OXFORD

Journal of Legal Analysis, 2025, **17**, 2–13

https://doi.org/10.1093/jla/laaf001
Advance access publication 27 March 2025
Article

Lifting the American Supreme Court Veil: Identifying Authorship in Unsigned Opinions

Ronen Avraham[1,*], Rami Nasser[2], Itamar Kohn[3], Tamar Kricheli-Katz[3], and Roded Sharan[2]

[1]School of Law, The University of Texas at Austin, Austin, TX, United States
[2]School of Computer Science and AI, Tel Aviv University, Tel Aviv, Israel
[3]Buchman Faculty of Law, Tel Aviv University, Tel Aviv, Israel

[*]Corresponding author. Ronen Avraham, Buchman Faculty of Law, Tel Aviv University, Israel. E-mail: Ronen3112@gmail.com

Abstract

The Supreme Court of the United States (SCOTUS) issues 10–15 % of its opinions unsigned, concealing authorship. Traditionally, unveiling authors required the posthumous release of Justices' personal papers. We trained our AI algorithm to achieve real-time authorship probabilistic identification, encompassing 17 Justices and 4,069 opinions from 1994 to 2024. Our algorithm identified the likely authors of the March 2024 Trump v. Anderson case, which enabled Donald Trump to run for office. Moreover, our algorithm unveiled the likely authorship in significant unsigned COVID-19 era cases, estimated with high probability individual parts of the joint dissent in the Obamacare Case (2012), and discerned the likely authors of the landmark cases of Bush v. Gore (2000). Applications range from legal research to decoding SCOTUS internal dynamics. Compared to prior methods, our study demonstrates a substantially higher accuracy rate of 91 per cent over a much longer period of time, offering timely insights into the nuances of SCOTUS decision-making. To facilitate further research, we provide a public web server at https://raminass.github.io/SCOTUS_AI/.

1. Introduction

Approximately 10–15 per cent of the official opinions of the Supreme Court of the United States (SCOTUS) remain unsigned, leaving the authorship unknown. Traditionally, the main method for determining the authorship of these opinions involved waiting for the release of personal papers of Justices, typically several years posthumously. This delay is exemplified in the recent unveiling of the personal papers of Justice John Paul Stevens, who retired in 2010 and passed away in 2019. These papers shed light on the deliberations surrounding the 2000 decision in Bush v. Gore,[1] a significant case where the majority opinion granting George W. Bush the presidency over Al Gore, was unsigned. However, such documents face extended public exposure delays—more than 22 years in the case of Justice John Paul Stevens. Furthermore, they are considered private property, and Justices are not obligated to share or publicize them after retiring. In fact, only about 1 in 3 Justices have chosen to donate their papers to the Library of Congress (Gresko 2023).

Few studies have explored SCOTUS opinion authorship through classical statistical and machine learning methods. Rosenthal and Yoon (2011a, 2011b) utilized function words with a linear or naïve Bayes classifier to distinguish between only 2 justices. Chandler, Muenster, and Lichtblau (2023) focused

[1] Bush v. Gore, 531 U.S. 98 (2000). Ronen Avraham, Buchman Faculty of Law, Tel Aviv University, Israel. This work was supported by a grant from the Tel Aviv University Center for AI and Data Science (TAD).

on a single case—**Martin v. Hunter's Lessee** (1816)—employing various machine learning methods to attribute the opinion to Chief Justice John Marshall or Joseph Story. Li et al. (William et al. 2013) addressed author attribution for 12 justices over a period of 8 years, introducing a probabilistic model based on single, pairs, or triplets of consecutive words (called *n*-grams). In this study, we leverage the latest developments in large language models (LLMs), presenting the first LLM-based deep learning approach for author attribution in unsigned SCOTUS opinions, encompassing 17 Justices over 30 years, from 1994 to 2024. This methodology enhances historical insights into the American Supreme Court's inner workings.

Covering 17 Justices and 196 unsigned opinions in the period 1994–2024, our approach identifies likely authorship in numerous unsigned landmark cases, including estimating the authors in the recent case of **Trump v. Anderson**,[2] where the Supreme Court determined Donald Trump can run for office, the authors of the majority opinion in **Bush v. Gore**, where the Supreme Court propelled George W. Bush to the oval office, as well as the authors in the 2012 dissent in **National Federation of Independent Business v. Sebelius**,[3] commonly known as the Obamacare Case, which upheld the constitutionality of the Affordable Care Act (ACA). We unveiled the authorship of two recent Covid-era cases.

Importantly, our algorithm can estimate likely authorship not only in full opinions but also in specific portions, which proves valuable when dealing with opinions collaboratively written by multiple justices, such as the well-known 1992 abortion case of **Planned Parenthood v. Casey**.[4]

We introduce a user-friendly website, enabling researchers and the public to unveil authorship in any unsigned opinion or its components from those years.

1 . A BRIEF HISTORY OF UNSIGNED OPINIONS

The recent March 2024 Supreme Court opinion in **Trump v. Anderson**, which determined that former President Trump can run for office, has been delivered in a per curiam fashion. While this occurrence is not particularly common, it is not unusual either. Typically, most SCOTUS opinions are authored by individual Justices, with other Justices having the option to join these opinions. However, over the past hundred years or so, a notable phenomenon has emerged: unsigned opinions.

There are two primary types of unsigned opinions: Per Curiam and Joint Opinion. A per curiam opinion is an unsigned opinion issued in the name of the entire Court, usually written by a single Justice whose identity is not revealed. It is often employed for unanimous decisions or when the Court aims to present a unified voice.

Initially reserved for straightforward cases lacking legal complexity, during the twentieth century, per curiam opinions began addressing matters of significant legal and cultural import. This evolution included cases such as the defense of press freedom (**New York Times Co. v. United States**, 1971)[5] or the desegregation in public housing (**New Orleans City Park Improvement Ass'n v. Detiege,** 1958).[6] Over time, their use increased, constituting approximately 12 per cent of cases between 2021 and 2023.

The use of unsigned per curiam opinions also helped the Court issue decisions quickly and efficiently on urgent issues, such as the 2020 COVID-19 related case of **Roman Catholic Diocese of Brooklyn v. Cuomo**,[7] where the Supreme Court ruled that New York may not restrict religious gathering in Catholic churches and Orthodox synagogues because the restrictions likely discriminate against religion in violation of the First Amendment, or the 2021 COVID-19 case of **Tanden v. Newsom**,[8] where the Supreme Court ruled that California may not bar the meeting of more than three families to worship in a private home because they violate the First Amendment free exercise clause.

In addition, the Court employed these opinions in deciding controversial cases, where anonymity could shield Justices from blame or accusations of partisanship and display unanimity (Robbins 2012), as seen in multiple cases dealing with segregation following **Brown v. Board of Education**.[9]

[2] Trump v. Anderson, 601 U.S. (2024).
[3] National Federation of Independent Business v. Sebelius, 567 U.S. 519 (2012).
[4] Planned Parenthood v. Casey, 505 U.S. 833 (1992).
[5] New York Times Co. v. United States, 403 U.S. 713 (1971).
[6] New Orleans Park Improvement Ass'n. v. Detiege, 358 U.S. 54 (1958).
[7] Roman Catholic Diocese of Brooklyn v. Cuomo, 592 U.S (2020).
[8] Tanden v. Newsom, 593 U.S. 61 (2021).
[9] Mayor and City Council of Baltimore City v. Dawson, 350 U.S. 877 (1955) (desegregation of public beaches); Holmes v. Atlanta, 350 U.S. 879 (1955) (desegregation of public golf courses); Johnson v. Virginia, 373 U.S. 61 (1963) (segregated courtrooms are unconstitutional); Holmes County Board of Education, 396 U.S. 19 (1969) (desegregation of public schools in the South).

Sometimes, however, per curiam opinions serve not as a reflection of the entire court but as a unified voice of the majority, as was the case in **Bush v. Gore**, or to express the joint perspectives of dissenting Justices, as seen in the Obamacare Case. We will unveil the likely authors of these two opinions later.

The other type of unsigned opinion is a joint opinion (Ray 2012). Joint opinions are rarer than per curiam opinions, and unlike the latter, the authors of joint opinions are explicitly identified, reflecting the shared perspectives of those Justices only. However, the specific assignment of parts within the joint opinion to individual Justices often remains a mystery.

One of the most renowned instances of joint opinions occurred in the majority opinion in the famous 1992 abortion case of **Planned Parenthood v. Casey**, authored by Justices O'Connor, Kennedy, and Souter and joined by Justices Blackmun and Stevens. As we show, our algorithm yielded unexpected predictions regarding the distribution of authorships among the justices in this notable case.

2. MATERIALS AND METHODS

2.1 Data Retrieval and Preprocessing

We used the CourtListener website[10] to download the text of Supreme Court opinions. The opinions were downloaded with full text that includes citations and other legal elements (syllabus, bibliography, etc.). We focused on opinions that were reported between September 1994 and December 2023 written by a total of 17 Justices (Supplementary Table S1). We included all types of opinions (majority, dissenting, and concurring) for training and testing (Supplementary Table S2). We preprocessed the opinion texts to clean citations from them using the Eyecite tool (Cushman, Dahl, and Lissner 2021).

2.2 Modeling and Training

Our classification model was based on a pre-trained language model, legal-BERT (Chalkidis et al. 2020) that employs the large-scale BERT language model for text representation and fine-tuned to extract informative semantic and stylistic features from legal documents. We use the complete BERT architecture, which consists of six transformer blocks, each of which includes a multi-head self-attention mechanism, layer normalization, a feed-forward neural network, and residual connections.

As the model is limited to input sequences of length up to 512 words, we split every opinion to chunks of up to 420 words with 10-word overlap between chunks. Chunk statistics appear in Supplementary Table S3.

Legal-BERT is pre-trained to produce an embedding (CLS token) that represents each input sequence. To transform this embedding into a probability, we fine-tuned legal-BERT with an additional classification head that includes a pooling layer to aggregate the token-level outputs and a dense layer followed by a softmax activation for final classification. The fine-tuning process involves the entire model, including all transformer blocks. This allows the model to adapt its internal representations to our specific task of opinion authorship attribution. In order to address imbalances between justices, we used a resampling approach to obtain 2,000 chunks for each Justice.

At the test stage, the inferred chunk probabilities were averaged to obtain predictions also at the opinion level. Specifically, the scores of each Justice across the different chunks were averaged to produce an opinion score. These latter scores were then normalized to sum to 1. To constrain the predictions to a pre-specified subset of Justices, we normalized only the scores of the Justices included in this subset. Note that the resulting scores, while normalized to be between 0 and 1, should not be interpreted as probabilities. In fact, modern neural networks are known to produce poorly calibrated outputs (Chuan et al. 2017). Thus, we evaluate below the accuracy implied by these scores using a cross-validation technique.

2.3 Model Evaluation

We evaluated each candidate model using five-fold cross validation. In each iteration of the cross validation, we hid 20 per cent of the opinions, trained the model on the chunks of the other 80 per cent of the opinions, and tested the accuracy of the trained model on the chunks of the hidden opinions. Through this process, we ensured that the test chunks were derived from different opinions than those used for the training chunks.

[10] https://www.courtlistener.com/.

2.4 Model Selection

The reported model initially targets 13 Justices with a record of authoring more than 100 signed opinions each. To include in the model also the four most recent SCOTUS Justices: Gorsuch, Kavanaugh, Barrett, and Jackson, we added up to 100 of their most recent opinions while still serving in lower courts. We trained a total of ten models starting from different random initializations to mitigate the risk of being trapped in a local optimum. We evaluated each of the models in five-fold cross validation, saving those (six in total) that attained an accuracy of at least 90 per cent. We then retrained those models on the full data using the same starting points and measured the distance between pairs of models according to the sum of KL-divergences between their predictions on per-curiam cases in the 1994–2023 period. Finally, we selected the model that minimized the sum of distances to all other models (regardless of its classification performance). The selected model demonstrates an accuracy of 91 per cent in five-fold cross-validation and its detailed results are reported in Supplementary Table S4.

3. Results

We developed a large language-based model for authorship attribution of SCOTUS opinions (Fig. 1). The model receives as input an opinion, the year it was issued and optionally a subset of Justices that should be excluded from the prediction. The opinion is split into chunks of up to 420 words (see Section 2) and authorship is predicted for each chunk. These predictions are aggregated to form opinion-level predictions. The model outputs author predictions with confidence scores (summing to 100 per cent). We utilized our model to identify the authorship of 17 Justices during the period 1994–2024. This selection was based on having a minimum of 100 signed opinions suitable for training and validation, which includes all 17 justices.

We trained the model for chunk-level prediction using all types of opinions for which the author was uniquely identified in the 1994–2023 period (see statistics for these opinions and their chinks in Supplementary Tables S1 and S3). In five-fold cross validation, the model attained high accuracy of 75 per cent at the chunk level (Supplementary Fig. S1). These translated to very high accuracy of 91 per cent at the opinion level (Fig. 2, Supplementary Fig. S2, and Supplementary Table S4).

One can refine these observations based on the prediction scores obtained in each case. As explained in Section 2, these scores do not correspond to probabilities but can be calibrated based on the cross-validation results. In particular, when we trained the algorithm, we noticed that if the top prediction score is greater than 40 per cent (50 per cent, 60 per cent), our accuracy in predicting the authoring justice increases to 93 per cent (95 per cent, 96 per cent, respectively). Similarly, the original accuracy further improves to 95 per cent when considering the top two predictions per opinion, rather than a single one. If the sum of the top two prediction scores exceeds 50 per cent, the accuracy increases to 98 per cent.

Overall, the per-justice precision in our signed-opinions dataset ranged from 80 per cent (Justice Alito) to perfect 100 per cent (Justices Ginsburg, Kennedy, and Souter) with very high per-justice recall

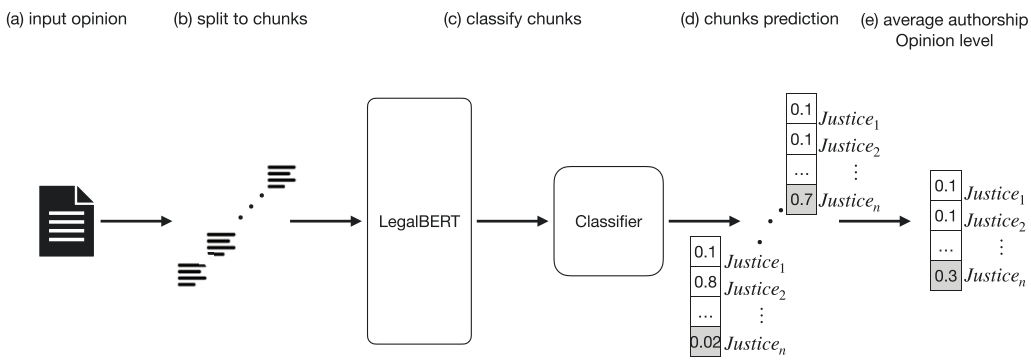

Figure 1. An overview of the authorship prediction model.

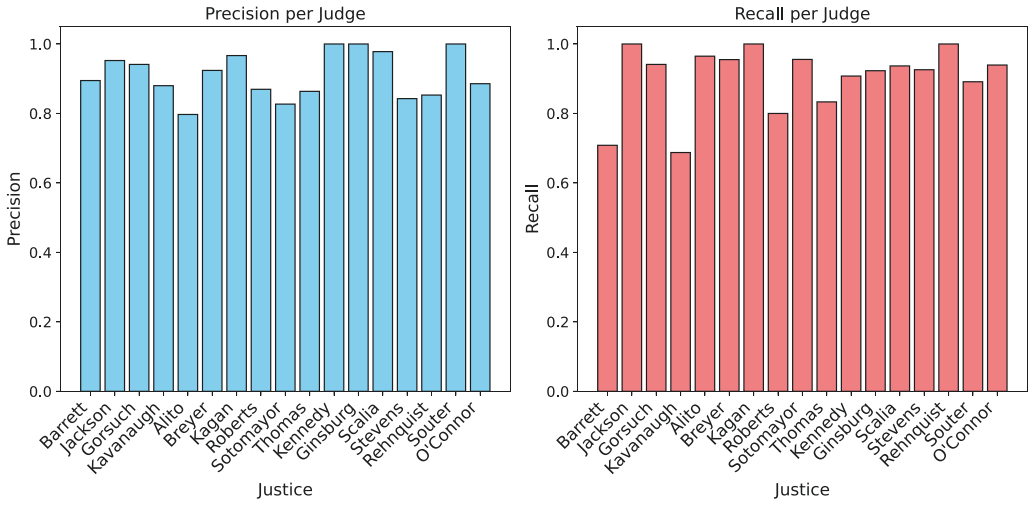

Figure 2. Opinion-level performance (precision & recall) per Justice in 5-fold cross validation.

(representing the percentage of opinions authored by that justice that were correctly predicted) and F1-score (harmonic mean of precision and recall).

Previous work by Li et al. (2013) analyzed unsigned opinions in the period 2005–2011 using *n*-gram features. For comparison purposes, we retrained our algorithm on these years only, spanning nine Justices who have served during the tenure of Chief Justice John Roberts (Alito, Breyer, Ginsburg, Kagan, Kennedy, Roberts, Scalia, Sotomayor, and Thomas). We extracted the same types of opinions (majority and dissenting) as used in Li et al. and used them for training and evaluation in a ten-fold cross-validation setting. The previous study reported an accuracy of 81.2 per cent, while our method achieved a higher accuracy of 84 per cent.

Finally, we compared our results to those of Rosenthal and Yoon (2011a, 2011b), who attempted to distinguish between pairs of Justices from Rehnquist's court (1994–2005). We followed the same evaluation methodology, retraining our algorithm using only majority opinions from that period in a five-fold cross-validation setting. To imitate the binary decisions in Rosenthal and Yoon's method, we focused each time on the probabilities assigned to the two Justices in question only. Despite facing a more challenging multi-label prediction scenario involving nine justices, our accuracy surpasses theirs (Supplementary Table S5), with an average accuracy improvement of 6.2 per cent over all pairs.

3.1 Identifying Authorship in Per Curiam Opinions

We start by predicting the author of **Trump v. Anderson** (2024). We then embark on the task of predicting authorship in two pivotal COVID-19 era cases and two of the most contentious Supreme Court cases in history: **Bush v. Gore** (2000) (1) and the Obamacare Case (8), and two pivotal Covid-19 era cases.

To aid in this endeavor, we have developed a dedicated online tool for SCOTUS authorship prediction.[11] This tool allows the application of our method to any opinion within the years 1994–2024. To enhance accuracy, users can manually exclude justices known not to have participated in the writing, such as those who authored dissenting opinions. To facilitate easy access, we added a table in Supplementary Appendix S8 that lists the top two predictions for all unsigned opinions in our dataset. These are raw predictions before manually eliminating justices known to not have authored the opinions.

The following findings were obtained after manually excluding irrelevant justices—a task that can only be performed on the accompanying website. Detailed results, both before and after manual exclusion of justices, are provided in Tables 1–5.

[11] https://raminass.github.io/SCOTUS_AI/.

3.1.1 *Trump v. Anderson (2024)*

In **Trump v. Anderson** (2024), the US Supreme Court unanimously ruled, in a Per-Curiam opinion, that states cannot determine eligibility for federal office, including the presidency. This case stemmed from the Colorado Supreme Court's rejection of former President Donald Trump's presidential eligibility due to his involvement in the January 6 Capitol attack. The Colorado court argued that Section 3 of the Fourteenth Amendment disqualifies presidential candidates who engage in insurrection against the United States. This decision marked the first time a presidential candidate was disqualified from office in a state based on the Fourteenth Amendment. However, the US Supreme Court reversed this decision through a per curiam ruling on March 4, 2024, asserting that the power to determine eligibility for federal office resides exclusively with the federal government.

There is speculation regarding the authorship of the opinion, with some suggesting Chief Justice Roberts as a potential author given the significance of the case (Kanthor and Liptak 2024). Others suggested Justice Kavanaugh. As far as we know, no one suggested Justice Alito wrote it.

The opinion consists of two main parts. Part I introduces the facts and history of the case at the Colorado Supreme Court. Part IIA discusses the legal framework for deciding the case, while Part IIB applies this framework to the facts of the case.

After manually eliminating the justices who wrote concurring opinions, our algorithm's top two prediction scores are Justices Alito (34 per cent) and Kavanaugh (29 per cent). Since the sum of these scores is larger than 50 per cent, we are 98 per cent sure one of them wrote it. Since the prediction score for each part that Justice Alito wrote is at least 40 per cent, we are 90 per cent sure that Justice Alito wrote both parts (see Table 1). Despite the importance of the case, it appears unlikely that Chief Justice Roberts was involved in the writing process. Notably, despite representing the *entire* court, none of the liberal justices contributed to the writing (see Table 1).

3.1.2 *Bush v. Gore (2000)*

Although just twenty-three years old, **Bush v. Gore** (2000) is one of the most important Supreme Court cases in American legal history. This high-profile 5:4 case decisively resolved the contested presidential election in favor of Governor Bush, sending him to the Oval Office. The case centered on the recount of presidential ballots in Florida, where the results were extremely close, leading to legal disputes and recounts. Unlike a regular per curiam decision, **Bush v. Gore** is a case where only the majority opinion (and not the entire Court's opinion) was delivered as per curiam. The identity of the anonymous author of the Court's per curiam majority opinion has generated some debate. Attention gravitates toward Justices Kennedy and O'Connor, conservative voices who abstained from affiliating with any separate opinions. Scholars point to Justice Kennedy's writing style as a possible indicator (Foley 2007; Schwartz 2002), while others attribute it to Justice O'Connor (Sheppard 2022). Some even imply a collaborative effort between the two (Stone 2020). Importantly, no one predicts Chief Justice Rehnquist participated in the writing.

Our algorithm identified the top three most likely authors of the majority per curiam opinion as follows: Justice Kennedy (prediction score of 57 per cent), Chief Justice Rehnquist (28 per cent), and Justice Scalia (14 per cent) as the most probable authors. When we separate Part I and Part II of the opinion, our algorithm predicts with a prediction score of 100 per cent that Chief Justice Rehnquist

Table 1. Authorship Trump v. Anderson.

	Before elimination	**After elimination (Kagan, Jackson, Sotomayor, and Barret)**
Full opinion	Alito 25%, Kavanaugh 25%, Gorsuch 15%, Barrett 9%, Sotomayor 9%, Roberts 8%, Thomas 7%, Kagan 2%, Jackson 1%	Alito 34%, Kavanaugh 29%, Gorsuch 19%, Roberts 11%, Thomas 7%
Part I	Alito 38%, Kavanaugh 23%, Barrett 11%, Thomas 11% Roberts 7%, Gorsuch 6%, Sotomayor 3%, Kagan 1%, Jackson 1%	Alito 40%, Kavanaugh 30%, Thomas 12%, Roberts 10%, Gorsuch 8%
Part II	Alito 47%, Kavanaugh 28%, Sotomayor 11%, Thomas 6%, Barrett 4%, Gorsuch 2%, Roberts 2%, Jackson 1%	Alito 41%, Kavanaugh 39%, Gorsuch 12%, Thomas 6%, Roberts 2%

Table 2. Authorship Bush v. Gore.

	Before elimination	After elimination (Breyer, Ginsburg, Souter, Stevens)
Full opinion	Kennedy 53%, Rehnquist 24%, Souter 14%, Scalia 5%, Stevens 3%	Kennedy 57%, Rehnquist 28%, Scalia 14%, O'Connor 1%
Part I	Rehnquist 100%	Rehnquist 100%
Part II	Kennedy 68%, Scalia 15%, Rehnquist 15%, Stevens 1%, Souter 1%	Kennedy 70%, Scalia 15%, Rehnquist 15%

wrote Part I, which addresses the background and facts of the case. More interestingly, our algorithm predicts (with a prediction score of 70 per cent, which translates to more than 96 per cent accuracy) that Justice Kennedy wrote Part II, which focuses on the legal analysis and reasoning applied by the Court in reaching its decision (see Table 2).

Notably, the algorithm did not recognize Justice O'Connor as the author of any of the parts of the opinion, assigning her a score of 0 per cent (both before and after manual elimination).

3.1.3 *The Obamacare Case (2012)*

National Federation of Independent Business v. Sebelius (2012) was a pivotal legal challenge to the ACA, a comprehensive health care reform law enacted in the USA in March 2010. The primary goals of the ACA were to increase the quality and affordability of health insurance. **Sebelius** focused on the constitutionality of the ACA's individual mandate, requiring most Americans to obtain health insurance under penalty, and the expansion of Medicaid to cover more low-income individuals. In a 5–4 decision, the Supreme Court upheld the individual mandate as a valid exercise of Congress's taxing power while limiting the Medicaid expansion. This ruling had far-reaching legal and policy implications, confirming federal authority in regulating commerce and influencing the trajectory of healthcare policy in the USA.

Chief Justice Roberts, a conservative, joined the four liberal justices in upholding the law. However, the case also featured a joint dissenting opinion signed by Justices Kennedy, Scalia, Thomas, and Alito. The authorship of this dissent, or its specific parts, remains a mystery.

There are three theories concerning the undisclosed authorship of the dissent. The most intriguing one suggests Chief Justice Roberts as the writer, proposing that he transitioned to the majority side only after completing what has later become the dissent opinion. According to this theory, the dissent was initially intended to function as a majority opinion, declaring Obamacare unconstitutional (Campos 2023; Long and Christensen 2013; Roy 2023). Under this theory, Chief Justice Roberts wrote both the dissent and the majority opinions. The second theory envisages a more prominent role for Justices Kennedy and Scalia, positioning them as principal authors, possibly utilizing an initial template provided by Chief Justice Roberts before he switched sides (Campos 2023; Roy 2023). Another claim is that Justice Scalia was possibly the sole author of the dissent (Long and Christensen 2013). Importantly, no one predicts Justice Alito participated in the writing.

Focusing on the four Justices who signed the dissent, results show a prediction score of 41 per cent for Justice Scalia, 33 per cent for Justice Alito, 20 per cent for Justice Kennedy, 4 per cent for Justice Thomas, and only 1 per cent for Chief Justice Roberts. Since it appears there is a likelihood of a collective effort between Scalia and Alito, possibly involving Kennedy too, we explored authorship for various sections. Our algorithm predicts that Part I on the Individual Mandate was most likely authored by Justice Scalia. Since the prediction score is larger than 40 per cent, we are more than 93 per cent convinced he wrote it. Part II on the Taxing Power was most likely authored by Justice Scalia (top prediction score is larger than 60 per cent and, therefore, we are more than 96 per cent sure about it). Part III on the Anti-Injunction Act was solely written (with a 97 per cent prediction score) by Justice Scalia. Part IV on Medicaid Expansion was probably solely authored by Justice Alito (with a prediction score of 74 per cent which translates to more than 96 per cent chance he wrote it). Part V on Severability is divided into the Act's Major Provisions, which was written by Justice Kennedy (prediction score of 85 per cent), and the Act's Minor Provisions, which was most probably authored by Justices Scalia (prediction score of 48 per cent which translate to more than 93 per cent chance he wrote it) (see Table 3).

Table 3. Authorship Sebelius ("Obamacare").

	Before elimination	After elimination (Breyer, Sotomayor, Kagan, Ginsburg)
Full opinion	Scalia 39%, Alito 30% Kennedy 18%, Breyer 5%, Thomas 3%, Kagan 2%, Roberts 1%, Sotomayor 1%	Scalia 41%, Alito 33%, Kennedy 20 %, Thomas 4%, Roberts 1%
Part I: The Individual Mandate	Scalia 44%, Alito 37% Kennedy 12%, Breyer 7%	Scalia 47%, Alito 40%, Kennedy 12%, Thomas 1%
Part II: The Taxing Power	Scalia 63%, Alito 29%, Breyer 8%	Scalia 68%, Alito 31%, Thomas 1 %
Part III: The Anti-Injunction Act	Scalia 97%, Alito 3%	Scalia 97%, Alito 3 %
Part IV: The Medicaid Expansion	Alito 69%, Thomas 13%, Scalia 7%, Kennedy 4%, Breyer 3%, Roberts 2%, Sotomayor 1%	Alito 74%, Thomas 14%, Scalia 7%, Kennedy 5%
Part V: Severability (Until "The Act's Minor Provisions")	Kennedy 65%, Sotomayor 19%, Breyer 6%, Alito 4%, Scalia 4%, Roberts 1%, Kagan 1%	Kennedy 85%, Alito 6%, Scalia 6%, Roberts 2%, Thomas 1%
The Act's Minor Provisions	Scalia 46%, Breyer 29%, Alito 12%, Kagan 4 %, Sotomayor 3%, Roberts 2%, Thomas 2%, Kennedy 1%	Scalia 48%, Alito 38%, Roberts 6%, Thomas 5%, Kennedy 4%

As it comes out, **Sebelius** is not really a per-curiam opinion; instead, it is a joint opinion co-authored by three of the four justices who signed it (Justice Thomas did not participate in the writing). Most notably, against a once prevailing theory, we can almost definitively conclude that Chief Justice Roberts did not participate in the writing of the dissenting opinion.

3.1.4 The Covid-19 Cases

During the Covid-19 pandemic, numerous states implemented restrictive measures, including those affecting religious gatherings. Two specific cases underscored the tension between public health considerations and the religious liberties safeguarded by the Free Exercise Clause of the First Amendment. In both instances, plaintiffs contended that the state's actions violated the Clause by treating religious activities unfairly, compared to secular ones. These cases were ultimately adjudicated by the Supreme Court, resulting in controversial per curiam rulings.

3.1.5 Roman Catholic Diocese of Brooklyn v. Cuomo (2020)

In **Roman Catholic Diocese of Brooklyn v. Cuomo** (2020), the Supreme Court of the United States (SCOTUS), responding within only two weeks of the injunctive relief application, issued a per curiam majority opinion restraining the application of New York Gov. Cuomo's executive order limiting attendance in churches and synagogues. To this date, the author of the 5:4 opinion has not been publicly revealed. Justices Gorsuch and Kavanaugh issued a concurring opinion. Justices Breyer, Kagan, and Sotomayor dissented. Interestingly, the results do not change much when we remove the dissenting justices (as well as Justice Ginsburg who passed away). Justice Alito remained the most likely author with a 78 per cent prediction score which translates to more than 96 per cent accuracy (see Table 4).

3.1.6 Tandon v. Newsome (2021)

In **Tandon v. Newsom** (2021), California's Governor Newsom issued restrictions barring the meeting of more than three families to worship in private homes. The plaintiff argued that larger groups were allowed to gather in some commercial settings under the executive order, questioning the fairness of limiting attendance in churches and synagogues. The Supreme Court of the United States (SCOTUS), responding within a week of the injunctive relief application, issued a per curiam majority opinion overturning a 9th Circuit decision that upheld California's COVID-19 restrictions, deeming them a violation of the First Amendment free exercise clause. Justice Kagan dissented, joined by Justice Breyer and Justice Sotomayor. Chief Justice Robers was reported to also deny the application.

Table 4. Authorship in Roman Catholic Dioceses of Brooklyin v. Cuomo.

	Before elimination	After elimination (Breyer, Ginsburg, Kagan, Sotomayor)
Full opinion	Alito 76%, Barrett 16%, Gorsuch 3%, Breyer 1%, Roberts 1%, Thomas 1%, Sotomayor 1%, Kavanaugh 1%	Alito 78%, Barrett 16%, Gorsuch 3%

Table 5. Authorship Tundon v. Newsome.

	Before elimination	After elimination (Breyer, Kagan and Sotomayor)
Full opinion	Sotomayor 32%, Thomas 20%, Barrett 18%, Alito 11%, Kavanaugh 9%, Gorsuch 4%, Roberts 4%, Kagan 1%	Barrett 29%, Thomas 23%, Kavanaugh 20%, Alito 16%, Gorsuch 6%, Roberts 6%

After removing the dissenting Justices, Justice Barrett emerged as the most likely author, with a prediction score of only 29 per cent, followed by Justice Thomas at 23 per cent, and Kavanaugh at 20 per cent. Since the sum of the top two prediction scores is larger than 50 per cent, we are 98 per cent sure that Justice Barret or Thomas wrote the opinion. It is also possible that the short length of the opinion (3.5 pages) prevented our algorithm from accurately predicting the authors.

4. SUMMARY

We develop an automated scheme for authorship probabilistic prediction of unsigned SCOTUS opinions. Our work introduces several novel contributions with respect to prior work: (i) We implemented state-of-the-art NLP models using AI, surpassing previous work that focused on specific and short time periods or individual cases, utilizing traditional (non-deep) feature-based approaches; (ii) our algorithm enables real-time author identification with a record high baseline accuracy of 91 per cent and over a much longer period than ever before; (iii) our algorithm allows for the probabilistic prediction of multiple authorships in various portions of a single opinion written collaboratively, sometimes referred to as "joint opinions." In particular, we discovered that the dissent in **Sebelius**, a lawsuit that challenges Obamacare, is such a case. In Supplementary Materials, we delineate joint authorship in two other pivotal joint opinion cases: **Planned Parenthood v. Casey** (1992), a landmark abortion case, and **McConnell v. Federal Election Commission** (2003),[12] a significant campaign reform case; (iv) we implement person–machine collaboration by allowing for the manual elimination of justices whom the users know or suspect could not have written the unsigned opinions, thus further improving our accuracy well beyond the 91 per cent baseline; and (v) We provide a public website to facilitate research on the hidden world of SCOTUS's unsigned opinions.

Using our algorithm, we predicted the likely authors of the majority unsigned opinion in **Bush v. Gore** (2000) as Justice Kennedy. We refuted the persistent theory as if Chief Justice Roberts wrote both the dissent and majority opinion in the Obamacare Case (2012), and identified the likely authors of the joint dissent as a collective effort, led by Justice Scalia, with some parts written by Justices Alito and Kennedy. We also revealed the likely authors of two COVID-19 era cases. In **Roman Catholic Diocese of Brooklyn v. Cuomo** (2020), the author is Justice Alito. In **Tandon v Newsome**, a very short opinion, we predict that Justice Barret or Thomas likely wrote it. Perhaps most interestingly to some, our algorithm revealed that the author of **Trump v. Anderson** is probably Justice Alito and not Chief Justice Roberts.

Our research is important for lawyers and litigators, who, for the first time ever, can receive this information for still-living and serving Justices, can leverage it when preparing cases, tailoring arguments to align with these specific Justices' perspectives.

Our research is also relevant for various academic fields. Specifically, revealing the authorship of unsigned opinions in real-time not only expedites historical insights for legal historians but also provides a valuable resource for researchers in diverse fields, including law, political science, and sociology. This

[12] McConnell v. FEC, 540 U.S. 93 (2003).

process enables the analysis of court dynamics and the development of legal doctrines, understanding the influence of individual Justices, assessing the effects of justices' race, gender, and experience, and tracing the evolution of judicial thought and group thinking.

For example, future research can use our algorithm to help uncover interesting dynamics in the court, such as whether the justice "closest" to the opposing group (from an ideological perspective) authored the opinion or whether, in contrast, it was one of the more extreme voices.

A related point is whether unsigned opinions are more frequently authored by the median justice, who often casts the pivotal vote in closely contested cases. The median justice, ideologically positioned at the center of the court, plays a crucial role in forming majorities and shaping decisions (Maltzman, Spriggs, and Wahlbeck 2000; Martin, Quinn, and Epstein 2005; Rhode and Spaeth 1976). Our algorithm can determine if the median justice also plays a pivotal role in crafting unsigned opinions, thereby revealing compromise dynamics among the justices. This remains relevant even though the court has not consistently been split along a "five against four" dynamic in recent years, as our research extends back to 1994.

One way to assess justices' political ideologies is by examining their Martin-Quinn scores (Martin and Quinn 2002). Developed by political scientists Andrew D. Martin and Kevin M. Quinn, these scores provide a dynamic measure of the ideological positions of U.S. Supreme Court justices based on their voting behavior. Using our algorithm, we can determine whether the authors of unsigned majority opinions, for example, align more closely with the dissenters compared to other majority justices.

In addition, our algorithm enables exploration of whether justices write differently and express different opinions when their opinions are signed versus unsigned. This inquiry addresses whether justices express more or less moderate views in anonymous opinions compared to signed ones, believing that unsigned opinions would not be easily attributed to them. To investigate this, one would need to develop a score at the opinion level that, unlike the Martin-Quinn score, considers the content of the decision and not just whether a ruling is affirmed or reversed (Farnsworth 2007).

Finally, our algorithm can help explore the reliance of justices on clerks. Using our tool, researchers can examine when and in which cases justices rely more on clerks, and whether this reliance decreases or increases with tenure on the Court. President Biden's recent tenure limit reform can more easily be justified if, as many suspect, Justices' reliance on clerks increases significantly as they age.

As AI becomes an integral part of our lives, it prompts other intriguing questions: Will the number of unsigned opinions decrease because of the ability to identify authorship? Will the potential for reputational damage encourage justices to invest more time and thought in enhancing the quality of their unsigned opinions, or perhaps in masking their identity using AI applications such as ChatGPT? If that happens, our ability to probabilistically identify authors of future unsigned opinions will be hampered. One thing is clear—the future holds a captivating exploration of how justices adapt to this technological shift.

Our design has some limitations, with the most notable one being that our algorithm was trained on signed individual opinions of Justices but was later used to predict authorship in unsigned opinions. If some Justices tend to write differently when composing a signed opinion compared to a joint unsigned one, our ability to identify them as the authors of the unsigned one could be compromised. A related point is the involvement of clerks who assist Justices in drafting their opinions. At first blush, this factor might seem to impact the accuracy of the model in predicting authorship by justices who rely more on clerks. However, if the justices' individual tendency to rely on clerks for drafting opinions is similar in both signed and unsigned opinions, our ability to discern authorship for both types of opinions will already be reflected in our justice precision scores.

Another challenge we encountered was the limited number of opinions written by the four newest Justices in the Court (Gorsuch, Kavanaugh, Barret, and Jackson). To overcome this obstacle, we included up to 100 of their most recent opinions from lower courts. This approach arguably provides a better reflection of their writing style, especially considering their relatively limited clerk assistance compared to their tenure in the Supreme Court. This adjustment improved the accuracy of the model.

Notably, the revelation of authorship comes with its own set of ethical implications. In the case of per curiam opinions, which represent the collective decision of the Court, some argue that the focus should be on the Court as a unified body rather than individual Justices. The tradition of sometimes issuing unsigned opinions often aims to underscore the unity of the Court, and revealing authors may divert attention to individual perspectives, potentially diluting the intended message of a collective decision. Our algorithm that uncovers likely authors of unsigned opinions could, according to this line of thought, potentially affect the perceived legitimacy of the court and its decisions. Indeed, a 2019 law in

France criminalizes data scientists who make predictions about individual judges' rulings (Livermore and Rockmore 2019).

We start by noting that our research is not the first to attempt to uncover the authors of per curiam opinions. This field of scholarly research is well established. What is new are the technological tools we employ to improve the accuracy of these crucial answers.

Next, we note that the tradition in France differs from common law jurisdictions like the USA. In France, opinions have historically never been signed and were delivered in the name of the entire court. In contrast, in the USA, the default is signing opinions.

Finally, we note that the legitimacy of the court can actually be improved by our research. Using our algorithm, future studies can shed light on court dynamics that have often remained hidden from public view. By identifying the likely individual justices behind anonymous decisions and sections, we can analyze their personal and professional backgrounds, ideological leanings, and political beliefs. This level of insight promotes accountability among justices. The court, along with its leadership, shapes the fates of individuals and nations. With such immense power comes an equally immense responsibility.

Be that as it may, the integration of AI technology to reveal the authors of unsigned Supreme Court opinions is now a reality, rendering the debate on the ethical implications of knowing authorship a theoretical one.

SUPPLEMENTARY DATA

Supplementary data is available at *JLA* online.

DATA AVAILABILITY

The dataset used in this study is publicly available on Zenodo at https://doi.org/10.5281/zenodo.14559402. The code for processing the data and training the models is available on Zenodo at https://doi.org/10.5281/zenodo.14559370. These resources provide full access to the data and code employed in this research, enabling reproducibility and further exploration by interested parties.

REFERENCES

Campos, P. (2023). *Roberts Wrote Both Obamacare Opinions*. Salon. http://www.salon.com/2012/07/03/roberts_wrote_both_obamacare_opinions/. Accessed Jan. 26, 2025.

Chalkidis, I, *et al*. (2020). *LEGAL-BERT: The Muppets Straight Out of Law School*. Ass'n for Computational Linguistics 2898. https://aclanthology.org/2020.findings-emnlp.261/

Chandler, S. J., Muenster, E. A., and Lichtblau, D. (2023). 'Who Really Wrote Martin v. Hunter's Lessee? Applying a Machine Learning Ensemble to Uncover a Chief Justice's Unethical Behavior', in Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law, 62–70. New York, NY: Association for Computing Machinery. https://dl.acm.org/doi/proceedings/10.1145/3594536

Cushman, J., Dahl, M., and Lissner, M. (2021) 'Eycite: A Tool for Praising Legal Citations', *Open Source Software*, 6: 3617.

Farnsworth, W. (2007). *The Legal Analyst: A Toolkit for Thinking About the Law*. Chicago, IL: The University of Chicago Press.

Foley, E. B. (2007) 'The Future of Bush v. Gore?', *68 Ohio St. L.J.*, 925: 972.

Gresko, J. (2023). *Supreme Court Justice John Paul Stevens' Private Papers Made Open to the Public*. PBS NewsHour. https://www.pbs.org/newshour/politics/supreme-court-justice-john-paul-stevens-private-papers-made-open-to-the-public. Accessed Jan. 26, 2025.

Guo, C., *et al*. (2017). 'On Calibration of Modern Neural Networks'. Proc. 34th Int'l Conf. Mach. Learning 1321–30. https://proceedings.mlr.press/v70/

Kantor, J., and Liptak, A. (2024) 'Sept. 15. How Roberts Shaped Trump's Supreme Court Winning Streak', *The New York Times*, https://www.nytimes.com/2024/09/15/us/justice-roberts-trump-supreme-court.html?unlocked_article_code=1.LE4.y8Ca.DZjE22XQI94t&smid=url-share. Accessed Jan. 26, 2025.

Li, W., *et al*. (2013) 'Using Algorithmic Attribution Techniques to Determine Authorship in Unsigned Judicial Opinions', *Stan. Tech. L. Rev.*, 16: 503.

Livermore, M., and Rockmore, D. (2019). *France Kicks Data Scientists Out of Its Courts*. https://slate.com/technology/2019/06/france-has-banned-judicial-analytics-to-analyze-the-courts.html. Accessed Jan. 26, 2025.

Long, L. N., and Christensen, W. F. (2013) 'When Justices (Subconsciously) Attack: The Theory of Argumentative Threat and the Supreme Court', *91 Or. L. Rev.*, 933: 957–58.

Maltzman, F., Spriggs, J. F. and Wahlbeck, P. J. (2000). *Crafting Law on the Supreme Court: The Collegial Game*. Cambridge: Cambridge University Press.

Martin, A. A., Quinn, K. M., and Epstein, L. (2005) 'The Median Justice on the United States Supreme Court', *N.C. L. Rev.*, 83: 1275.

Martin, A. D., and Quinn, K. M. (2002) 'Dynamic Ideal Point Estimation via Markov Chain Monte Carlo for the U.S. Supreme Court, 1953–1999. 10 Pol', *Analysis*, 134: 145–51.

Ray, L. K. (2012) 'Circumstance and Strategy: Jointly Authored Supreme Court Opinions', *Nev. L.J.*, 12: 727.

Rhode, D. W., and Spaeth, H. J. (1976). *Supreme Court Decision Making*. San Franciso: W. H. Freeman.

Robbins, I. P. (2012) 'Hiding Behind the Cloak of Invisibility: The Supreme Court and Per Curiam Opinions', *Tul. L. Rev.*, 86: 1197–205.

Rosenthal, J. S., and Yoon, A. H. (2011a) 'Judicial Ghostwriting: Authorship on the Supreme Court', *Cornell L. Rev.*, 96: 1307–43.

Rosenthal, J. S., and Yoon, A. H. (2011b) 'Detecting Multiple Authorship of United States Supreme Court Legal Decisions Using Function Words', 5 *Annals Applied Statistics*, 5: 283–308.

Roy, A. (2023) 'The Inside Story on How Roberts Changed His Supreme Court Vote on Obamacare', *Forbes*. https://www.forbes.com/sites/theapothecary/2012/07/01/the-supreme-courts-john-roberts-changed-his-obamacare-vote-in-may/?sh=512577ed701d. Accessed Jan. 26, 2025.

Schwartz, P. M. (2002) 'Voting Technology and Democracy', *77 N.Y.U. L. Rev.*, 625: 657–58.

Sheppard, B. (2022) 'Contemplating Arbitration in Disputed Congressional Elections: A Case Study with the Closest Senate Election in U.S. History', *20 U.N.H. L. Rev.*, 229: 244–45.

Stone, G. R. (2020) 'A Four-Decade Perspective on Life Inside the Supreme Court', *Harv. L. Rev.* 133: 1010, 1028.

William, L., *et al.* (2013) 'Using Algorithmic Attribution Techniques to Determine Authorship in Unsigned Judicial Opinions', *Stan. Tech. L. Rev.* 16: 503.