

# Legal and political stance detection of SCOTUS language

Noah J. Bergam, Emily Allaway, Kathleen McKeown (Natural Language Processing Lab)

Columbia University. Laidlaw Scholars Foundation (2022).

## SC-stance: the first legal stance detection dataset

→ We match written opinion to legal questions from a legal education website  
 → *Task*: figure out whether the opinion rejects or affirms the question.

**Case:** School District of Abington Township v. Schempp (ID 1962-148), Majority Opinion.

**Target:** Did the Pennsylvania law requiring public school students to participate in classroom religious exercises violate the religious freedom of students as protected by the First and Fourteenth Amendments?

**Text:** Once again we are called upon to consider the scope of the provision of the First Amendment to the United States Constitution which declares that "Congress shall make no law respecting an establishment of religion, or prohibiting the free exercise thereof" [...] In light of the history of the First Amendment and of our cases interpreting and applying its requirements, we hold that the practices at issue and the laws requiring them are unconstitutional under the Establishment Clause, as applied to the States through the Fourteenth Amendment. [...]

**Label:** pro (text affirms the target)

Sample data point from SC-stance.

How political are Supreme Court justices?  
**Idea: look at their language!**

we collect:

**Written Opinions**  
 33k documents  
 (1789 to 2020)

**Oral Arguments**  
 3.8m lines of dialogue  
 (1955 to 2020)

We analyze each corpus using **automated stance detection**, e.g.

$D_1 =$ Once the Court starts looking to the currents of public opinion regarding a particular judgment, it enters a truly bottomless pit from which there is simply no extracting itself. (Rehnquist, 1992)	
$D_2 =$ Will this institution survive the stench that this creates in the public perception that the Constitution and its reading are just political acts? (Sotomayor, 2022)	
$T =$ The Supreme Court ought to make decisions with the public opinion in mind.	
<b>stance(<math>D_1, T</math>) = con</b>	<b>stance(<math>D_2, T</math>) = pro</b>

Table 1: A relevant, sophisticated example of stance detection.

We train stance detectors using **machine learning**, i.e. we allow the model to *learn patterns from examples*. Most of our models rely on bidirectional encoder representations from transformers, more popularly known as BERT (Devlin 2018).

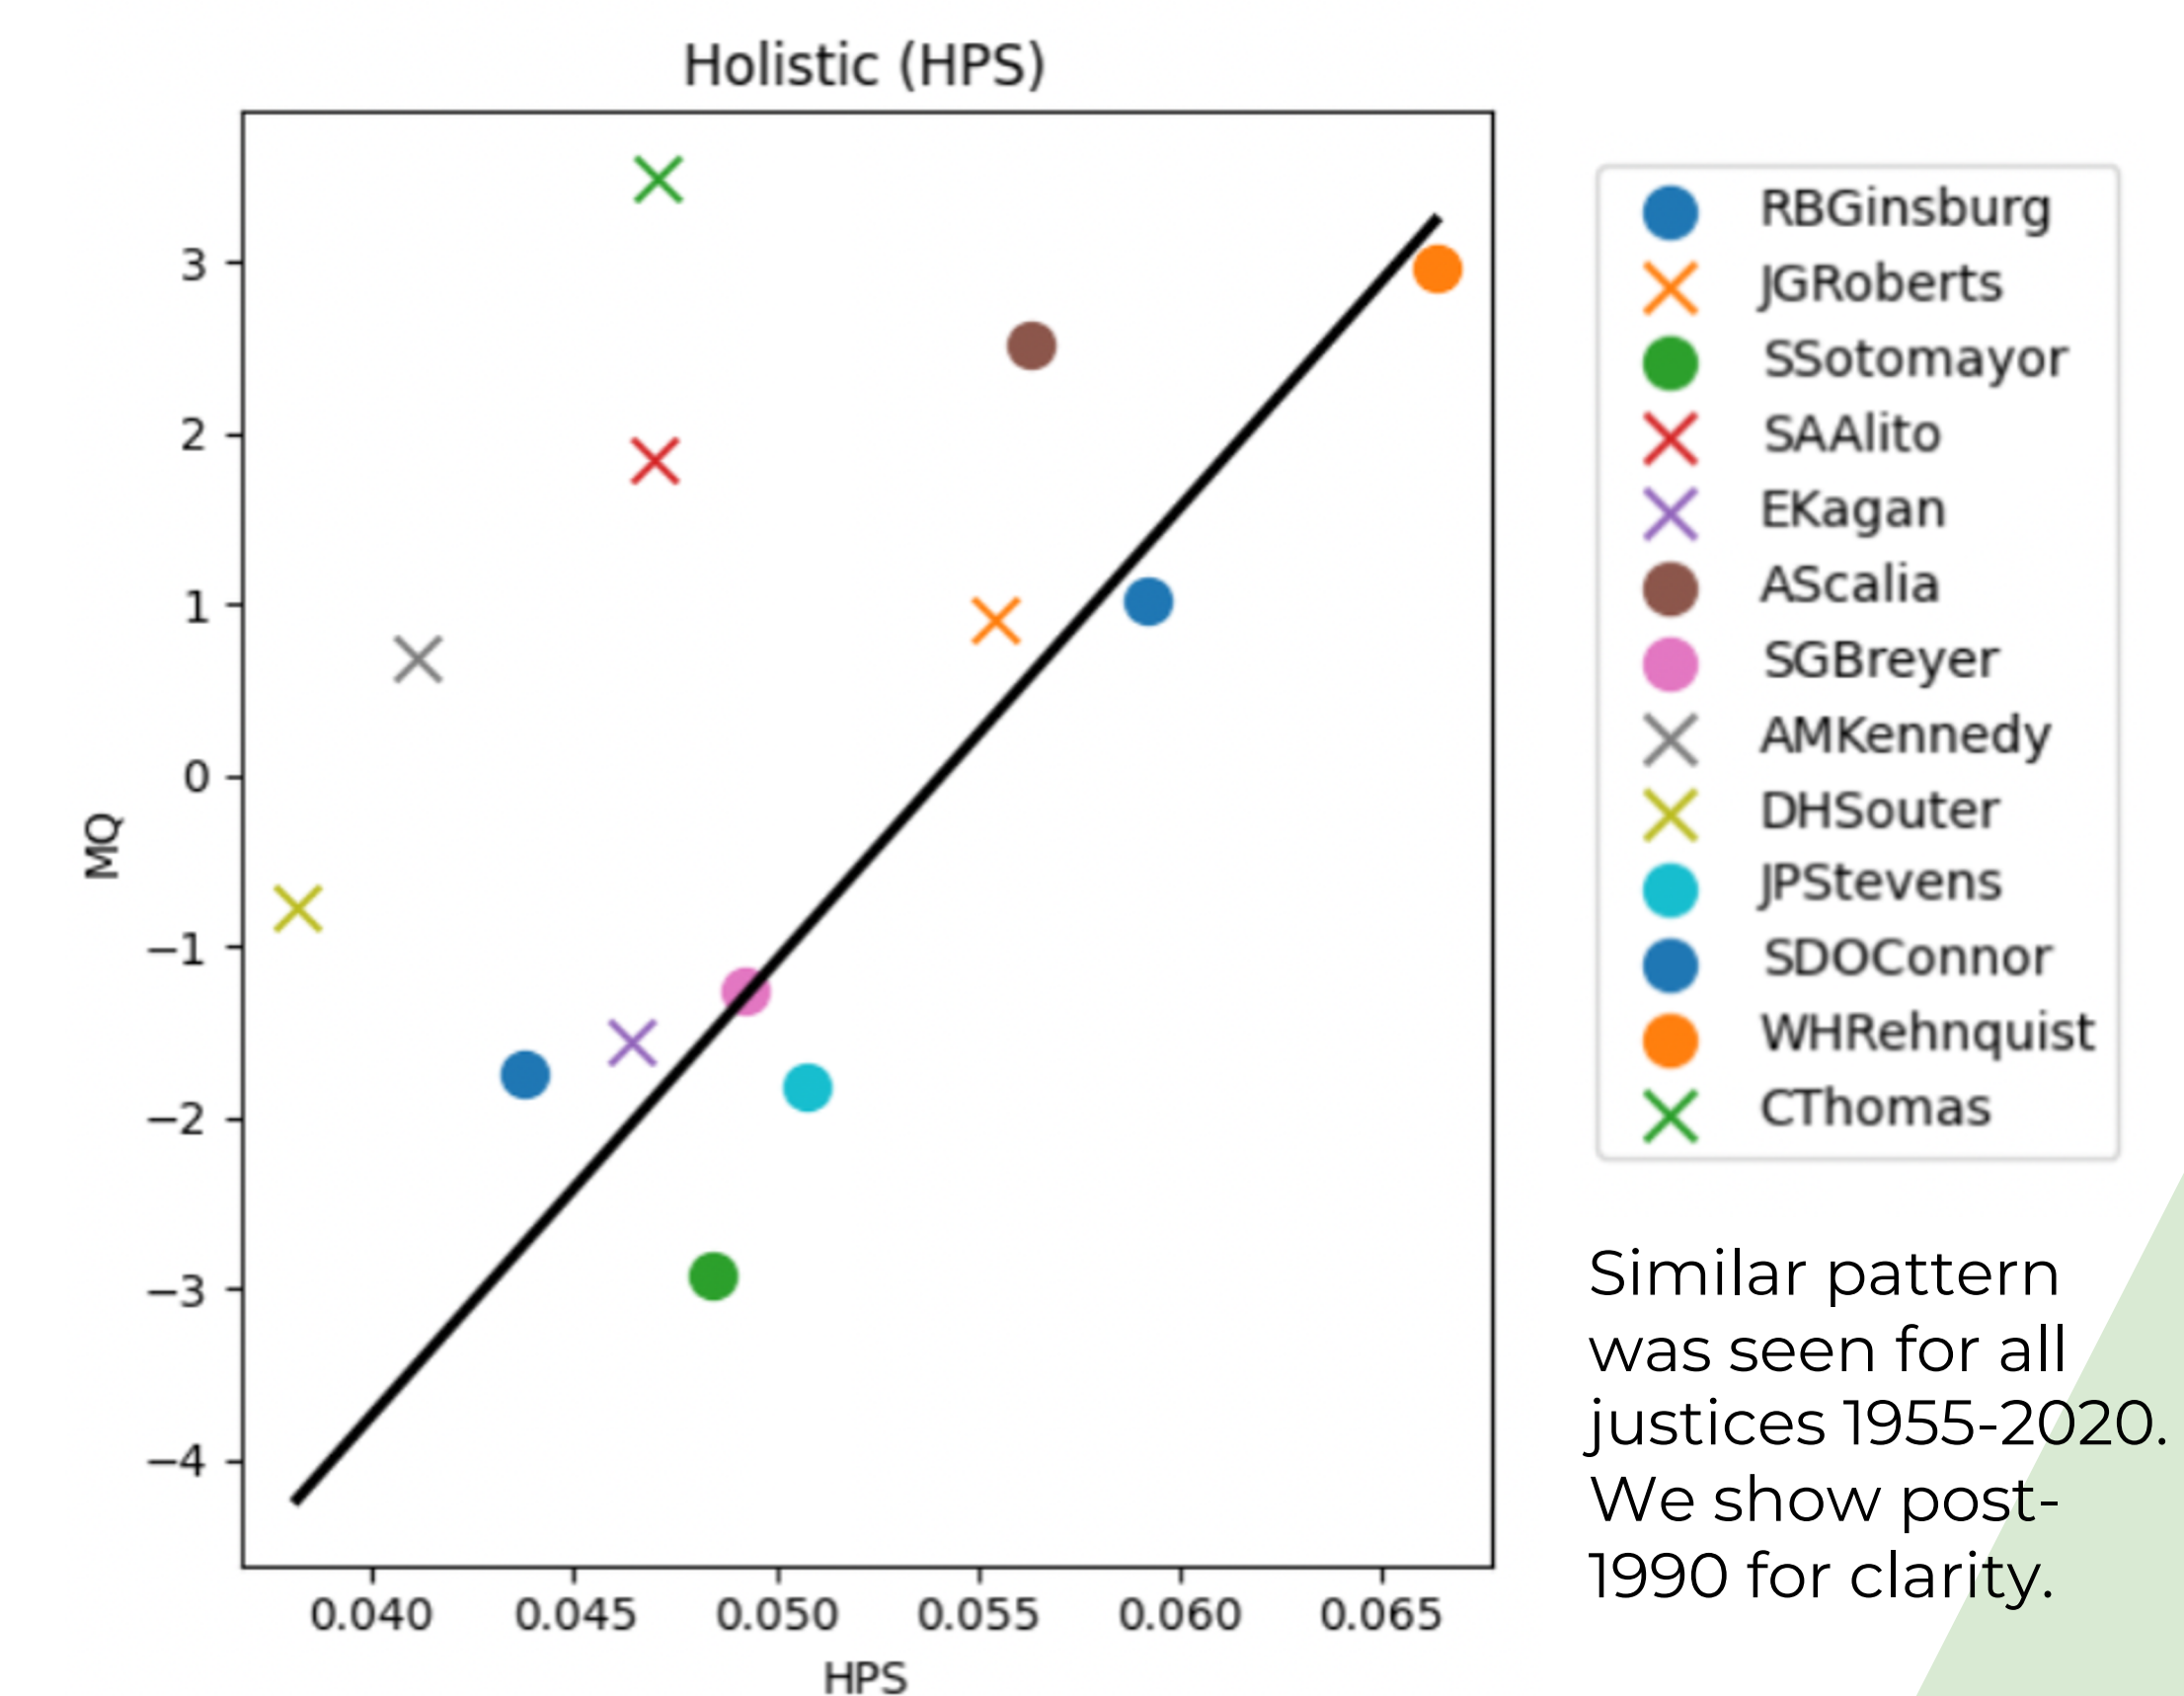
Two perspectives.

**SCOTUS language as a training ground for stance detection.**  
 We build and evaluate a dataset.

**SCOTUS language as the data for quantitative social sciences study.**  
 We track and interpret historical patterns.

## New evidence of the attitudinal change hypothesis

**Experiment:** We plot justices by our stance detection-based ideology score (HPS)\* and the voting-based Martin-Quinn score



Similar pattern was seen for all justices 1955-2020. We show post-1990 for clarity.

Circles denote justices with a statistically significant correlation between their Martin-Quinn score and the Stimson public policy mood (i.e. they are "responsive" to public opinion).

**Justices who are more responsive to public opinion express their politics in their public-facing language!**

## Stance Detection F-1 Scores

Majority	39.6
tf-idf (LR)	41.4
tf-idf (MLP)	50.0
BERT	50.4
CaseLaw-BERT	47.6
Legal-BERT	52.8
Legal Adapter	<b>55.6</b>

"dummy" classifier: predicts most common class

Term & doc frequency vectors

Transformer neural net embedding

BERT w/ legal pre-training


BERT w/ legal adapter plugin

**Our new legal language embedding (BERT+legal adapter) is strong!**

The standard approach for this kind of hard legal language understanding task is to pre-train a language model like BERT on specialized (legal) text. **Our approach involves using a generalist model (BERT) and plugging in a legal adapter.** One advantage of this new method: flexibility and fewer parameters needed to switch domains.



## Linguistic Ideology Metrics



This result can be explained by the **attitudinal change hypothesis** (Casillas 2010): if individual attitudes determine justice behavior, these attitudes would jointly affect voting results *and* language.

**Experiment:** We construct two stance-detection based indicators of political ideology and apply to >60 years of SCOTUS dialogue.

### Holistic Political Stance (HPS)\*

Evaluates how much the justices talk like Republicans/Democrats in Congress. Pre-trained on (Thomas 2006)'s Convote dataset.

### Issue-Specific (ISS)

Evaluates whether justices express liberal or conservative opinions, by topic. Pre-trained on (Allaway and McKeown 2020)'s VAST dataset.

### ISS and HPS correlate! (p < .0005)

⇒ Affirms intuition behind the design of these metrics  
 ⇒ Suggests the existence of an partisan ideological signal in the speech of the (officially apolitical) Supreme Court