki, D. Hegde, and S. Petrov. Temporal analysis of language through neural language models. *arXiv preprint arXiv:1405.3515*, 2014.

A. Kutuzov, L. Øvrelid, T. Szymanski, and E. Velldal. Diachronic word embeddings and semantic shifts: a survey. *arXiv preprint arXiv:1806.03537*, 2018.

L. Wendlandt, J. K. Kummerfeld, and R. Mihalcea. Factors influencing the surprising instability of word embeddings. *arXiv preprint arXiv:1804.09692*, 2018.