

ARTICLE

# Measuring Media Criticism with ALC Word Embeddings

Christopher Barrie<sup>1</sup> , Neil Ketchley<sup>2</sup> , Alexandra Siegel<sup>3</sup>  and Mossaab Bagdouri<sup>4</sup>

<sup>1</sup>Department of Sociology, New York University, New York, NY 10012, USA; <sup>2</sup>Department of Politics and International Relations, University of Oxford, Oxford OX1 2JD, UK; <sup>3</sup>Department of Political Science, University of Colorado, Boulder, CO 80309, USA; <sup>4</sup>BagTek, Tetouan 93020, Morocco

**Corresponding author:** Christopher Barrie; E-mail: [christopher.barrie@nyu.edu](mailto:christopher.barrie@nyu.edu)

(Received 4 December 2024; revised 11 March 2025; accepted 23 March 2025)

## Abstract

The ability of news media to report on events and opinions that are critical of the executive branch of government is central to media freedom and a marker of meaningful democratization. Existing indices use scoring criteria or expert surveys to develop country year measures of media criticism. In this article, we introduce a computationally inexpensive and fully open-source method for estimating media criticism from news articles using à la carte (ALC) word embeddings. We validate our approach using Arabic-language news media published during the Arab Spring. An applied example demonstrates how our technique generates credible estimates of changes in media criticism after a democratic transition is ended by a military coup. Experiments demonstrate that the method works even with sparse data. Analyses of synthetic news media demonstrate that the method extends to multiple languages. Our approach points to new possibilities in the monitoring of media freedom within authoritarian and democratizing settings.

**Keywords:** Text as data; word embeddings; media freedom; MENA; democracy

**Edited by:** Daniel J. Hopkins and Brandon M. Stewart

## 1. Introduction

How can we measure the core characteristics of democracy? This question has animated political science for nearly a quarter of a century (see, e.g., Bush 2017; Claassen *et al.* 2024; Munck and Verkuilen 2002; Przeworski *et al.* 2000). It has also motivated a number of large-scale data collection projects—including Freedom House and Varieties of Democracy (V-Dem)—that are increasingly used by scholars and policymakers to understand trends in democratization. A recent development in this field is the use of surveys to generate disaggregated measures of democracy, where country experts score regimes on a range of theoretically-relevant features (Coppedge *et al.* 2011). As well as being very resource intensive, these techniques have generated considerable debate about the extent to which bias from human coders exaggerates the extent of democratic backsliding worldwide (Knutsen *et al.* 2024; Little and Meng 2024; Widmann and Wich 2022). In this context, some democratization scholars have argued for the development of more “objective” measures of democracy, derived from repeated empirical observations of regime behavior (Little and Meng 2024).

We contribute to this debate by showing how news articles can be utilized to develop measures of media freedom, which is a key component of the indices that are used to study and monitor democratization and autocratic backsliding. Specifically, we focus on measuring the media’s ability to report events and opinions that are critical of the political executive—an essential feature of democratic life that appears in all major democracy datasets and is often the most important component of empirical

indexes measuring media freedom. This item is particularly important given the changing nature of authoritarianism. With the rise of new types of autocratic governance, there has emerged a new form of media capture and control. Contemporary “informational autocrats” continue to police limits on acceptable political reporting, but also derive benefits from allowing certain forms of media coverage and critical commentary (Egorov, Guriev, and Sonin 2009; Guriev and Treisman 2019; Walker and Orttung 2014). In these contexts, while some degree of free speech is permitted, reporting on events or opinions that criticize regime leaders or the top of the political power structure often constitutes a red line (Lorentzen 2014). It follows that the degree to which this red line is enforced offers a tangible measure of media freedom.

To capture changes in media criticism, we introduce a technique that builds on a recent advance in unsupervised word-embedding approaches: “A la Carte” (ALC) word embeddings (Arora *et al.* 2018; Khodak *et al.* 2018; Rodriguez, Spirling, and Stewart 2023). Our technique requires no human input or financial investment beyond the collection of media articles and the minimal computational costs of training an embedding layer and is applicable to any context with a national news media. It is also more granular and responsive to changes in political context than traditional methods for measuring media freedom, e.g., expert surveys, which typically measure developments in media criticism at the country year level. To implement our method, we measure the distance in semantic space between a vector of target words, i.e., the names of political leaders or the titles of their offices, and language found in news media connoting either support or opposition. Drawing on both real and synthetic news media, we show how the proximity of our target words to language connoting opposition is interpretable as a robust measure of criticism. This innovation enables us to recover the level of critical news or opinion in the media at units of varying scales (e.g., articles, publications, or countries) and measures of time (e.g., days, weeks, and months), thus providing considerably more flexibility and granularity than the country-year measures that are currently available.

To validate our approach, we first draw on a large corpus of 8.5 million Arabic-language news media published over the period 2008–2019 generated from five countries in the Middle East and North Africa (MENA). This period, which coincides with the 2011 Arab Spring, witnesses democratization processes, sustained anti-regime protests, a military coup, and authoritarian backsliding. During our analysis period, three of the countries in our sample (Algeria, Morocco, and Saudi Arabia) maintained persistent and deeply entrenched autocratic politics, while two (Egypt and Tunisia) experienced democratic transitions. The variation in our cases allows us to determine whether our approach accurately recovers changes in political reporting that follow from structural political change. We demonstrate our approach using local-language media as this is the most relevant when understanding the degree to which print news media can openly report on events and opinions that are critical of regime elites.

The article proceeds in five parts. First, we outline how to construct our media criticism measure, and then compare our scores to V-Dem. As we show, our approach to quantifying criticism of the executive in national news media closely tracks the values recorded in expert surveys during periods of substantive political change. V-Dem performs less well in stable autocracies, missing time-varying changes in the level of media criticism. This suggests, that expert surveys may capture large changes in easy-to-recognize cases, but can miss less dramatic developments in authoritarian contexts. Second, we demonstrate with a series of experiments that the technique can recover reliable estimates with sparse data. Third, we demonstrate the utility of media criticism scores derived from news media for both descriptive and causal research. For descriptive research, we demonstrate that changepoint analyses of our media criticism scores recover shifts in media freedom that align with case knowledge. For causal research, we demonstrate that comparison cases can be used to generate credible estimates of backsliding events such as military coups on media criticism. Fourth, we generate a series of synthetic articles across seven additional languages to demonstrate that the method extends to multiple linguistic domains. A battery of additional checks, including human validation and design-based supervised learning, underscore the validity and robustness of our approach. We conclude with a discussion of how the method may extend to multiple domains beyond media criticism.

## 2. Media Freedom and Its Measurement

Control of the media constitutes one of the most powerful weapons in the authoritarian arsenal (McMillan and Zoido 2004). Media capture by the state leads to censorship of unfavorable news and events, the distortion of facts, and pro-government agenda setting (Field *et al.* 2018; Woo 1996). Research from diverse contexts suggests that media capture has important real-world consequences, including shifting policy attitudes to favor government positions, boosting party membership, increasing the vote share for pro-regime parties, inciting violence against political opponents, stifling collective action, and reducing aggregate political knowledge (Adena *et al.* 2015; Chen and Yang 2019; Enikolopov, Petrova, and Zhuravskaya 2011; King, Pan, and Roberts 2013; Yanagizawa-Drott 2014).

Many modern authoritarian regimes employ both direct and indirect means of control (Guriev and Treisman 2019). These measures include preventing outlets from reporting on critical content by arresting journalists and editors, prosecuting media owners under the guise of national security laws, conducting punitive tax audits, manipulating government advertising, and imposing “seemingly reasonable” content restrictions (Simon 2014). Contemporary “informational autocrats” also derive benefits from allowing some media criticism as this aids in the functioning of government and provides a veneer of political freedom (Guriev and Treisman 2020; Walker and Orttung 2014). However, direct criticism of the executive branch of government—either in the form of hostile editorials, or coverage of events that directly criticize regime leaders such as protests—is rarely tolerated. Infractions can incur severe penalties, including sizable fines, imprisonment, and state-sanctioned violence (Carter and Carter 2021; Lorentzen 2014).

Against this backdrop, the ability of media outlets to report on events and opinions critical of regime elites has become a key variable for the construction of composite indices of both media freedom and democracy (AMB 2022; FreedomHouse 2017; RSF 2022; Whitten-Woodring and James 2012). To date, social scientists have relied mainly on panels of expert survey respondents to develop measures of media freedom. The V-Dem dataset is one of the most widely used and sophisticated examples of this approach—and provides survey responses for the measurement of media criticism specifically (Coppedge *et al.* 2021; Lührmann, Marquardt, and Mechkova 2020). In their analysis of different indicators of media freedom, Solis and Waggoner (2021) find that the V-Dem variable measuring the ability of media outlets to criticize the government, contributes most information to their latent measure of media freedom. For these reasons, we focus on media criticism as a key determinant of media freedom overall. Given its wide uptake, we can also use the V-Dem measures of media criticism as an initial reference against which to compare the text-based estimates of media criticism.<sup>1</sup>

## 3. Measuring Media Criticism

Our main approach to measuring media criticism exploits newly developed word-embedding approaches to project words that appear close to the mention of any leader onto a vector index of opposition and support. This enables us to determine whether the leader(s) are the subject of news of events or opinions that are more or less critical over time and means we can capture a core feature of media freedom—the degree to which criticism of the political executive is permitted.

To validate our approach, we begin by analyzing media reporting in five MENA countries—Algeria, Egypt, Morocco, Saudi Arabia, and Tunisia—over the period from 2010 to 2019, coinciding with the Arab Spring. Two of these countries—Egypt and Tunisia—saw substantial political change over our observation period; the other three saw relative stability. We refer to the former as our “change cases”

<sup>1</sup>The question in the expert survey asks: “Of the major print and broadcast outlets, how many routinely criticize the government?” The response options are: 0: None; 1: Only a few marginal outlets; 2: Some important outlets routinely criticize the government, but there are other important outlets that never do; 3: All major media outlets criticize the government at least occasionally. The confidence intervals are derived from additional questions asking respondents how confident they are in the accuracy of their response. Although other indices are available, none run over the entirety of our observation period and do not provide disaggregated information for the specific index components on media criticism.

and to the latter as our “stability cases.” We provide background information about happenings in each case in Section A of the Supplementary Material.

Researchers can increasingly access new media from autocratic contexts at scale. We draw on a set of Arabic-language news articles taken from news aggregation websites for each of the countries in the sample: <https://www.djazairess.com/> (Algeria), <https://www.masress.com/> (Egypt), <https://www.maghress.com/> (Morocco), <https://www.sauress.com/> (Saudi Arabia), and <https://www.turess.com/> (Tunisia). Articles date from 2008–2019 for each of the countries in our sample. We only include sources that can credibly be characterized as news sources.<sup>2</sup> In total, across all five countries, we have 335 newspaper sources and a subsample of 8.5m unique news articles. We provide a full list of the newspaper sources and number of articles sampled for each country in Table B.1 in the Supplementary Material.

The smallest size sample in our data is the Tunisia corpus, numbering around 1.7m news articles. To ensure comparable overall sample sizes for our analytical samples, we therefore set the total size of the Tunisia sample as our upper limit for sampling other countries. To identify passages related to the political executive, we use the names of the leaders in each of the countries during the time period of their rule. A table detailing each leader and the period of their rule is provided in Table B.2 in the Supplementary Material.

### 3.1. Word Embedding

Word-embedding techniques represent an important recent advance in the large-scale analysis of text and, in particular, semantic meaning (Caliskan, Bryson, and Narayanan 2017; Charlesworth, Caliskan, and Banaji 2022; Garg *et al.* 2018). The basic requirement for training a word embedding layer is to convert a corpus of text into a term co-occurrence matrix. With this matrix we are then able to exploit pre-packaged algorithmic architectures to learn the pattern of co-occurrences and derive a distributional representation of each word in the corpus in vector space. To date, most practitioners use one of the GloVe (Pennington, Socher, and Manning 2014) or Word2Vec (Mikolov *et al.* 2013) modeling approaches. In the following, we use GloVe to train our embedding layer. We do so by sampling a maximum of 1.5m news articles across all countries combined and estimating a single embedding layer using the combined data from all countries. We also detail below experiments in the minimum effective training data sample size for this procedure.

While a promising agenda, studying semantic change over time in this way is confronted with two problems: 1) computational inefficiency; and 2) identification. Training an embedding layer is computationally expensive. As such, examining shifts in the relationship between words (proximity in vector space) over covariates of interest (such as time) requires a large amount of compute power, especially for large corpora. A corollary problem is that when embedding layers are trained over different temporal units, this means that over-time comparisons are no longer robust due to lack of identification in the underlying vector space (Hamilton, Leskovec, and Jurafsky 2016; Rodriguez *et al.* 2023).

### 3.2. ALC Word Embeddings

A recent innovation by Khodak *et al.* (2018), and implemented and extended by Rodriguez *et al.* (2023), helps solve these problems. The technique—“ALC on Text”—provides a computationally efficient way to identify semantic change over time. The advantage of this technique is that we are able to use a single pre-trained embedding layer, and accompanying transformation matrix, to induce embeddings for a given target word over time without having to retrain an embedding layer for each unit of time.

The efficiency gains of the ALC approach come from the realization that embeddings for a particular (even very rare) target word may be derived by averaging the vectors of embeddings for words within its (here: six-word) context window from a pre-trained embedding layer (Arora *et al.* 2018; Khodak

<sup>2</sup>That is, we exclude outlets that focus exclusively on sports and/or celebrity gossip.

*et al.* 2018; Rodriguez *et al.* 2023). Once we have the embeddings of context words, we can then take the average of these vectors to derive our distributional representation of the target word. A transformation matrix—required to downweight words (such as stop words) that appear with high frequency—is then computed using the term co-occurrence matrix used to generate the embedding layer as well as the embedding layer itself (Khodak *et al.* 2018; Rodriguez *et al.* 2023). Unlike other, e.g., dictionary-based methods, then, ALC embeddings do not rely on specific words appearing within the text corpus to derive a measurement nor does it rely on our target word appearing with high frequency in the embedding layer.<sup>3</sup>

In our application, we train an embedding layer across a combined sample of all newspaper sources in all countries that make up our sample. We refer to this as our “reference embedding.” We pre-process the text by removing numbers, stopwords, and punctuation. We then use the GloVe algorithm, the R packages `quanteda` (Benoit *et al.* 2018) and `text2vec` (Selivanov, Bickel, and Wang 2025). We set vector dimensionality to length 300, and use a window size of six. The maximum number of iterations for training the embedding layer was set to 100, and the models all converged under this threshold for each country. We pruned the vocabulary over which to train the embedding layer such that the overall dimensionality of the resulting co-occurrence matrices was  $\sim 30000 \times 30000$  (i.e., 30000 unique words). We then compute the transformation matrix required for the ALC approach using the R package `conText` developed by (Rodriguez *et al.* 2023). This is used to reweight words appearing with high frequency in the corpus. We also conduct experiments to determine the size of the feature space required to reliably detect signal.

### 3.3. Criticism Index

Unlike word frequency or topic modelling approaches, which use a bag of words as their foundation, word-embedding techniques retain the context and order of the text. One advantage of this is that the embedding layers retain information on the semantic associations between words, which means we can use matrix arithmetic to perform analogy tasks or derive index (vector) representations of concepts of interest (Bolukbasi *et al.* 2016). Our target leader words are detailed in Table B.2 in the Supplementary Material. We then calculate a criticism dimension by subtracting the vector for the word “opposition” (معارضة) from the vector for the word “support” (دعم) in our reference embedding. The rationale for subtracting one word from the other is that this means the both poles of the index are interpretable. Ultimately, it also reduces the number of operations as it means projecting target words onto one index rather than two. This gives us a single “criticism index,” which will be used to capture coverage of events or editorial opinions that are critical of the political executive. We infer that text with high cosine similarity to the opposition pole is likely to be critical.

By criticism, we mean both news of events that are critical of the leader and editorial articles that are directly critical of the executive and its policies.<sup>4</sup> Here, our index captures one or all of three different things, which we understand to denote criticism of the executive:

1. Reporting on events that target the figure of the leader, e.g., protests against the leader or their policies.
2. Opinion articles and editorials detailing failings and allocating blame to the leader or his/her government.
3. Second-hand criticism of the figure of the leader, e.g., reporting on public opinion and soundbites of citizens or other figures critical of the leader.

<sup>3</sup>See also Rodriguez *et al.* (2023) for a full discussion of the advantages of ALC embedding compared to, for example, dictionary-based methods of text analysis.

<sup>4</sup>See also the definition we provided to human coders for the validation steps we describe in the Supplementary Material.

### 3.4. Projecting Words Over Time

We can observe temporal trends by calculating the cosine similarities between our target words of interest and our criticism index. To recover the over-time cosine similarities, we first split our observation period into year-week slices, and then get the context words around our target leader words for each country week. Using the ALC approach, we then estimate a time-period-specific embedding for the leader from the words appearing around their name over this time period. We do so by taking the average of the vectors of surrounding context words from our pre-trained references embedding layer for each of the leaders in each country in our sample respectively. We then combine these context words and apply the transformation matrix to downweight commonly appearing words. The weighting specified for the transformation matrix determines the extent to which commonly appearing words are penalized. A larger weighting means fewer words are downweighted.<sup>5</sup> Here, we set our transformation matrix weighting at 100, which is similar to other published work (Rodriguez *et al.* 2023). Below, we also detail experiments to determine the influence of the transformation matrix weighting on results.

From this procedure we are able to induce a single period-specific embedding for each leader over each time period. Once we have recovered these embeddings, we can then project them onto our criticism index by calculating the (l2-normalized) cosine distance between the vectors for each of our leaders and our criticism index over time. To aid applied researchers implement our approach, we have summarized this process in Figure 1.

## 4. Observational Diagnostics and Causal Effects

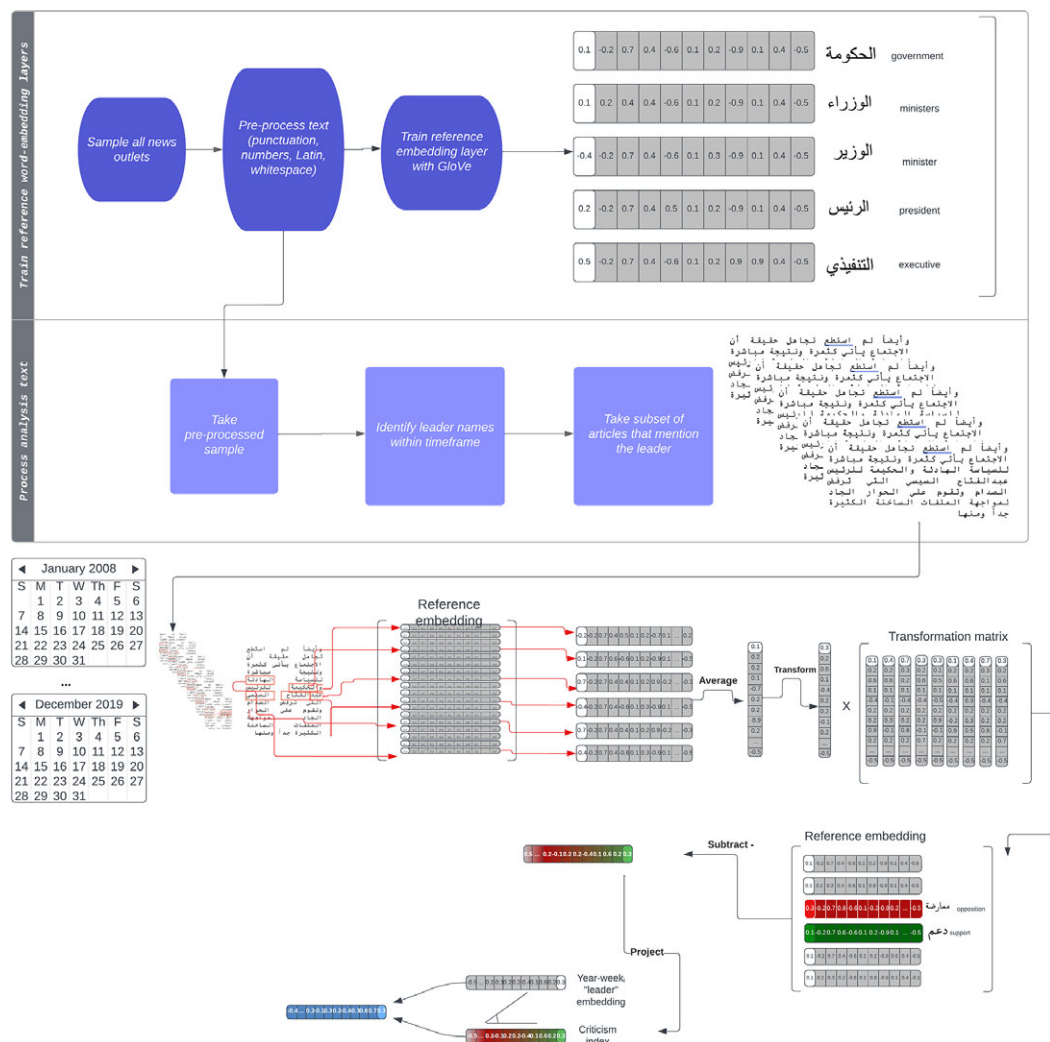
To illustrate the validity and potential use cases for our approach, we focus on media criticism in our change cases: Egypt and Tunisia. We do so for two reasons: 1) to demonstrate the value of the media criticism scores in applied observational and causal settings; 2) to demonstrate how we might benchmark the substantive importance of an observed change in media criticism scores. The first, changepoint, technique provides diagnostics of what constitutes statistically significant change in observed levels of media criticism—and tells us whether such a change aligns with case knowledge. The second, synthetic difference-in-differences, technique benchmarks the size of any change to another case in order to provide counterfactual causal estimates of the effect size of an event in time.

### 4.1. Changepoint Analysis

To provide a diagnostic routine for detecting signal of abrupt change in the levels of observed media criticism, we use a conventional cumulative sum (CUSUM) changepoint approach to detect structural changes in the time-series data (Zeileis *et al.* 2002). The CUSUM approach works by estimating model residuals as a function of the time parameter of interest. It does so by estimating an OLS model of the outcome of interest, calculating the cumulative sum of standardized residuals over time, and comparing these to a null hypothesis of no change. Here, our specification is:  $\cos\_sim_t = \beta_0 + \beta_1 \text{week}_t + \epsilon_t$  where  $\cos\_sim_t$  is the dependent variable (cosine similarity) at time  $t$ ;  $\beta_0$  is the intercept;  $\beta_1$  is the slope coefficient for the predictor  $\text{week}_t$ ; and  $\epsilon_t$  is the error term at time  $t$ . The F-statistic in this approach, provides us with an over-time estimate of model fit under two competing hypotheses: one of no structural change and another of structural change. A high F-statistic at a given point in time provides evidence of improved model fit when accounting for some structural change in over-time variation.<sup>6</sup>

<sup>5</sup>Technically, this weighting procedure is downweighting common directions in the embedding space. We thank Pedro Rodriguez for this clarification.

<sup>6</sup>This is provided by comparing observed residuals to those expected under a null of a random Wiener or Brownian process. That is, the F-statistic provides no indication of the direction of any structural change. In order to determine the direction, we refer to the original criticism score estimates.



**Figure 1.** Data analysis pipeline for measurement of media criticism in Arabic-language news media. This figure, illustrates the key steps in our methodology. (1) We collect news articles from publicly available sources across multiple countries and preprocess the text by removing stopwords, punctuation, and irrelevant content. (2) We identify mentions of political leaders and extract context words appearing within a six-word window around each mention. (3) Using pre-trained GloVe embeddings, we generate a reference embedding layer for all articles, which forms the basis for estimating media criticism. (4) We apply the ALC embedding method to construct time-specific word embeddings for each leader, allowing us to track changes in media discourse over time. (5) We compute a criticism index by projecting leader embeddings onto a semantic dimension spanning words associated with support and opposition. This pipeline enables us to estimate media criticism dynamically and at a granular temporal scale.

## 4.2. Synthetic Difference-in-Differences

We envisage that researchers will want to use our estimation approach to ask counterfactual questions. In particular, they may look to make causal estimates of the effect size of political events on media criticality of the executive. For our cases, the most obvious is the 2013 coup in Egypt. Here, the counterfactual question is: what would media criticism in Egypt have looked like had a coup not happened? The natural comparison case is Tunisia. Both Egypt and Tunisia experienced democratic breakthroughs in early 2011. Both are Arabic-speaking, Muslim-majority republics where entrenched dictators were overthrown through street-level mobilization within a few months of the other (Brownlee, Masoud, and

Reynolds 2015; Ketchley and Barrie 2020). In both cases, precarious and highly polarized democratic transitions unfolded, with secular political forces competing for electoral power against organized Islamist movements (Nugent 2020). Crucially, both cases also saw a proliferation of new independent media organizations and the lifting of long-standing restrictions on media reporting during the post-breakthrough democratic transitions (El- Issawi 2016). However, unlike in Tunisia, Egypt's democratic transition was abruptly ended in mid-2013, when a military coup overthrew the country's first democratically-elected president, sparking thousands of anti-government street protests and a cycle of contention that targeted Egypt's post-coup leadership (Ketchley 2017, chapter 6).

For our main counterfactual analysis, we exploit the availability of news media from Tunisia to implement the synthetic difference-in-differences estimation procedure as described in Arkhangelsky *et al.* (2021). Our analysis uses a panel of ten Egyptian and ten Tunisian newspapers ( $i$ ) observed at weekly periods ( $t$ ) beginning at the start of Egypt's democratic transition in February 2011.<sup>7</sup> Following the case literature, we assume that Egyptian and Tunisian newspapers observed prior to the coup are operating in transitional democratizing contexts where they are more able to report on news and opinion that criticizes the executive, while newspapers in Egypt after the coup were more constrained in their reporting. Our econometric specification is thus:  $\cos\_sim_{it} = L_{it} + \tau_{it}W_{it} + \epsilon_{it}$ , where  $\tau_{it}$  is the effect of the coup on the cosine similarity score of newspaper ( $i$ ) at week ( $t$ ), and we estimate the average of  $\tau_{it}$  over the observations where  $W_{ij}=1$ . The matrix  $L_{it}$  are simple two-way fixed effects at the unit and week level. Unit weights ( $\hat{\omega}$ ) match the pre-trend of the treated newspapers with the untreated controls (here: Tunisian newspapers), while time weights ( $\hat{\lambda}$ ) minimize the differences between the pre and post-treated periods for the controls.<sup>8</sup> In Section H of the Supplementary Material, we also estimate an interrupted time series model. This strategy is useful when applied researchers want to estimate the effect of an event on media criticality, but lack media articles from a comparison case.<sup>9</sup> We can also imagine that researchers might use newspaper media criticality scores from different contexts to estimate a comparative interrupted time series.

### 4.3. Synthetic Data Simulation

Applied researchers will also want to implement our proposed technique in other languages. To assess this, we innovate by generating synthetic data using two OpenAI large language models (LM) (specifically, gpt-3.5-turbo and gpt-4o). The prompts and code we used to generate these data are in code block 1 in the Supplementary Material. We used a limited prompting design, asking the model to generate a series of 500 articles that were “critical” and 500 “not critical” articles of a political figure we refer to as POLITIFIG.<sup>10</sup> We use this neutral denotation to mitigate against activating any biases baked into the training data. We iterate over seven additional languages as well as Arabic. These data are useful for two key reasons: 1) they provide evidence of the generalizability of our technique to other languages; 2) they are designed specifically to include articles that are variously “critical” or “not critical” meaning we are able to determine whether our criticism index is actually capturing this concept. We translate words for support and opposition into each of these languages (see Figure C.1 in the Supplementary Material for the translations used). Instead of re-estimating embedding layers for these languages, we use the pre-trained embeddings for each language provided by Wirsching *et al.*

<sup>7</sup>To ensure a balanced panel, we use a random forest imputation algorithm implemented in the *missRanger* R package (Mayer 2023) to recover units with missing media criticism scores at the beginning of the time series. This naturally will raise concerns about whether values are missing at random. We are confident that the majority are random and are due to breakdowns in the automated crawlers that power the news aggregation platforms. See also the Robustness section for details of further validation checks that account for potential measurement error.

<sup>8</sup>We implement our analysis using the R package *synthdid* (Arkhangelsky 2023).

<sup>9</sup>As we show, this approach also recovers credible estimates of the effect of the coup on media criticism in Egyptian newspapers (see Figure I.1 in the Supplementary Material).

<sup>10</sup>For gpt-4o, we generate 50 of each article type for cost reasons.

(2025). We select those languages that Wirsching *et al.* (2025) have validated with human coders: Arabic, Chinese, English, French, Japanese, Korean, Russian, and Spanish.<sup>11</sup>

## 5. Estimating Media Criticism

Using the pipeline described in Figure 1, we first generate country-level descriptive trends, which we benchmark to expert survey scores. Our word-embedding estimates of media criticism in our change cases, closely track those reported in V-Dem (see Figure 2). Spearman's  $\rho$  ranges from .85 to .93 for Egypt and Tunisia, respectively—the two countries that underwent democratic transitions during the observation window.<sup>12</sup> In Egypt, we see an increase in media criticism in the aftermath of Mubarak's ousting in 2011 uprising followed by a sharp decrease in the aftermath of the coup of 2013. In Tunisia, we see a sharp increase in media criticism in the aftermath of the 2010–11 uprising that then stays at approximately the same level for the ensuing years to 2019. For our stability cases, we see relatively flat lines throughout the observation period. We do nonetheless observe that there is more substantial variation in media criticism scores in our stable cases than the expert survey scores would suggest.

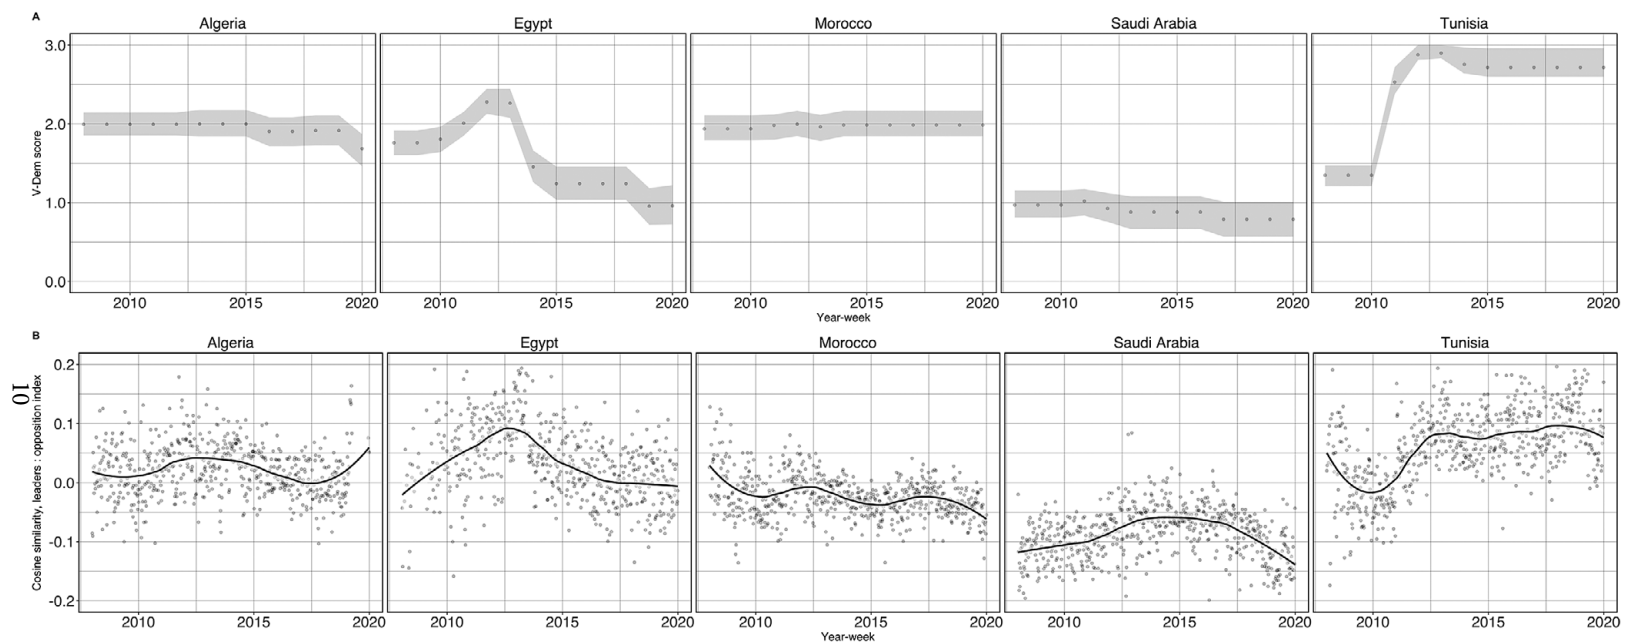
To aid applied researchers adopting this technique, we run several tests to explore the effects of several parameter choices when: a) determining the minimum number of leader words required in each time unit; b) determining the size of the training data necessary for our reference embeddings; c) determining the vocabulary (feature) size of the reference embedding layer; and d) specifying the weighting of the transformation matrix. Full results of these experiments are detailed in the Supplementary Material. Overall, we see that the same basic trends obtain across most iterations of  $n$  of leader words, corpus (training data) size, vocabulary size, and transformation matrix. While we caution readers against placing too high trust in results when data are particularly sparse for a given time period, it appears that just two occurrences of a leader word is actually sufficient to extract a reasonably reliable signal. Even training data of size 10k news articles seems to recover the basic trends observed above, with the exception of Tunisia. Perhaps surprisingly, a vocabulary size of 1k also picks up a comparable signal across all countries. For the transformation matrix, only the most severe weight setting (100k) exhibits variation that departs markedly from the other parameter settings. Taken together, these checks are good news for applied researchers, as they demonstrate that our approach is able to reliably detect known signals even when the size of the data is comparatively small. It also means the computational cost of the procedure can be lower. That said, to train a word embedding layer with even the largest vocabulary size, applied researchers need no more computational power than provided by a modern personal computer. After the initial training, the estimation of criticism scores takes seconds.

### 5.1. Detecting Changepoints

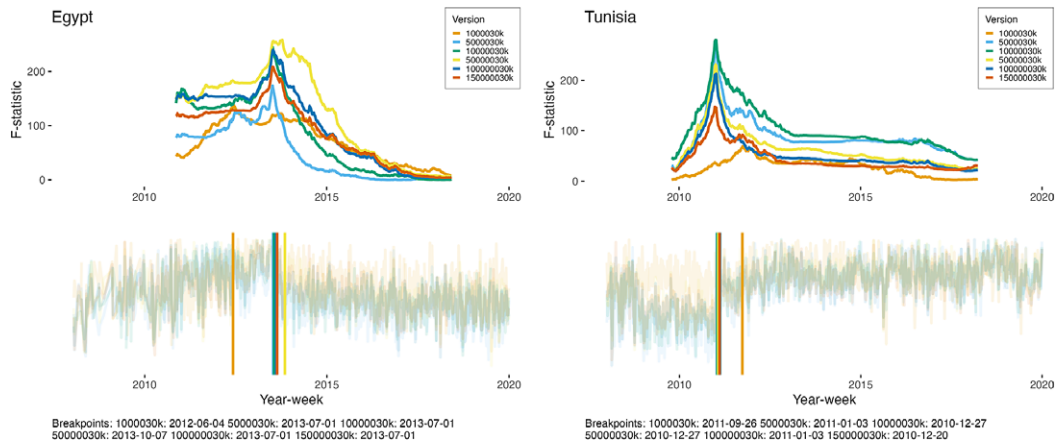
The above estimates demonstrate a high correlation between our text-based measures of media criticism in our change cases and those deriving from expert surveys measured at the country year level. If we are to incorporate text-based indicators into standard metrics of media freedom, we need a method to detect changes and a measure of the uncertainty associated with these changes. These rule-of-thumb metrics are central to other contributions using large text data as the foundation for early-warning systems (Balashankar, Subramanian, and Fraiberger 2023; Stoleran *et al.*, 2023). To achieve this, we implement our changepoint procedure that estimates an F-statistic of model fit under assumptions of structural change in the level of media criticism. The point where the F-statistic peaks can be understood as the most probable changepoint.

<sup>11</sup>We acknowledge that the latest large LM are less performant in low-resource and non-Western languages (Jiao *et al.* 2023; Nasution and Onan 2024). However, for this task we are less interested in the model producing synthetic reporting that would pass a Turing test and much more in generating grammatically correct text that uses language in a way that approximates media reporting.

<sup>12</sup>To make this comparison, we calculate a yearly average media criticism score and 95% confidence intervals.



**Figure 2.** **A:** V-Dem media criticism scores over time across all countries; **B:** Normalized cosine similarity criticism scores over time across all countries. Lines in panel **B** is smoothed (LOESS) curve with span ( $\alpha$ ) set to 0.5. Confidence intervals for the V-Dem scores in panel **A** are based on cross-coder aggregations calculated according to the Bayesian item response theory measurement model described in Coppedge *et al.* (2019).



**Figure 3.** *Top panel:* F-statistics over time for changepoint procedure across versions; *Bottom panel:* text-based criticism scores for Egypt and Tunisia and breakpoints for each version. Breakpoints displayed with slight offset for visibility.

Figure 3 displays the weeks with the highest probability across our two change cases. The most pressing concern for applied researchers will likely be the size of training data required to detect signal. As such, we estimate our changepoint models over all six versions of our training data, i.e., from 10k to 1.5m unique news articles. We keep the feature size constant at 30k.<sup>13</sup> For both Egypt and Tunisia, we see that all versions, with the exception of the smallest in Tunisia, recover estimates of structural change points that align with case knowledge. In Egypt, this is at the beginning of July 2013 following a military coup; in Tunisia this is in late-December of 2010 and January of 2011 when Ben Ali's dictatorial regime was ousted from power.

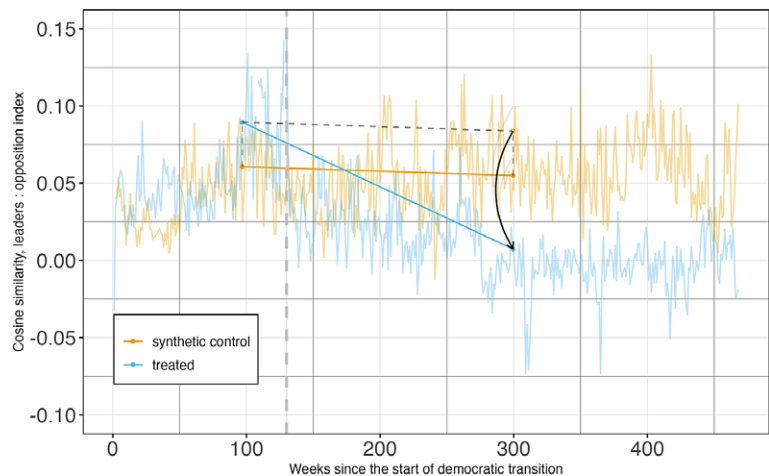
## 5.2. Counterfactual Estimation

Figure 4 shows the results of our synthetic difference in difference analysis to estimate the effect of the military coup in July 2013 in Egypt on criticism of the executive. As noted, to generate a plausible pre-coup trend, we use media criticism scores from newspapers from nearby Tunisia—which was also undergoing a democratic transition during this period—to construct unit and time weights in a two-way fixed effects model. The outcome measure is a newspaper's criticism score assigned to the year-week from the period of democratic breakthrough in early 2011 through to 2019. The results suggest that the treatment in Egypt led to a substantive, enduring, and statistically significant diminution in media criticism, roughly equivalent to a one standard deviation decrease relative to pre-coup media criticism scores ( $p < .001$ ). This marked reduction in media criticism of the post-coup executive comes despite Egypt experiencing rampant inflation, currency devaluation and a foreign exchange crisis, thousands of anti-coup protests that continued for years after the military's seizure of power, other episodes of street-level mobilization including against food prices and unpopular foreign policy decisions, and protracted insurgencies in the country's border provinces (Grimm 2019; Ketchley and El-Rayyes 2017; Ketchley 2017; Nugent and Siegel 2024).

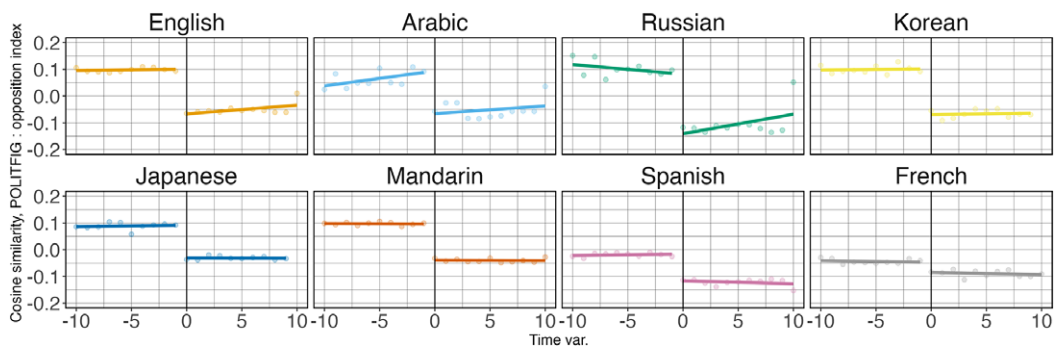
## 5.3. Synthetic Data

Figure 5 displays the results of re-estimating our main analysis across seven additional languages as well as Arabic using synthetic data. The text data is split into ten time units corresponding to 100 articles

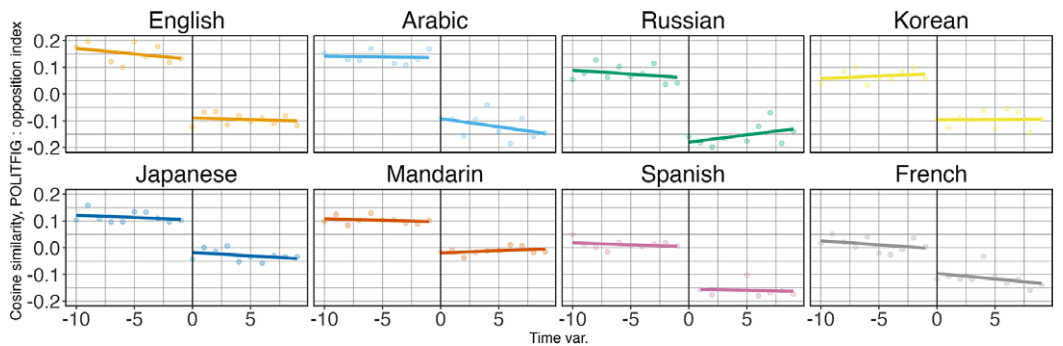
<sup>13</sup>We provide the full distribution of F-statistics in Figure H.1 in the Supplementary Material.



**Figure 4.** Treatment effect of the coup (black arrow) on media criticism in Egypt using Tunisian newspapers as a counterfactual. Dashed line marks the coup.



**Figure 5.** Normalized cosine similarity criticism scores over time across eight languages using synthetic articles generated by gpt-3.5-turbo. Coloured lines represent the linear best fit over each period. The line at 0 is the point at which the LLM shifts to producing “not critical” articles.



**Figure 6.** Normalized cosine similarity criticism scores over time across eight languages using synthetic articles generated by gpt-4o. Coloured lines represent the linear best fit over each period. The line at 0 is the point at which the LLM shifts to producing “not critical” articles.

each. The call to the OpenAI API specified that language should become less critical after 500 of the 1000 total runs. The imagined time point at which the API shifts to producing “not critical” articles is displayed as 0 in Figures 5 and 6. Reassuringly, we observe that across all languages, there is a marked and statistically significant change at the midway point that is detected by our ALC word embedding

approach for both LM. These results demonstrate that our approach is not sensitive to the input language and may be applied to other cases of authoritarianism and democratization.

#### 5.4. Robustness

In the Supplementary Material, we provide a number of additional tests to ensure the robustness of our approach. These include comparing scores generated by our criticism index to human-labelled values for a sample of news articles, design-based supervised learning, estimating alternative criticism indices, alternative target words for the political executive, alternative normalization procedures when estimating cosine similarity, and additional checks that our scores are not an artefact of our use of “opposition” to connote criticism.

### 6. Discussion and Conclusion

In this article, we propose and empirically validate a computationally inexpensive—and completely unsupervised—approach to scoring a key indicator of media freedom: the level of media criticism directed at the political executive. To date, researchers have relied on expert surveys and composite indices to measure the health of the fourth estate across countries and across time. Drawing on news media from autocracies, as well as synthetic media reports, we build on innovations in the computational analysis of text to demonstrate a method that convincingly recovers estimates of media criticism across transitional and stable contexts. For applied researchers, we demonstrate that these scores can be used in both descriptive and causal settings. A series of experiments show that the technique recovers sensible measures even with sparse data. Using synthetic data, we demonstrate that the technique travels to multiple other languages.

Importantly, the technique we propose is not limited to the study of media criticism alone. A key benefit of our method is that we are able to recover over-time estimates of text-based trends even when the target construct or individual is rare (in the text).<sup>14</sup> Applied researchers might use a version of the method for e.g.: estimating the over-time targets of populist speech through the generation of an index of populism and enumerating a set of targets (e.g., institutions, groups, other countries); estimating the issue positions of individual legislators by estimating an index of support versus opposition and enumerating a set of issue targets (e.g., abortion, gun ownership, and healthcare reform); or estimating changes in the targets and level of hostility online by generating a hostility index and identifying a set of targets (e.g., political groups, individuals, and ideas).

Despite its advantages, our method has several possible limitations. First, while our approach effectively measures media criticism, it does not directly capture media freedom, as criticism may also capture changes in leader popularity or economic performance rather than shifts in press autonomy. That said, even very popular leaders will still be subject to criticism from political opponents, just as all policy platforms inevitably create winners and losers. Thus, there is good reason to believe that dramatic declines in criticism are more likely to result from reduced media freedom rather than increased leader popularity. Indeed, we expect that in environments with free media, criticism of popular leaders will be especially prevalent, as opponents seek to undermine their support. Future research could integrate additional indicators, such as independent assessments of media restrictions, to further disentangle these factors. Second, our method relies on publicly available news sources, which may introduce selection biases if certain types of outlets are underrepresented or disproportionately censored. The increasing availability of news aggregators should help to ameliorate this problem. Third, while the ALC embedding technique allows for efficient estimation of media criticism over time, it does not account for nuanced rhetorical strategies, such as self-censorship or coded dissent, which may be important in authoritarian contexts. Finally, our approach assumes that language associated with opposition and

<sup>14</sup>See Rodriguez *et al.* (2023) for a full discussion.

support is relatively stable over time, though shifts in political discourse could affect our estimates. Despite these limitations, our method provides a scalable and replicable tool for tracking media criticism across diverse settings, offering a valuable complement to expert-coded and survey-based measures.

A number of extensions naturally follow from this advance. Given the temporal granularity we can now rely on, we will be able to incorporate fine-grained measures of media criticism as variables within survey research or into commonly used indices of media freedom. Extensions of the approach might also involve the use of our granular measures as features in a supervised machine-learning context to detect widening or narrowing media freedoms worldwide (Balashankar *et al.* 2023; Mueller and Rauh 2018).

For counterfactual designs, we demonstrate how to use outlet-level measurements of media criticism to estimate the effect of major political events on media reporting. This opens the door to estimation of the causal effects of a wide variety of events on media freedom, including new entrants in media markets, violent episodes, and the advent of alternative media (Guriev and Treisman 2020; Hale 2018; Shirky 2011). Because our approach can be adapted to many languages, it also facilitates comparative analysis. We hope similar methods will be applied in diverse global contexts to augment existing measures and improve our understanding of the dynamics of media freedom in democratic and authoritarian societies alike.

**Acknowledgements.** Versions of the article were presented at conferences and workshops for Academia Sinica, Taiwan, the Talking Methods Seminar at the University of Edinburgh, the Dealing with Messy Data workshop at the University of Edinburgh, the Social Data Science Hub workshop at the University of Edinburgh, the Washington University in St. Louis Comparative Politics Annual Conference, and the 2020 annual conference of the American Political Science Association Conference. For their feedback on previous drafts, we are particularly grateful to: Zachary Steinert-Threlkeld, Mohammad Dhia Hammami, Aybuke Atalay, Maddi Bunker, Laurence Rowley-Abel, Ala' Alrababa'h, and Thoraya El-Rayyes.

**Funding Statement.** The research benefited from an internal departmental grant at the University of Oslo.

**Data Availability Statement.** Data and code required to reproduce the analyses in this article are available at Barrie (2025). A preservation copy of the same code and data can also be accessed via Dataverse at <https://doi.org/10.7910/DVN/NXESQQ>. Due to the data sharing agreement underwriting this research, we are unable to share the raw news text data.

**Author Contributions.** C.B., N.K., A.S. conceived of project and wrote article; C.B. developed measurement technique; N.K. and A.S. developed additional analyses; M.B. collected original data.

**Competing Interest.** The authors declare none.

**Ethical Standards.** The project received Research Ethics Approval from the University of Edinburgh Ref No: ID 286780.

**Supplementary Material.** For supplementary material accompanying this paper, please visit <https://doi.org/10.1017/pan.2025.10012>.

## References

- Adena, M., R. Enikolopov, M. Petrova, V. Santarosa, and E. Zhuravskaya. 2015. "Radio and the Rise of The Nazis in Prewar Germany\*" *The Quarterly Journal of Economics* 130 (4): 1885–1939.
- Allsop, J. 2022. "A Big Step Backward for Tunisia's Press." *Columbia Journalism Review – The Media Today*. [https://www.cjr.org/the\\_media\\_today/tunisia\\_referendum\\_press\\_freedom.php.cjr.org](https://www.cjr.org/the_media_today/tunisia_referendum_press_freedom.php.cjr.org).
- Arkhangelsky, D. 2023. "Synthdid: Synthetic Difference-in-Difference Estimation." R package version 0.0.9.
- Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager. 2021. "Synthetic Difference-in-Differences." *American Economic Review* 111 (12): 4088–4118.
- Arora, S., Y. Li, Y. Liang, T. Ma, and A. Risteski. 2018. "Linear Algebraic Structure of Word Senses, with Applications to Polysemy." *Transactions of the Association for Computational Linguistics* 6: 483–495. Cambridge MA: MIT Press. [https://doi.org/10.1162/tacl\\_a\\_00034](https://doi.org/10.1162/tacl_a_00034).
- Balashankar, A., L. Subramanian, and S. P. Fraiberger. 2023. "Predicting Food Crises Using News Streams." *Science Advances* 9 (9): eabm3449.
- Barrie, C. 2025. "The Process of Revolutionary Protest: Development and Democracy in the Tunisian Revolution." *Perspectives on Politics*. 23 (1): 103–121. <https://doi.org/10.1017/S1537592723002062>.

- Barrie, C. 2025. "Replication Data for: Measuring Media Criticism with ALC Word Embeddings." Harvard Dataverse. <https://doi.org/10.7910/DVN/NXESQQ>.
- Barrie, C., and N. Ketchley. 2018. "Opportunity without Organization: Labour Mobilization in Egypt after the 25th January Revolution." *Mobilization* 23 (2): 181–202.
- Benoit, K., et al. 2018. "Quanteda: An R Package for the Quantitative Analysis of Textual Data." *Journal of Open Source Software* 3 (30): 774.
- Bolukbasi, T., K.-W. Chang, J. Zou, V. Saligrama, and A. Kalai. 2016. "Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings." In *Advances in Neural Information Processing Systems 29 (NeurIPS 2016)*, 4356–4364. Red Hook: Curran Associates Inc.
- Brownlee, J., T. Masoud, and A. Reynolds. 2015. *The Arab Spring: Pathways of Repression and Reform*. Oxford: Oxford University Press.
- Bubeck, S., et al. 2023. "Sparks of Artificial General Intelligence: Early Experiments with GPT-4." arXiv preprint [arXiv:2303.12712](https://arxiv.org/abs/2303.12712).
- Bush, S. 2017. "The Politics of Rating Freedom: Ideological Affinity, Private Authority, and the Freedom in the World Ratings." *Perspectives on Politics* 15 (3): 711–731.
- Caliskan, A., J. J. Bryson, and A. Narayanan. 2017. "Semantics Derived Automatically from Language Corpora Contain Human-like Biases." *Science* 356 (6334): 183–186. Publisher: American Association for the Advancement of Science Section: Reports.
- Carter, E. B., and B. L. Carter. 2021. "Propaganda and Protest in Autocracies." *Journal of Conflict Resolution* 65 (5): 919–949.
- Charlesworth, T. E. S., A. Caliskan, and M. R. Banaji. 2022. "Historical Representations of Social Groups Across 200 Years of Word Embeddings From Google Books." *Proceedings of the National Academy of Sciences* 119 (28): e2121798119.
- Chen, Y., and D. Y. Yang. 2019. "The Impact of Media Censorship: 1984 or Brave New World?" *American Economic Review* 109 (6): 2294–2332.
- Claassen, C., et al. 2024. "Conceptualizing and Measuring Support for Democracy: A New Approach." *Comparative Political Studies* 58 (6): 1171–1198. <https://doi.org/10.1177/00104140241259458>.
- Coppedge, M., et al. 2011. "Conceptualizing and Measuring Democracy: A New Approach." *Perspectives on Politics* 9 (2): 247–267. <https://doi.org/10.1017/S1537592711000880>
- Coppedge, M., et al. 2021. "V-Dem Country-Year/Country-Date Dataset V11."
- Coppedge, M., et al. 2019. "V-dem Methodology v9." *Varieties of Democracy (V-Dem) Project Working Paper*. Gothenburg: University of Gothenburg.
- Egami, N., M. Hinck, B. Stewart, and H. Wei. 2023. "Using Imperfect Surrogates for Downstream Inference: Design-based Supervised Learning for Social Science Applications of Large Language Models." *Advances in Neural Information Processing Systems* 36: 68589–68601.
- Egorov, G., S. Guriev, and K. Sonin. 2009. "Why Resource-poor Dictators Allow Freer Media: A Theory and Evidence from Panel Data." *American Political Science Review* 103 (4): 645–668.
- el Issawi, F. 2012. "Tunisian Media in Transition." Technical report, Carnegie Endowment for International Peace.
- el Issawi, F., and B. Cammaerts. 2016. "Shifting Journalistic Roles in Democratic Transitions: Lessons from Egypt." *Journalism* 17 (5): 549–566.
- Enikolopov, R., M. Petrova, and E. Zhuravskaya. 2011. "Media and Political Persuasion: Evidence from Russia." *American Economic Review* 101 (7): 3253–3285.
- Field, A., D. Klinger, S. Wintner, J. Pan, D. Jurafsky, and Y. Tsvetkov. 2018. "Framing and Agenda-Setting in Russian News: A Computational Analysis of Intricate Political Strategies." In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP 2018)*, 3570–3580. Brussels: Association for Computational Linguistics.
- Freedom House. 2017. "Freedom of the Press Research Methodology." *Freedom House*. Accessed January 26, 2023. <https://freedomhouse.org/freedom-press-researchmethodology>.
- Friedrich-Ebert-Stiftung, fesmedia Africa. 2022. "African Media Barometer (AMB): A Home-Grown Analysis of the Media Landscape in Africa." Windhoek: Friedrich-Ebert-Stiftung, 2004 – present. <https://fesmedia-africa.fes.de/themes/african-mediabarometer>.
- Garg, N., L. Schiebinger, D. Jurafsky, and J. Zou. 2018. "Word Embeddings Quantify 100 Years of Gender and Ethnic Stereotypes." *Proceedings of the National Academy of Sciences* 115 (16): E3635–E3644.
- Grewal, S. 2024. "Military Repression and Restraint in Algeria." *American Political Science Review* 118 (2): 671–686.
- Grimm, J. 2019. "Egypt is not for Sale! Harnessing Nationalism for Alliance Building in Egypt's Tiran and Sanafir Island Protests." *Mediterranean Politics* 24 (4): 443–466.
- Guriev, S., and D. Treisman. 2019. "Informational Autocrats." *Journal of Economic Perspectives* 33 (4): 100–127.
- Guriev, S., and D. Treisman. 2020. "A Theory of Informational Autocracy." *Journal of Public Economics* 186: 104158.
- Hale, H. E. 2018. "How Crimea Pays: Media, Rallyinground the Flag, and Authoritarian Support." *Comparative Politics* 50 (3): 369–391.
- Hamilton, W. L., J. Leskovec, and D. Jurafsky. 2016. "Diachronic Word Embeddings Reveal Statistical Laws of Semantic Change." In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL 2016)*, Vol. 1: Long Papers, 1489–1501. Berlin: Association for Computational Linguistics.

- El- Issawi, F. 2016. *Arab National Media and Political Change*. Basingstoke: Palgrave Macmillan.
- Issawi, F. el. 2020. "Egyptian Journalists and the Struggle for Change Following the 2011 Uprising: The Ambiguous Journalistic Agency Between Change and Conformity." *International Communication Gazette* 82 (7): 628–645.
- Jiao, W., W. Wang, J.-T. Huang, X. Wang, and Z. Tu. 2023. "Is Chatgpt a Good Translator? A Preliminary Study." arXiv preprint arXiv:2301.08745
- Ketchley, N. 2017. *Egypt in a Time of Revolution: Contentious Politics and the Arab Spring*. Cambridge University Press.
- Ketchley, N., and C. Barrie. 2020. "Fridays of Revolution: Focal Days and Mass Protest in Egypt and Tunisia." *Political Research Quarterly* 73 (2): 308–324.
- Ketchley, N., and T. El-Rayyes. 2017. "On the Breadline in Sisi's Egypt." Middle East Report Online.
- Khodak, M., N. Saunshi, Y. Liang, T. Ma, B. Stewart, and S. Arora. 2018. "A La Carte Embedding: Cheap but Effective Induction of Semantic Feature Vectors." In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL 2018)*, Vol. 1: Long Papers, 12–22. Melbourne: Association for Computational Linguistics.
- Kilavuz, M. T., S. Grewal, and R. Kubinec. 2023. "Ghosts of the Black Decade: How Legacies of Violence Shaped Algeria's Hirak Protests." *Journal of Peace Research* 60 (1): 9–25.
- King, G., J. Pan, and M. E. Roberts. 2013. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review* 107 (2): 18.
- Knutsen, C. H., et al. 2024. "Conceptual and Measurement Issues in Assessing Democratic Backsliding." *PS: Political Science & Politics*. 57 (2): 162–177. <https://doi.org/10.1017/S104909652300077X>.
- Lawrence, A. K. 2017. "Repression and Activism Among the Arab Spring's First Movers: Evidence from Morocco's February 20th Movement." *British Journal of Political Science* 47 (3): 699–718.
- Little, A. T., and A. Meng. 2024. "Measuring Democratic Backsliding." *PS: Political Science & Politics*. 57 (2): 149–161. <https://doi.org/10.1017/S104909652300063X>.
- Lorentzen, P. 2014. "China's Strategic Censorship." *American Journal of Political Science* 58 (2): 402–414.
- Lührmann, A., K. L. Marquardt, and V. Mechkova. 2020. "Constraining Governments: New Indices of Vertical, Horizontal, and Diagonal Accountability." *American Political Science Review* 114 (3): 811–820.
- Masoud, T. 2014. *Counting Islam: Religion, Class, and Elections in Egypt*. New York: Cambridge University Press.
- Mayer, M. 2023. "missRanger: Fast Imputation of Missing Values." R package version 2.3.0.
- McMillan, J. and P. Zoido. 2004. "How to Subvert Democracy: Montesinos in Peru." *Journal of Economic Perspectives* 18 (4): 69–92.
- Mikolov, T., K. Chen, G. Corrado, and J. Dean. 2013. "Efficient Estimation of Word Representations in Vector Space." In *Proceedings of the 1st International Conference on Learning Representations (ICLR 2013)*, Scottsdale: Workshop Track.
- Mueller, H., and C. Rauh. 2018. "Reading Between the Lines: Prediction of Political Violence Using Newspaper Text." *American Political Science Review* 112 (2): 358–375.
- Munck, G. L., and J. Verkuilen. 2002. "Conceptualizing and Measuring Democracy: Evaluating Alternative Indices." *Comparative Political Studies* 35: 5–34.
- Nasution, A. H., and A. Onan. 2024. "Chatgpt label: Comparing the quality of human-generated and llmgenerated annotations in low-resource language nlp tasks." *IEEE Access*. 12: 71876–71900.
- Nugent, E. R. 2020. *After Repression: How Polarization Derails Democratic Transition*. Princeton, NJ: Princeton University Press. <https://doi.org/10.23943/princeton/9780691203058.001.0001>
- Nugent, L., and A. Siegel. 2024. "How Exiles Mobilize Domestic Dissent." *Journal of Politics*.
- Pennington, J., R. Socher, and C. Manning. 2014. "GloVe: Global Vectors for Word Representation." In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, edited by A. Moschitti, B. Pang and W. Daelemans, 1532–1543. Doha: Association for Computational Linguistics.
- Przeworski, A., M. E. Alvarez, J. A. Cheibub, and F. Limongi. 2000. *Democracy and Development: Political Institutions and Well-Being in the World, 1950–1990*. Cambridge: Cambridge University Press.
- Rodriguez, P. L., A. Spirling, and B. M. Stewart. 2023. "Embedding Regression: Models for Context-Specific Description and Inference." *American Political Science Review* 117 (4): 1255–1274.
- RSF. 2022. "World Press Freedom Index: Questionnaire 2022." Technical report, Reporters Without Borders.
- Schraeder, P. J., and H. Redissi. 2011. "The Upheavals in Egypt and Tunisia: Ben Ali's Fall." *Journal of Democracy* 22 (3): 5–19. Publisher: Johns Hopkins University Press.
- Selivanov, D., M. Bickel, and Q. Wang. 2025. text2vec: Modern Text Mining Framework for R. R package version 0.6.4. <https://doi.org/10.32614/CRAN.package.text2vec>
- Shirky, C. 2011. "The Political Power of Social Media: Technology, the Public Sphere, and Political Change." *Foreign Affairs* 90 (1): 28–41.
- Simon, J. 2014. *The New Censorship*. New York, NY: Columbia University Press.
- Solis, J. A., and P. D. Waggoner. 2021. "Measuring Media Freedom: An Item Response Theory Analysis of Existing Indicators." *British Journal of Political Science* 51 (4): 1685–1704.
- Stolerman, L. M., et al. "Using Digital Traces to Build Prospective and Real-Time County-Level Early Warning Systems To Anticipate Covid-19 Outbreaks in the United States." *Science Advances* 9 (3): eabq0199.

- TIMEP. 2022. "Freedom of Expression Under Attack by Tunisia's Kais Saied." March 21, 2022. <https://timep.org/2022/03/21/freedom-of-expression-under-attack-by-tunisia-kais-saied/>.
- Vaswani, A., et al. 2017. "Attention Is All You Need." [arXiv:1706.03762](https://arxiv.org/abs/1706.03762) [cs].
- Walker, C., and R. W. Orttung. 2014. "Breaking the News: The Role of State-Run Media." *Journal of Democracy* 25 (1): 71–85.
- Whitten-Woodring, J., and P. James. 2012. "Fourth Estate or Mouthpiece? A Formal Model of Media, Protest, and Government Repression." *Political Communication* 29 (2): 113–136.
- Widmann, T., and M. Wich. 2022. "Creating and Comparing Dictionary, Word Embedding, and Transformer-Based Models to Measure Discrete Emotions in German Political Text." *Political Analysis* 31 (4): 1–16.
- Wirsching, E. M., P. Rodriguez, A. Spirling, and B. M. Stewart. 2025. "Multilanguage Word Embeddings for Social Scientists: Estimation, Inference and Validation Resources for 157 Languages." *Political Analysis*. 33(2): 156–163. <https://doi.org/10.1017/pan.2024.17>.
- Woo, J. 1996. "Television News Discourse in Political Transition: Framing the 1987 and 1992 Korean Presidential Elections." *Political Communication* 13 (1): 63–80.
- Yanagizawa-Drott, D. 2014. "Propaganda and Conflict: Evidence from the Rwandan Genocide." *The Quarterly Journal of Economics* 129 (4): 1947–1994.
- Zeileis, A., F. Leisch, K. Hornik, and C. Kleiber. 2002. "Strucchange: An R Package for Testing for Structural Change in Linear Regression Models." *Journal of Statistical Software* 7 (2): 1–38.