

The Geopolitical Threat Index: A Text-Based Computational Approach to Identifying Foreign Threats

PETER TRUBOWITZ 

London School of Economics and Political Science, UK

AND

KOHEI WATANABE 

Waseda University, Japan

Few concepts figure more prominently in the study of international politics than threat. Yet scholars do not agree on how to identify and measure threats or systematically incorporate leaders' perceptions of threat into their models. In this research note, we introduce a text-based strategy and method for identifying and measuring elite assessments of international threat from publicly available sources. Using semi-supervised machine learning models, we show how text sourced from newspaper articles can be parsed to discern arguments that distinguish threatening from non-threatening states, and to measure and track variation in the intensity of foreign threats over time. To demonstrate proof of concept, we use news summaries from *The New York Times* from 1861 to 2017 to create a geopolitical threat index (GTI) for the United States. We show that the index successfully matches periods in US history that historians identify as high and low threat and correctly identifies countries that have posed a threat to US security at different points in its history. We compare and contrast GTI with traditional indicators of international threat that rely on measures of material capability and interstate behavior.

Pocos conceptos ocupan un lugar tan importante en el estudio de la política internacional como el de “amenaza.” Sin embargo, los estudiosos no se ponen de acuerdo acerca de cómo identificar y medir las amenazas o incorporar sistemáticamente las percepciones de los líderes sobre las amenazas en sus modelos. En esta nota de investigación, presentamos una estrategia y un método basados en textos para identificar y medir las evaluaciones de las élites sobre la amenaza internacional a partir de fuentes disponibles públicamente. Mediante el uso de modelos de aprendizaje automático semisupervisado, demostramos cómo se puede analizar el texto de los artículos de prensa para discernir los argumentos que distinguen los estados amenazantes de los no amenazantes, y para medir y rastrear la variación de la intensidad de las amenazas extranjeras a lo largo del tiempo. Para demostrar la prueba de concepto, utilizamos artículos de *The New York Times* desde 1861 hasta 2017 con el fin de crear un índice de amenazas geopolíticas (Geopolitical Threat Index, GTI) para Estados Unidos. Demostramos que el índice coincide con éxito con los períodos de la historia de Estados Unidos que los historiadores identifican como de alta y baja amenaza, e identifica correctamente los países que han supuesto una amenaza para la seguridad de Estados Unidos en diferentes momentos de su historia. Comparamos y contrastamos el GTI con los indicadores tradicionales de amenaza internacional que se basan en medidas de capacidad material y comportamiento interestatal.

Peu de concepts occupent une place aussi importante que celui de menace dans l'étude de la politique internationale. Pourtant, les chercheurs ne s'accordent pas sur la manière d'identifier et de mesurer les menaces ou d'intégrer systématiquement la perception des menaces par les dirigeants à leurs modèles. Dans cet exposé de recherche, nous présentons une stratégie et une méthode basées sur le texte pour identifier et mesurer la façon dont les élites évaluent les menaces internationales à partir de sources accessibles au public. Nous utilisons des modèles reposant sur du machine learning semi-supervisé pour montrer la manière dont le texte issu d'articles de journaux peut être analysé de façon à discerner les arguments distinguant les États menaçants des États non menaçants et à mesurer et suivre la variation d'intensité des menaces étrangères au fil du temps. Pour démontrer la validité de ce concept, nous avons utilisé des articles parus dans le *New York Times* entre 1861 et 2017 et créé un indice des menaces géopolitiques pour les États-Unis. Nous montrons que cet indice concorde bien avec les périodes de l'histoire des États-Unis que les historiens désignent comme à faibles ou fortes menaces et qu'il permet d'identifier correctement les pays qui ont représenté une menace pour la sécurité des États-Unis à différents moments de leur histoire. Nous comparons et opposons cet indice aux indicateurs traditionnels de menace internationale qui reposent sur des mesures de capacité matérielle et de comportement entre États.

Introduction

Few concepts figure more prominently in the study of international politics than threat. Many theories of international politics consider variation in the international threat environment facing countries to be decisive in explaining their foreign policies and behavior—that is, in explaining leaders' decisions to spend precious resources on the military, to sacrifice autonomy by allying with other countries, and so on. Yet as essential as threat is to the study of world politics, scholars do not agree on how to identify and measure

Peter Trubowitz is a Professor of International Relations and Director of the US Centre at the London School of Economics and Political Science.

Kohei Watanabe is a computational social scientist affiliated with Waseda University's Institute for Advanced Study and the US Centre at the London School of Economics and Political Science.

Author's notes: For helpful comments, we are indebted to Kenneth Benoit, Catherine Boone, Koji Kagotani, and the anonymous reviewers at *International Studies Quarterly*. For research assistance, we thank Gokhan Ciflikli. The data underlying this article are available on the ISQ Dataverse at <https://dataverse.harvard.edu/dataverse/isq>.

Trubowitz, Peter, and Kohei Watanabe. (2021) The Geopolitical Threat Index: A Text-Based Computational Approach to Identifying Foreign Threats. *International Studies Quarterly*, doi: 10.1093/isq/sqab029

© The Author(s) (2021). Published by Oxford University Press on behalf of the International Studies Association. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

threats or systematically incorporate leaders' perceptions of threat into their models (Leeds and Savun 2007; Yarhi-Milo 2013). In this research note, we introduce a text-based strategy and method for identifying and measuring elite perceptions of foreign states' capabilities and intentions from publicly available sources. Using machine learning models, we show how text sourced from newspaper articles can be efficiently parsed to discern arguments that distinguish hostile states from peaceful countries, and to measure and track variation in elite perceptions of foreign threats over time.

To demonstrate proof of concept, we use articles from *The New York Times* from 1861 to 2017 to create a geopolitical threat index (GTI) for the United States. In addition to spanning more than one hundred and fifty years of American history, the analysis includes over 385,000 news stories from the "newspaper of record" in the United States about foreign nations' military capabilities, military threats and use of force, and about military force postures, doctrines, and motivations. We also create separate indices for fifteen countries that US policymakers have defined as urgent or potential threats or risks to national security at various points in the nation's history. On the basis of this text data, we are able to track continuities and changes in American perceptions about the source and level of foreign threat facing the United States over time. Because of the language versatility of the machine learning models that we use here (*Latent Semantic Scaling* [LSS] and *Newsmap*), the GTI model can be adapted to create similar threat indices for non-English speaking countries with comparable publicly available data sources.

We evaluate the accuracy of our GTI in three ways. First, we consider whether variations in the index capture and match well-known historical events (e.g., the attack on Pearl Harbor, the launch of Sputnik) and periods in US history that international relations scholars, foreign policy analysts, and diplomatic historians consider "security rich" (e.g., post-bellum America) or "security poor" (e.g., the Cold War). Second, we disaggregate our GTI measure by individual states to see whether it accurately "post-dicts" which great powers posed a discernible threat to US security, and when, and correctly identifies periods when fears about the strategic importance of smaller states (e.g., Vietnam and Cuba) are evident in news media accounts of the period. Finally, we compare and contrast our GTI with other quantitative indices that treat defense spending, interstate disputes, alliance portfolios, and UN voting behavior as proxies for foreign threat to show where the GTI measure and other quantitative indices converge and diverge in assessing America's international threat environment over time.

The paper is organized into three sections. In the first section, we briefly review the main conceptual approaches in international relations to identifying and measuring foreign threats. The second section describes the semi-supervised machine learning model we use to distinguish threatening from non-threatening states and to measure and track variations in the intensity of foreign threats over time. In the third section, we describe the results and assess them using the three tests or benchmarks mentioned above. We show that our model performs well in each of these tests: variations in our GTI measure correspond to well-known events and periods of relative security and insecurity in US history; the GTI index clearly distinguishes between threatening and non-threatening countries, and recognizes that over time some of America's friends have become foes, and vice versa; and the GTI index is more granular in how it assesses threats, and responsive to sudden shifts in the international environment, than latent-threat indicators. We

conclude by discussing how semi-supervised machine learning models can be used to exploit the full potential of newspaper and other text-based data that international relations scholars rely on to understand political leaders' foreign policy choices.

Geopolitical Threat Index

The idea that states and leaders worry first and foremost about security is central to many theories of international politics (Trubowitz 2011). It is thus not surprising that international relations scholars have invested a great deal of time and effort in developing and testing measures of international threat. Many scholars model threats indirectly based on capability or behavioral indicators such as patterns of military spending (Nordhaus, Oneal, and Russett 2012), the propensity to engage in militarized interstate disputes (MIDs) (e.g., Bennett 1997), and states' overall foreign policy ideological orientation (Bueno de Mesquita 1981; Signorino and Ritter 1999). Other international relations scholars emphasize the importance of subjective factors, especially elite perceptions of potential adversaries' power and ambitions (e.g., Jervis 1976; Wohlforth 1993; Yarhi-Milo 2013). These scholars draw heavily on textual materials (e.g., newspaper articles, diplomatic cables, personal papers, and parliamentary proceedings) to reconstruct political elites' views about foreign states' capabilities and intentions, and to determine how widely those views are shared by intelligence analysts, elected officials, and other opinion makers (e.g., journalists and business leaders).

Each approach has strengths and weaknesses. Capability and behavioral indicators provide insight into states' latent or potential ability to launch an attack or engage in coercive diplomacy, and are generally quantifiable. They thus permit statistical, cross-national analysis, and provide a large universe of cases (observations) for testing theories about the possible effects of external threats on alliance formation, arms races, the use of force, and other international behavior. However, these indicators are less helpful for gaining insight into when and how political elites "code" these international signals as actual, manifest threats to the national interest. Diplomatic cables, intelligence reports, newspaper stories, and parliamentary proceedings, which provide a written record of political leaders', policymakers', and opinion makers' views of foreign actors' intentions, can shed considerable light on these questions. However, in the past it was not feasible to factor elite opinions about threats into large-N research in systematic and parsimonious ways.¹ The tools needed to convert such texts into usable form for statistical analysis did not exist (Benoit 2019).

Rapid advances in computational text analysis have now made it possible for international relations scholars to mine text sources for elite and mass opinion and sentiments. Indeed, today there is a growing body of empirical work that employs these methods to study a wide variety of political phenomena. Here, we show how computational text analysis can be used to model opinion makers' characterization of foreign actions and policies as friendly or hostile, cooperative or confrontational, and peaceful or belligerent. We restrict our analysis to states here and follow Wallander and Keohane (1999) in defining threats as a positive probability that one state has the capability and intent to harm the security of another state. In the pages below, we rely on the

¹Yarhi-Milo (2013, 9) analyzed more than 30,000 archival documents and intelligence reports in her analysis of three case studies of elite assessments of adversary's intentions.

judgments of *The New York Times* (*NYT*) editors and foreign correspondents about states' material power, international behavior, and strategic intentions.

The *NYT* has one of America's largest international reporting divisions, regularly reporting on other nations' international ambitions and diplomacy, as well as their military and capabilities. Boasting the highest circulation of American metropolitan newspapers, the *NYT* is widely considered the most authoritative source of international (and domestic) news in the United States. It commands the attention of political, government, and business leaders (Chernomas and Hudson 2015) and enjoys similar standing internationally, where it is considered "mandatory first reading in newsrooms across the world" (Dell'Orto 2013). The "newspaper of record," the *NYT* is also an important intermedia agenda setter in the United States at the national, regional, and local levels. As one of the only newspapers in the world with digital archives extending back to the nineteenth century, the *NYT* provides a single continuous source of news reporting.²

Computational Text Analysis

A variety of approaches and algorithms are available for large-scale computational text analysis. Broadly speaking, these approaches differ in terms of how much human supervision (i.e., manual coding) they require (Nelson 2017; Benoit 2019). At one end of the spectrum are fully supervised machine learning models, which seek to generate inferences or predictions about some universe of texts using a "top-down" approach. The user attaches *ex ante* labels to a sample set of training texts. The machine uses the coding of the training texts to identify matches in the larger corpus of texts that has not been manually coded. At the other end of the spectrum are unsupervised machine learning approaches. These models do not rely on user-provided labels to help the machine to learn. Instead, the machine itself exploits differences in textual features in documents to create clusters or topics that the analyst then interprets drawing on their knowledge of clusters that emerge from the "bottom-up." Topic modeling is a well-known example of the unsupervised machine learning approach to text analysis (Roberts et al. 2014).³

Semi-supervised machine learning models lie in between these two general approaches.⁴ These models try to exploit the comparative advantages of both supervised and unsupervised approaches. Supervised learning models' main advantage is that they produce results that relate directly to the analyst's substantive interests. This is because these models learn based on information (labels) provided *ex ante* by the analyst. One major disadvantage is that they are impractical for long historical analyses due to the high cost of training a model.⁵ Unsupervised models suffer from the reverse problem: they are very cost-effective at classifying or scaling large

numbers of documents, but they require the analyst to impose *ex post* interpretations on results and often yield topics of little theoretical interest.⁶ Given our interest in identifying variations in geographically specific threats spanning a long time period and hundreds of thousands of news stories, we adopted a semi-supervised approach.

Classification and Steps

The construction of our GTI involved five phases that are summarized in Figure 1: data (article) collection, text pre-processing, geographic and thematic classification, data filtering, and finally, threat (GTI) scaling. The first collection phase involved culling and condensing all *NYT* news stories involving US security. We used the *NYT* Application Program Interface (API) to identify relevant stories in the newspaper's digital archival database. Searching the API using a simple Boolean query made up of keywords, we downloaded lead sentences summarizing the news stories.⁷ The corpus of news summaries ($N = 387,896$) was then pre-processed using *Quanteda*, an R-package for segmenting texts into "tokens" (words) and creating a statistically usable document-feature matrix from the frequency of tokens (Benoit et al. 2018).⁸ We then classified the pre-processed text data using two semi-supervised machine learning models in parallel: *Newsmap* and *LSS*.

Newsmap is a geographic document classification technique designed to identify the primary geographic (country) focus of news stories (Watanabe 2018). Unlike fully supervised models, *Newsmap* does not require a manually coded set of training texts. It relies instead on a dictionary of geographic seed words comprised of country and city names. The model identifies associations between words and places based on co-occurrences with the seed words.⁹ Using this historical geographical lexicon, we then trained a *Newsmap* model to determine the principal regional and country focus of the news summaries of *NYT* stories. News summaries that *Newsmap* classified as primarily about events and developments *inside* the United States were removed from the analysis.

The same corpus of pre-processed news story summaries was submitted to a semi-supervised thematic categorization model called *LSS*.¹⁰ *LSS* also relies on a small set of user-provided polarity words as "seed words," that is, words that

²See Eshima, Imai, and Sasaki (2020) for an analysis of the relative advantages of semi-supervised machine models over unsupervised models. As they show, adding only a few seed words substantially improves the performance of unsupervised topic models.

³The backend of the *NYT* API system searches the full text for these keywords, but searches only return news summaries' lead sentences to users. Our Boolean search included the following keywords: ("military" OR "soldier*" OR "air force" OR "navy" OR "army") AND ("threat*" OR "danger*" OR "fear*" OR "risk*").

⁴The selection phase of the data analysis reduced the average number of tokens (words) per text document in the corpus from 35.8 to 17.5. To remove the function words, we used the list of stop words available on CRAN (<https://cran.r-project.org/web/packages/stopwords/index.html>).

⁵Given the duration of the study period (1861–2017), we added historical names (e.g., Ottoman Empire; Leningrad) so that news summaries from earlier periods in US history could be classified with greater accuracy.

⁶*LSS* is based on a technique known as Latent Semantic Analysis (Deerwester et al. 1990). It performs singular vector decomposition (SVD) of a document-feature matrix to estimate the semantic proximity of words. Mathematically, the method is similar to other dimension-reduction techniques such as factor analysis. SVD is first applied to a document-term matrix D ,

$$D \approx U \Sigma V'$$

where V' contains word vectors, $v_1 \dots v_f$, for all the words. Then the average similarity, $g_1 \dots g_f$, between word vectors for seed words, $s \in S$, is computed to create

²To our knowledge, only three US newspapers' digital archives extend back to the nineteenth century: *The Chicago Tribune*, *The New York Times*, and *The Washington Post*.

³Many scholars have used topic modeling to study comparative politics. See, for example, Baerg and Lowe (2020); Blaydes, Grimmer, and McQueen (2018); Hagemann, Hobolt and Wratil (2017); and Pan and Chen (2018).

⁴See Watanabe (2020 and 2018) for a detailed description of semi-supervised machine learning, including the two semi-supervised models used here: *Newsmap* and *LSS*.

⁵Supervised models must "see" words multiple times to "understand" (estimate the parameters for them). Because substantively important words occur very infrequently, hundreds and sometimes thousands of documents must be labeled manually to train the model.

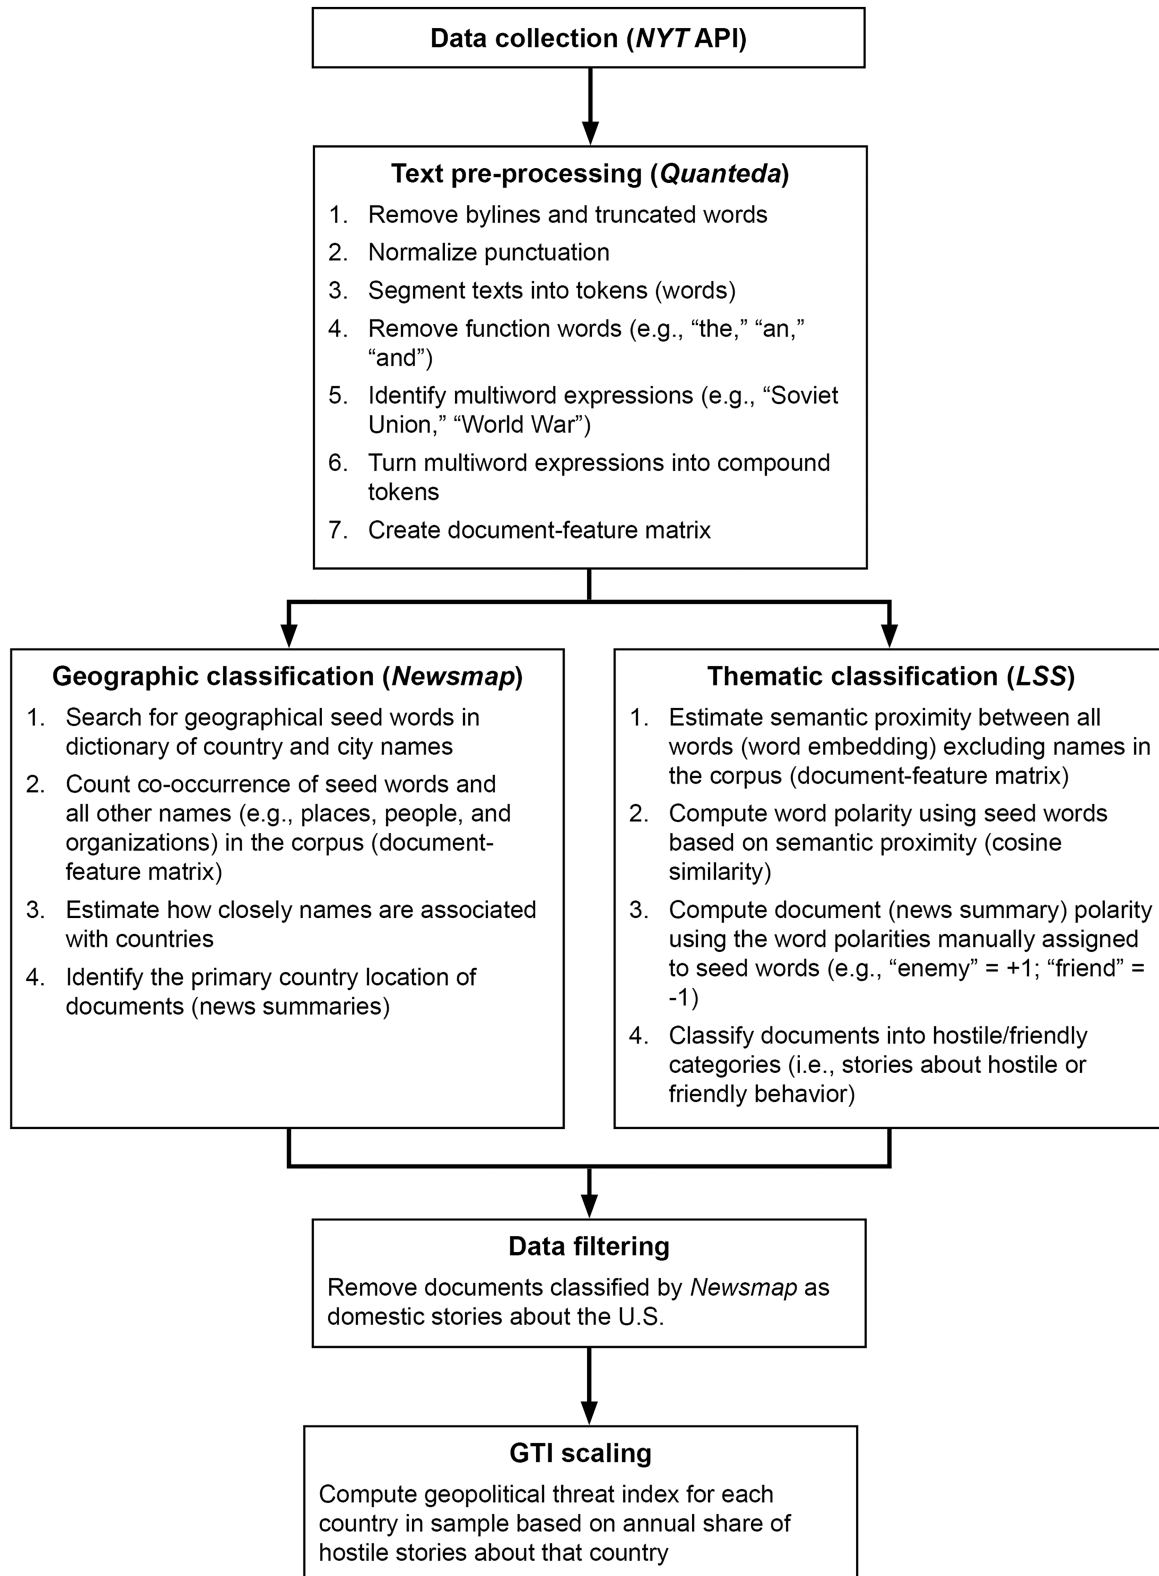


Figure 1. Phases of semi-supervised classification of text.

are positively or negatively associated with the substantive issues or topic under consideration (Watanabe 2020). The

polarity scores for all of the words p_s :

$$g_f = \frac{1}{|S|} \sum_{s \in S} \cos(v_s, v_f) p_s,$$

seed words are then used to assign polarity scores to other words in the corpus of texts (news summaries). The process

where $\cos(v_x, v_y)$ is the cosine similarity between word vector of words x and y . Following Turney and Littman (2003), we set the document dimension to 300 in the SVD.

in phase three (geographic and thematic classification) involves three steps. In the first step, the semantic proximity between all words in the corpus is computed using a word embedding technique to estimate cosine similarity between the words. These words are then assigned polarity scores based on their proximity to the user-defined seed words. Finally, a polarity score for each document is computed by weighting the individual words' polarity scores by how frequently they appear in the document. We take the average of the words' polarity scores for each document.¹¹

We used *LSS* to generate polarity scores for each word in *The New York Times* corpus of news summaries using seed words related, positively and negatively, to “hostility.”¹² Our positive (+1) seed words were “adversary,” “enemy,” “foe,” and “hostile,” including plural forms of each word. Our negative (−1) seed words were “aid,” “ally,” “friend,” and “peaceful,” and again, their plural forms. Figure 2 summarizes the polarity scores for each word in the corpus (horizontal axis) and their frequency (vertical axis). We have highlighted a sample of words that are either positively or negatively associated with hostility. Words close to zero on the horizontal axis are neutral, in the sense that they are at the same semantic distance from seed words on the right side and the left side of the space.¹³ The location of most of the words on the right or left side of the neutral point on the horizontal axis makes intuitive sense.¹⁴

In Figure 3, we plot the individual polarity scores of documents by year. The circles in the figure represent the polarity scores for 10,000 stories randomly drawn from the full sample of 387,896 stories. News summaries with a score of more than zero are stories about hostile or “antagonistic” behavior. News summaries with a score of less than zero are stories about non-hostile or “friendly” behavior. We have plotted the moving average.¹⁵ As Figure 3 indicates, the level of hostile behavior facing the United States varies considerably over time. During the so-called Awkward Years in the 1870s and 1880s, when America was on the periphery of the world economy and had few international ambitions, it faced comparatively few threats to its interests, internationally (Pletcher 1962). By contrast, *NYT* reporting of hostile behavior increases sharply during World War I and before and during World War II. For most of the Cold War, the polarity score in Figure 3 remains above zero. Not surprisingly, the main exception occurs during the era of détente in the 1970s, when tensions between the United States and the USSR abated and Washington began to normalize relations

¹¹We compute the polarity score of a summary article k , which comprises of features F , by taking the sum of the polarity scores g_f weighted by frequency of words h_f :

$$k = \frac{1}{N} \sum_{f \in F} g_f h_f,$$

where N is the total number of words in the document.

¹²Initially, we tried to measure threat directly using threat-related seed words in the text. However, we found that we achieved greater accuracy by aggregating the number of stories referring to *hostile* state behavior by year. This may be because measuring threat directly requires more information about the context of events than the short *NYT*ADI news summaries contain.

¹³Thus, frequently used words like “war,” “peace,” and “military,” which appear in the center of Figure 2 tend to have small polarity scores because they appear as often in news stories about allies and friends as they do in stories about adversaries and foes.

¹⁴Without access to the complete articles behind the API summaries, we cannot determine the precise accuracy of the polarity scores for words in the corpus. However, *LSS*'s ability to classify stories about hostile and non-hostile state behavior offers some degree of confidence, as does the intuitive placement of extreme and neutral words in the two-dimensional space in Figure 1.

¹⁵To make the results easier to interpret visually, the polarity scores are normalized by the standard deviation and centered around the global mean.

with Beijing. The polarity scores then increase in the 1980s, during the so-called Second Cold War. Since the mid-1990s, the polarity score averages less than zero, as one would expect during an era of “unipolarity” when the United States faced no peer rival.

The last two phases in the process described in Figure 1 involved data filtering and constructing the GTI. The data filtering stage removes news summaries that *Newsmap* classifies as principally about the United States. We are only interested in news stories about foreign nations.¹⁶ We then used these news stories to compute our GTI index for all foreign nations in our sample (225 in total over the study period) and separate indices for each of the following fifteen countries: Afghanistan, Britain, Canada, China, Cuba, France, Germany, Iran, Iraq, Japan, Mexico, Russia, Spain, Syria, and Vietnam. The GTI index refers to the proportion of summaries that are about a country's hostile military actions out of the total number of news stories about that country's military behavior in the year. We assume that the annual number of news articles published by the *NYT* about hostile foreign policy behavior and policies of each country was proportional to the level of perceived threat it posed to the United States in that year.

Validity Checks

We also conducted two validity and sensitivity checks to see how well *Newsmap* and *LSS* coded the *NYT* news summaries. The first check was a manual one, involving two steps. First, we randomly selected fifty foreign news summaries per decade over the study period (850 summaries in total) and manually classified each story by the principal country the story was about.¹⁷ We then compared the machine's coding of stories against our own coding of the same stories. *Newsmap* correctly classified the news summaries 76 percent of the time. Next, we compared *LSS*'s classification of the news summaries as hostile or peaceful behavior to our own assessment of the summaries' content. Following Young and Soroka (2012), we aggregated the classification results by decade and country because our primary interest is measuring threat levels annually for countries. As Figure 4 indicates, the proportion of hostile summaries computed for each decade ($r = 0.82$) and country ($r = 0.94$) was strongly correlated.¹⁸

We also wanted to check the reliability of the API summaries for classifying stories substantively. Although an invaluable source of information for historically oriented research, the API only provides short summaries of news articles. To determine whether the summaries contain sufficient information to correctly classify the story in terms of country focus and level of hostility, we collected a sample of full-text *NYT* articles ($N = 3,890$) from the Nexis database between 1980 and 2018, using the same Boolean

¹⁶Domestic stories that are filtered out account for roughly 40 percent of the total number of news summaries. These stories either contain names strongly associated with the United States or entirely lack names associated with foreign countries.

¹⁷Since the corpus has many domestic news summaries, we first separated domestic and foreign news summaries using the *Newsmap* classifier. We confirmed that the machine correctly distinguished between foreign and domestic news summaries 77 percent of time.

¹⁸Roughly 59 percent of the individual summaries classified as stories about hostile state behavior agreed between machine and manual classification. The level of agreement between computer and manual classification on the individual summaries is typically lower because human coders classify texts based on their understanding of historical context. By contrast, machines code solely on the basis of the words that appear in the texts.




Figure 2. Distribution of polarity scores and word frequencies.




Figure 3. Polarity scores for *The New York Times* news summaries, 1861–2017.

query we used to collect the API news story summaries.¹⁹ Using *Quanteda*, we segmented the full-text articles into segments of twenty tokens each, the average length of an API summary. We then processed the tokens using the *Newsmap* dictionary and the *LSS* model described above. The results are reported in Figure 5. It shows that while the median length of full-text articles in the *NYT* is 612 tokens, the first twenty-token segment summaries contain more of the key information about threats to the United States than the remaining twenty-token segments that make up the story. Countries that are the focus of the news stories are mentioned, on average, 1.5 times more in the first twenty-token

¹⁹ Unfortunately, the API summaries do not provide access to the original article to test for robustness. As a result, we relied on Nexis to collect the original, full-text articles.

segment. In addition, the polarity scores are higher in the first twenty-token segment than in the rest of the news story. Despite their short length, the *NYT* API summaries offer valuable information on both the source and level of threat facing the United States.

Finally, we ran a simple keyword frequency analysis as a sensitivity check. While we are principally interested here in the geographic location of perceived threats, we wanted to know whether the word frequency of well-known ideologies (e.g., communism, fascism) and warfare strategies (e.g., trench warfare, guerilla warfare) matched what one might expect, given the history of the United States, the rise and fall of different ideologies internationally, and changes in military technology and warfare. Figure 6 summarizes the normalized frequency of keywords for ideology and warfare by year from 1861 to 2017. The trends square with




Figure 4. Level of agreement between manual and *LSS* classification, by decade and country.




Figure 5. Position of countries and hostile actions mentioned in full-text articles.

common knowledge. We see in the top panel that “fascism” appears frequently in the *NYT* news stories during the 1930s and 1940s, while “communism” increases significantly in the 1940s and 1950s. “Imperialism” occurs frequently before World War I and peaks during the surge of anti-colonial independence movements after World War II. The keyword “terrorism” increases in the 1970s and 1980s and especially in the 2000s, following the September 11 attacks.

The trends we see in the lower panel in Figure 6 also square with the history of warfare and military strategy. We would expect the keyword “trench” to register a high frequency during World War I and then to drop off as trench warfare became an anachronism. As Figure 6 indicates, the term “Blitzkrieg” strategy is closely associated with World War II. While guerilla warfare has a long history, in the modern era the keyword “guerilla” is closely associated with the Cold War (e.g., Vietnam, Angola, Nicaragua, and Afghanistan). The keyword “nuclear” appears after the bombings of Hiroshima and Nagasaki. Its frequency varies over time, surging during periods when fears of nuclear war were high (e.g., Cuban Missile Crisis; US–Soviet tensions in early 1980s) as well as during periods of heightened concern about nuclear proliferation (e.g., Iraq, Iran, and North Korea in early 2000s). Given the growing concern about “cyber

warfare” in the United States (and elsewhere) in the past decade, it is not surprising that the term appears with great frequency in the 2010s.

America’s Changing Threat Environment

In this section, we evaluate our GTI in three different ways. We begin by considering how well the overall patterns of elevated threat gleaned from *NYT* news summaries for the full sample of 225 foreign nations. To assess robustness, we then disaggregate GTI by a sample of fifteen countries to judge how well it captures the ebb and flow of US relations with different nations: from those that the United States has generally enjoyed friendly relations with to those that have been America’s foes at one or more points in its history. These countries account for 43 percent of total foreign news coverage over the study period. Finally, we consider how well changes in elite assessments of manifest threat correlate with latent changes in the material capabilities and international behavior of the fifteen countries in our sample.

Aggregate Threat Patterns

Figure 7 summarizes the total GTI score for the United States for all foreign nations. The overall pattern of high




Figure 6. Frequency of keywords for international ideologies and military strategies.

and low threat corresponds closely to historical accounts of America's foreign relations (e.g., [Herring 2008](#)). During much of the nineteenth century, for example, when the United States faced few threats to its security from abroad, the threat index is relatively low. By contrast, our text-generated threat index is generally higher in the twentieth century, especially during periods when the United States confronted rising powers: Hitler's Germany, Imperial Japan in the 1930s and 1940s, and the Soviet Union during the Cold War. During the Cold War (1950–1991), the mean GTI score is 0.50. By contrast, from 1861 to 1913, the mean GTI score is 0.45. We see too that in wartime the threat index spikes (e.g., the Spanish-American War, World Wars I and II, the Persian Gulf War). Following the cessation of hostilities, it decreases rapidly.

GTI is also able to distinguish between gradual as well as sudden changes in the international threat environment. The mounting challenge posed by German and Japanese power during the 1930s is reflected in GTI's steady rise in [Figure 7](#), as are shocks associated with key events such as the sinking of the *Lusitania* in 1915, which presaged the US declaration of war two years later, and the September 11, 2001 attacks that led to America's wars in Afghanistan and Iraq. Within-period variation is also evident in the figure. During the Cold War, for example, the index is elevated in the 1950s and 1960s, when the rivalry was especially intense. It is considerably lower in the 1970s, when tensions eased during the era of US–Soviet détente.

America's Friends and Foes

We also created separate GTI measures for each of the fifteen countries in our sample. We divided the countries into two groups: states typically classified by international relations scholars as great powers for some or all of the past 150 years ([Snyder 1991](#); [Mearsheimer 2001](#)) and smaller states that have been a source or location of a perceived threat to US interests. Our list of great powers includes Britain, China, France, Germany, Japan, Russia (Soviet Union), and Spain. The smaller states were subdivided into two groups: those located or having interests in the Western Hemisphere (Canada, Cuba, and Mexico) and states located in other parts of the world (Afghanistan, Iraq, Iran, Syria, and Vietnam). In the late nineteenth century, when America was still a regional power, US policymakers were particularly concerned about states in the Western Hemisphere. Since America's emergence as a global power after World War II, states located in other parts of the world have been of growing concern to US policymakers.

We begin with Germany and Japan, once former adversaries of the United States and now among America's closest allies. [Figure 8](#) tracks changing perceptions of each country since the Civil War. There are few, if any, surprises here. In the case of Germany, the threat index rises and falls as one might expect: it spikes during World War I and again, during World War II. In other periods, the threat index is very low. This too conforms to expectations. Germany was




Figure 7. Baseline GTI for United States, 1861–2017. (Based on full sample of 225 countries and kernel smoothed by ± 1 year.)

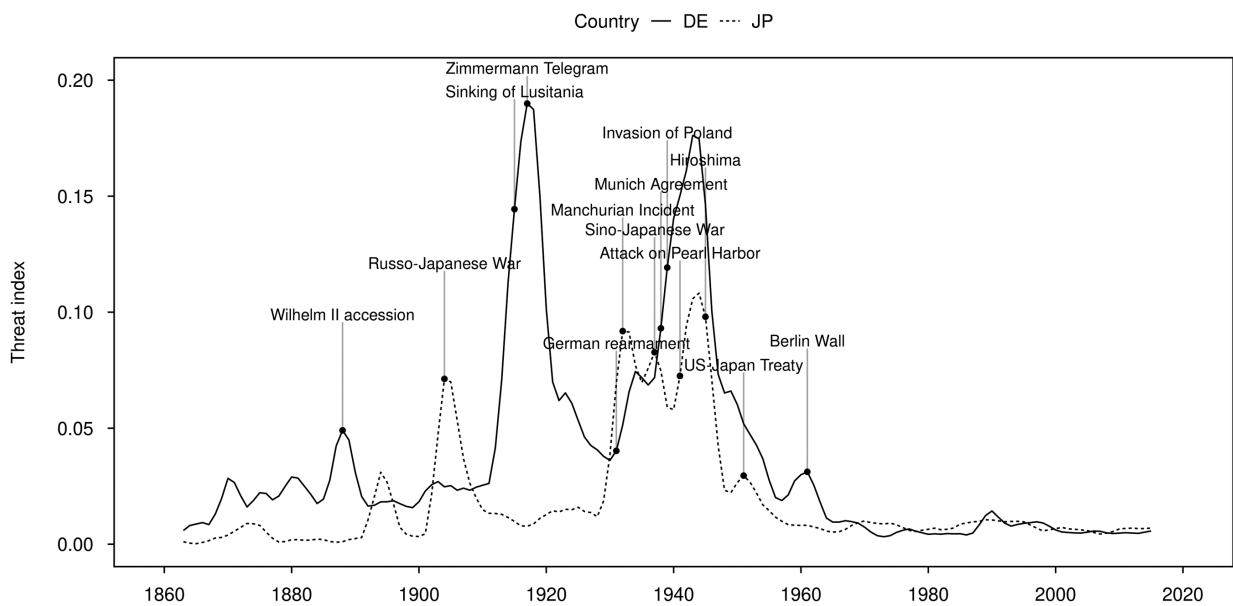


Figure 8. Threats to US interests by country: Germany and Japan, 1861–2017.

of little concern to the United States during the nineteenth century (Jonas 1985). During the Cold War, West Germany was closely allied to the United States. Indeed, from the mid-1960s to the current era, the GTI for Germany is close to zero.

The pattern we see in Japan’s case is similar. Here too America’s worries about Japanese power at different historical junctures are captured by the threat index. We see clear spikes in the GTI during the Russo-Japanese War (1904–1905) and Japan’s invasion of Manchuria in 1931, both of which caused concern in the United States (Green 2017). We also see a spike with Japan’s surprise attack at Pearl Harbor (1941). Meanwhile, Japan’s low threat score in the nineteenth century is consistent with scholarship on US–Japanese relations; relations were generally good (Green

2017). Since signing the Treaty of Peace with Japan (1951), formally ending the war and America’s occupation of Japan, and the US–Japanese Mutual Security Treaty (1951), Japan scores very low on the threat index. This too conforms to expectations. Like Germany, Japan has been considered a trusted ally and friend for decades by the vast majority of Americans.

Figure 9 focuses on China and Russia—two former adversaries of the United States who many Americans now view as competitors and potential adversaries. Of the two nations, Russia looms as the larger threat to American interests over most of the past 150 years. This is especially clear during the long Cold War, when there is sizable gap in the GTI between Russian and China. To be sure, there are periods when concerns about China eclipse worries about Russia.




Figure 9. Threats to US interests by country: China and Russia (Soviet Union), 1861–2017.

Most of these have to do with developments *inside* China: Boxer Rebellion (1899–1901), Xinhai Revolution (1911), the Chinese Civil War (1927–1949), and the Chinese Revolution (1949). More surprising, is the relatively low threat score for China, especially during the 1950s and 1960s when Beijing was widely considered an adversary in the United States. One plausible explanation for its comparatively low score is that US opinion makers considered Beijing to be the junior member of the “Sino-Soviet bloc.”²⁰

Concerns about Russian geopolitical ambitions have long influenced US foreign policy and likely explain the spikes we see during Russo-Japanese War of 1904–1905, the 1917 Russian Revolution, and the time of the signing of the 1939 Molotov–Ribbentrop Pact between Nazi Germany and the Soviet Union. During the Cold War, Soviet technological breakthroughs (the 1949 Soviet atomic blast of 1957 Sputnik satellite launch) set off alarm bells in the United States. So did the Soviet invasions of Czechoslovakia (1968) and Afghanistan (1979). In both cases, the GTI rises sharply. The reverse occurs in periods of reduced tensions: following the signing of the Limited Test Ban Treaty (1963) and the US–Soviet détente in the 1970s, the threat index decreases. We also see the index plummet when the Soviet Union collapses. This more positive view of Russia lasts until 2014, when Moscow invades the Ukraine and annexes Crimea.

Figure 10 focuses on two long-time allies of the United States: Britain and France. Overall, the level of threat posed by London and Paris to American interests is comparatively low. Concerns about British and French ambitions and power do surface from time to time. In the case of Great Britain, the threat index spikes on several occasions in the late nineteenth century and at the outbreak of World War I and World War II.²¹ In the case of France, the index surges in 1870, when America favored Germany over France in the

Franco-Prussian War (Schieber 1921), again in 1923 in opposition to France’s occupation of the German’s Ruhr Valley, and in 1954, when American-backed French forces went down to military defeat in Indochina.

Finally, Figures 11 and 12 track the threat index for our two groups of smaller states. Figure 11 includes countries in the Western Hemisphere. Figure 12 covers countries located in other parts of the world, principally the Middle East. Our expectation is that the GTI scores for the first group will be considerably higher in the nineteenth century, when America still defined its vital interests in regional terms. Conversely, those in the second group should be of greater concern to the United States since World War II, when the United States becomes a global power with far-reaching interests. This is what we see. The only exception is Cuba, which becomes a source of great concern in the United States following the Cuban Revolution and the establishment of a new government led by Fidel Castro in 1959.

Comparison of GTI and Other Threat Measures

In this section, we compare our GTI measure of elite threat perception to other measures of foreign threat. Most existing quantitative measures of foreign threat rely on indicators of material capability (e.g., military spending), international behavior (e.g., the propensity to threaten or use military force), or overall foreign policy orientation as proxies for foreign threat. Some well-known examples include Nordhaus, Oneal, and Russett’s (2012) Liberal-Realist Model (LRM) and Leeds and Savun’s (2007) measure of threat, as well as the Composite Index of National Capability (CINC) (Singer, Bremer, and Stuckey 1972) and MID (Palmer et al. 2015) measures of relative power and interstate conflict.²² Bailey, Strezhnov, and Voeten’s (2017)

²⁰ In the absence of full *NYT* news stories, it is hard to be certain.

²¹ As an ally of the United States, Britain’s elevated GTI in World War II appears to reflect concerns about Britain’s security in the face of German threats and attacks. We manually checked stories about Britain from 1939 to 1941. There is noticeable increase in stories on “The Blitz,” the German bombing campaign against the UK, including what Britain’s ability to defend itself means for the United States. This appears to explain the surge in Britain’s GTI score in 1940.

²² For each of these five measures, we constructed a threat index for the United States. LRM is the predicted probability that a country will become involved in militarized disputes based on a country’s relative power, distance from other countries, alliance relationships, its degree of democracy or autocracy, and the degree of integration into the international system. It covers the period from 1951 to 2001. MIDs refer to instances where the United States was involved in




Figure 10. Threats to US interests by country: Britain and France, 1861–2017.

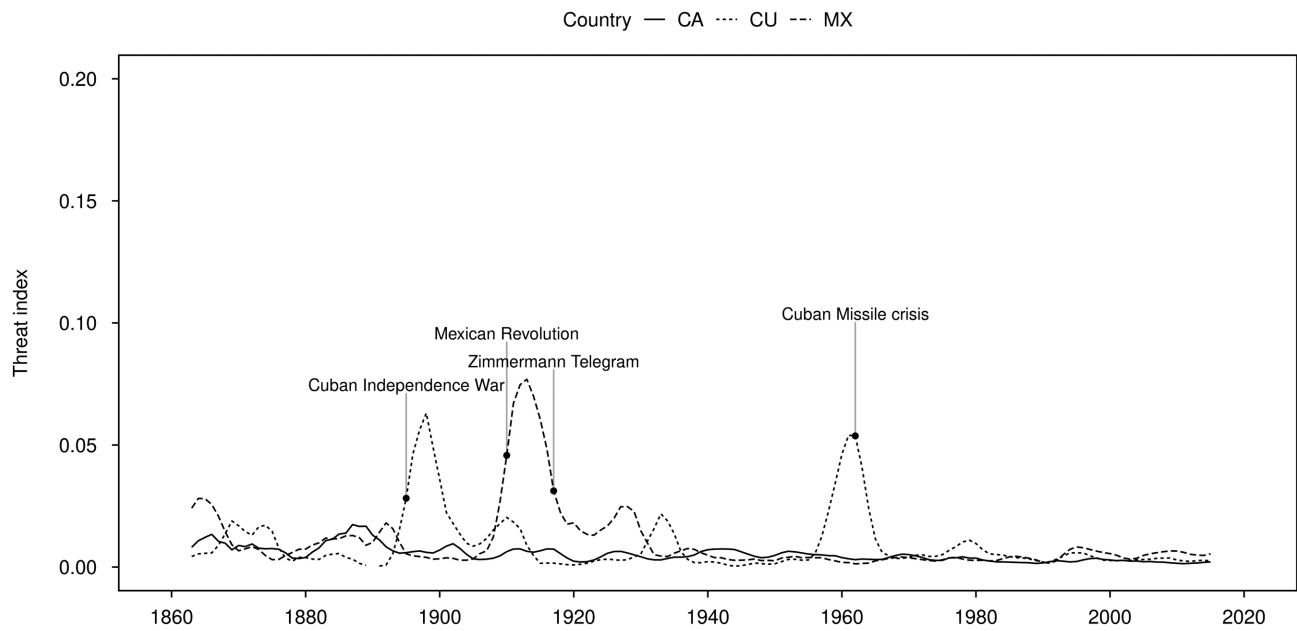


Figure 11. Threats to US interests by country: Canada, Cuba, and Mexico, 1861–2017.

measure of UN voting dissimilarity is another measure that can be used to model state’s threat environment.²³ We compare our GTI index to these five measures using Pearson’s correlation coefficients. The results are summarized in Tables 1 and 2, by major power and small states, respectively.

a militarized dispute short of war with one of these fifteen countries. We use version 4.0. CINC is the Composite Index of National Capability of the countries and a standard measure of national power. We use version 5.0. LS refers to the CINC score for non-US allies in the ATOP database. We include only countries that have an S score with the United States that is less than the median value of 0.458. UN is the ideal point distance from the United States measured by voting patterns at the United Nations.

²³We use the ideal point distance to calculate UN voting dissimilarity scores between the United States and the fifteen countries in our sample.

Table 1. Correlation analysis of GTI index of foreign threat for the United States, by major power

Country	MIDs	CINC	LS	UN	LRM
Britain	0.322***	0.353***	0.292***	-0.345**	0.709***
China	0.044	-0.156	-0.129	0.198	0.347*
France	0.102	0.502***	0.408***	-0.235*	0.698***
Germany	0.654***	0.586***	0.570***	-0.257	0.348*
Japan	0.380***	0.424***	0.307***	-0.230	0.068
Russia	0.168	0.442***	0.356***	0.358**	0.144
Spain	-0.148	0.354***	0.312***	0.001	0.149
Average	0.217	0.358	0.302	-0.073	0.352

* $p < .5$, ** $p < .01$, *** $p < .001$.

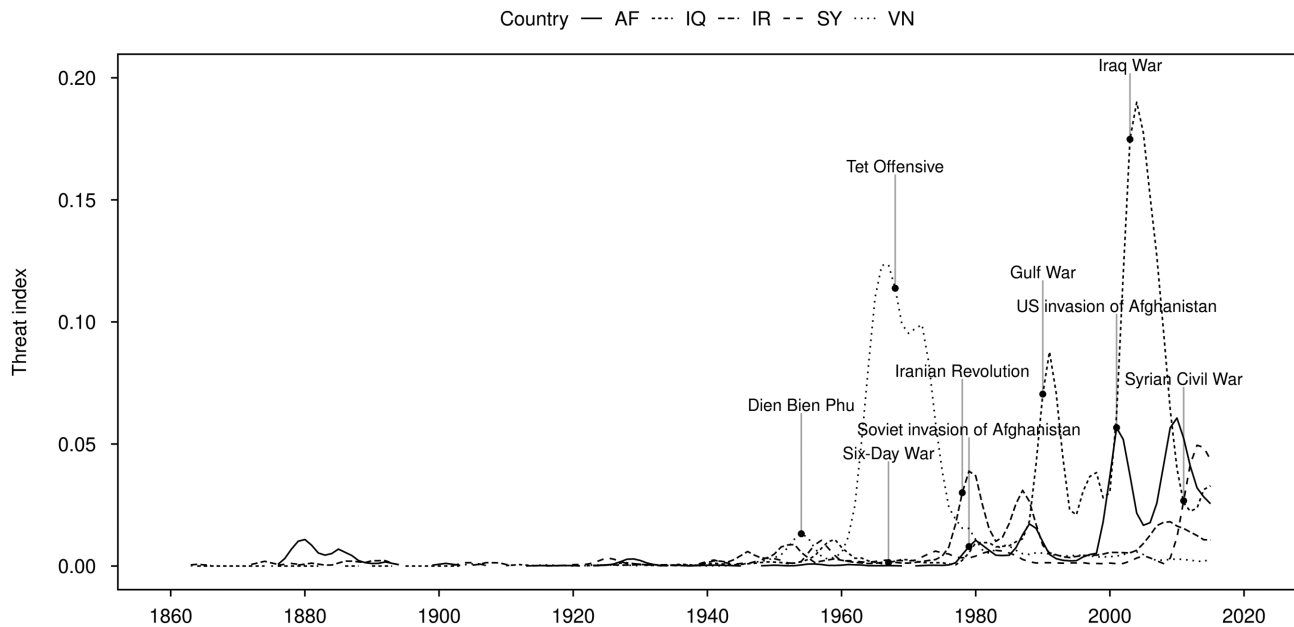


Figure 12. Threats to US interests by country: Afghanistan, Iraq, Iran, Syria, and Vietnam, 1861–2017.

Table 2. Correlation analysis of GTI index of foreign threat for the United States, by small states

	<i>MIDs</i>	<i>CINC</i>	<i>LS</i>	<i>UN</i>	<i>LRM</i>
Afghanistan	-0.106	-0.129	-0.129	0.243*	0.045
Canada	0.002	0.357***	0.391***	-0.408***	0.462***
Cuba	-0.145	0.196*	-0.187	0.007	-0.459***
Iraq	0.007	0.356**	0.359**	0.346**	0.104
Iran	0.583***	0.433***	0.429***	0.244*	-0.119
Mexico	0.707***	-0.242**	0.359***	0.307**	-0.290*
Syria	-0.051	0.138	0.138	0.204	-0.350*
Vietnam	-0.220	-0.457***	-0.457***	-0.300	-0.090
Average	0.097	0.082	0.113	0.080	-0.087

* $p < .5$, ** $p < .01$, *** $p < .001$.

As Table 1 indicates, the correlation between GTI and the other measures of foreign threat varies considerably.²⁴ On average, GTI correlates most strongly with CINC ($r = 0.35$), LRM ($r = 0.35$), and LS ($r = 0.30$). GTI's correlation is much weaker for MIDs ($r = 0.21$) and UN ($r = -0.07$). We also see that the correlation between GTI and the other foreign threat indicators varies considerably by major power. The only exception is CINC, where we see a modest but consistent correlation with GTI between potential adversaries' material capabilities and American elite perceptions of threat.²⁵ The more major powers invest resources in the capacity to project power (e.g., strengthening armies; expanding military-industrial capacity), the more likely American political leaders and opinion makers are to view those investments as threatening to US interests.

The results in Table 1 are suggestive. One plausible interpretation of the patterns we see is that in gauging major powers' geopolitical intentions, American political elites and opinion makers find capabilities (e.g., military spending) a more credible (costly) signal of foreign states' inten-

²⁴ In general, the five other measures of foreign threat correlate weakly with each other. Results are available upon request.

²⁵ Among the major powers, the only exception is China.

tions than militarized disputes, alliance commitments, or public support or disapproval. This may be true for major powers, but it does not appear to be the case for smaller states. As Table 2 indicates, GTI correlates weakly, or even negatively, with CINC. This may be because most of these smaller states have been viewed by American policymakers as proxies or surrogates for one of the major powers (e.g., Vietnam for the Soviet Union during the Cold War). The coefficients for Afghanistan, Cuba, and Vietnam are negative for at least three of the five threat indicators. US perceptions of foreign threat would appear to have more to do with their status as chess pieces in larger great power competition rather than a result of independent policy decisions or actions.

Figure 13 compares the different foreign threat indicators temporally to see how well they capture well-known periods of high and low threat in American history. The upper panel in the figure runs from 1940 to 2017 because the LS, LRM, and UN indicators do not cover as much historical ground as the GTI measure, one of GTI's comparative advantages. We see in the ebb and flow of the three indicators, LRM is closest to GTI, visually confirming the correlation coefficients in Table 1. In the figure's lower panel, GTI and MIDs are the most similar visually, especially before World War II. Looking at the two panels together, it seems clear that the GTI index is capturing a dimension of threat assessment that is different from non-text-based measures of foreign threat.

Conclusion

In this research note, we have introduced a text-based strategy and method to identify and measure foreign threats. International relations scholars recognize the importance of incorporating elite perceptions of foreign threat into the analysis of international politics and foreign policy. However, before the advent of large-scale computational text analysis, it was not possible to exploit the full potential of available data sources for understanding threat perceptions. Semi-supervised machine learning models like

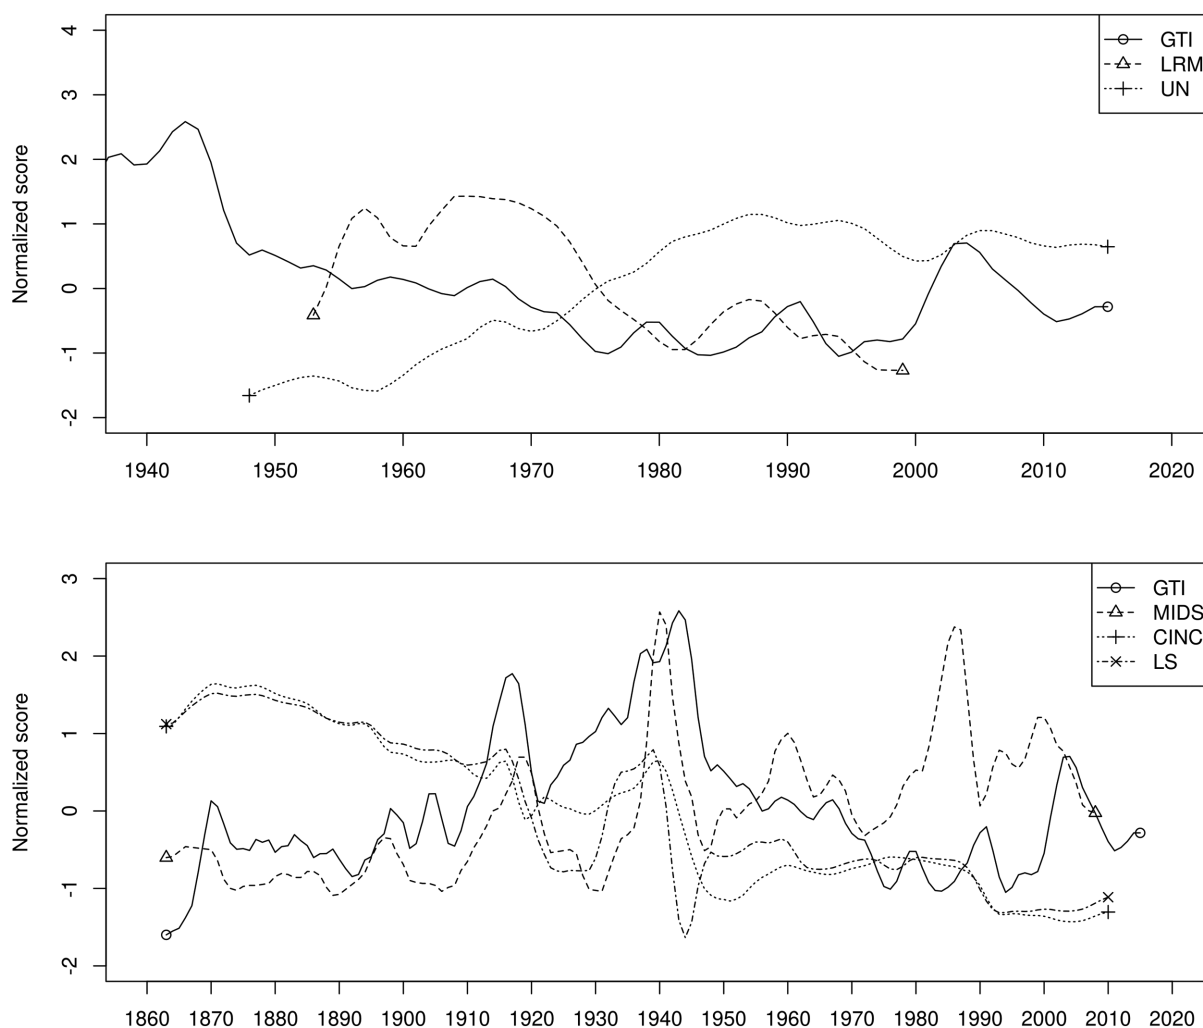


Figure 13. GTI and other measures of foreign threat-facing United States.

the ones used here offer a consistent, cost-effective way to systematically analyze how political leaders and opinion makers view their country's geopolitical circumstances. Using the American case for illustrative purposes, we trained a machine algorithm to parse *NYT* news summaries to generate a historical timeline measuring the ebb and flow of foreign threats to American interests since the Civil War. We selected states widely considered to be friends of the United States at some points and foes in other periods. We also included small countries as well as great powers, democratic as well as authoritarian countries, and distant as well as neighboring states.

Overall, the model efficiently distinguishes friend from foe and, importantly, adapts as elite opinion about America's international threat environment changes over time. In most instances, countries international relations scholars and diplomatic historians categorize as enemies in one era or another are captured by the GTI index. Meanwhile, countries categorized as enemies in one period were appropriately recategorized as friends in another period, and vice versa. Our machine learning text analysis model also discriminates between periods of "low threat," like the post-bellum period in America in the late 1800s, when US policymakers thought national security was plentiful, and other periods in American history such as the Cold War, when US policymakers operated on the assump-

tion that they had little room for error internationally. Finally, a comparison of various threat indicators suggests that a text-based approach to threat assessment captures facets missed by measures relying solely on states' capabilities (e.g., military spending), foreign policy behavior, or geographic proximity.

Our text-based approach to threat-assessment can be strengthened in several ways. As noted above, one limitation of our model is that it is based on *NYT* news summaries rather than full-text articles. This may help explain the anomalous cases (e.g., the relatively low threat score for China during the 1950s and 1960s). Another limitation is that our machine learning model is based on a single source of reporting. As more and more newspapers digitize their archives, it will be possible to expand the corpus of stories and control for possible variation in newspapers' international coverage as well possible regional (e.g., internal versus coastal) bias in US news reporting.²⁶ Because *Newsmap* and *LSS* can be easily adapted to other languages, it will also be possible to extend the analysis of threat percep-

²⁶The model could also be expanded to include other sources of online text data such as the *Foreign Relations of the United States* (FRUS). This would make it possible to see how similar media and government assessments of foreign threats are and how they might diverge (e.g., which is more forward-looking or reactive). FRUS is available at the History Lab, an interdisciplinary collective. See <http://history-lab.org>.

tion to other countries where major newspaper archives are available in digital form (e.g., Britain, China, Japan, and Israel).²⁷ Finally, in principle these models can be applied dyadically to study strategic interaction between states. Are threat perceptions as interdependent and mutually reinforcing as theories of alliance formation, security dilemmas, and great power transition suggest? Such questions are ripe for large-scale machine-readable text analysis of the type discussed here.

References

- BAERG, NICOLE, AND WILL LOWE. 2020. "A Textual Taylor Rule: Estimating Central Bank Preferences Combining Topic and Scaling Methods." *Political Science Research and Methods* 8 (1): 106–22.
- BAILEY, MICHAEL A., ANTON STREZHNEV, AND ERIK VOETEN. 2017. "Estimating Dynamic State Preference from United Nations Voting Data." *Journal of Conflict Resolution* 61 (2): 430–56.
- BENNETT, D. SCOTT. 1997. "Testing Alternative Models of Alliance Duration, 1816–1984." *American Journal of Political Science* 41 (3): 846–78.
- BENOIT, KENNETH. 2019. "Text as Data: An Overview." In *SAGE Handbook of Research Methods in Political Science and International Relations*, edited by Luigi Curini and Robert Franzese, 461–97. New York: SAGE.
- BENOIT, KENNETH, KOHEI WATANABE, HAIYAN WANG, PAUL NULTY, ADAM OBENG, STEPHAN MÜLLER, AND ADAM MATSUO. 2018. "Quanteda: An R Package for the Quantitative Analysis of Textual Data." *Journal of Open Source Software* 3 (30): 774. <https://doi.org/10.21105/joss.00774>.
- BLAYDES, LISA, JUSTIN GRIMMER, AND ALISON MCQUEEN. 2018. "Mirrors for Princes and Sultans: Advice on the Art of Governance in the Medieval Christian and Islamic Worlds." *The Journal of Politics* 80 (4): 1150–67.
- BUENO DE MESQUITA, BRUCE. 1981. *The War Trap*. New Haven, CT: Yale University Press.
- CHERNOMAS, ROBERT, AND IAN HUDSON. 2015. *Gatekeeper: 60 Years of Economics According to The New York Times*. New York: Routledge.
- DEERWESTER, SCOTT, SUSAN T. DUMAIS, GEORG W. FURNAS, THOMAS K. LANDAUER, AND RICHARD HARSHMAN. 1990. "Indexing by Latent Semantic Analysis." *Journal of the American Society of Information Science* 41 (6): 391–407.
- DELL'ORTO, GIOVANNI. 2013. *American Journalism and International Relations*. New York: Cambridge University Press.
- ESHIMA, SHUSEI, KOSUKE IMAI, AND TOMOYA SASAKI. 2020. "Keyword Assisted Topic Models." *Computation and Language* arXiv:2004.05964.
- GREEN, MICHAEL J. 2017. *By More than Providence: Grand Strategy and American Power in the Asia Pacific since 1783*. New York: Columbia University Press.
- HAGEMANN, SARA, SARA B HOBOLT, AND CHRISTOPHER WRATIL. 2017. "Government Responsiveness in the European Union: Evidence from Council Voting." *Comparative Political Studies* 50 (6): 850–76.
- HERRING, GEORGE C. 2008. *From Colony to Superpower: U.S. Foreign Relations since 1776*. New York: Oxford University Press.
- JERVIS, ROBERT. 1976. *Perception and Misperception in International Politics*. Princeton, NJ: Princeton University Press.
- JONAS, MANFRED. 1985. *The United States and Germany: A Diplomatic History*. Ithaca, NY: Cornell University Press.
- LEEDS, BRETT ASHLEY, AND BURCU SAVUN. 2007. "Terminating Alliances: Why Do States Abrogate Agreements?" *Journal of Politics* 69 (4): 1118–32.
- MEARSHEIMER, JOHN. 2001. *The Tragedy of Great Power Politics*. New York: Norton.
- NELSON, LAURA K. 2017. "Computational Grounded Theory: A Methodological Framework." *Sociological Methods and Research* 49 (1): 3–42.
- NORDHAUS, WILLIAM, JOHN R. ONEAL, AND BRUCE RUSSETT. 2012. "The Effects of the International Security Environment on National Military Expenditures: A Multi-Country Study." *International Organization* 66 (3): 491–513.
- PALMER, GLENN, VITO D'ORAZIO, MICHAEL KENWICK, MATTHEW, AND LANE. 2015. "The MID4 Dataset, 2002–2010: Procedures, Coding Rules and Description." *Conflict Management and Peace Science* 32 (2): 222–42.
- PAN, JENNIFER, AND KAIPING CHEN. 2018. "Concealing Corruption: How Chinese Officials Distort Upward Reporting of Online Grievances." *American Political Science Review* 112 (3): 602–20.
- PLETCHER, DAVID M. 1962. *The Awkward Years: American Foreign Relations under Garfield and Arthur*. Columbia, MO: University of Missouri Press.
- ROBERTS, MARGARET E., BRANDON M. STEWART, DUSTIN TINGLEY, CHRISTOPHER LUCAS, JETSON LEDER-LUIS, SHANA KUSHNER GADARIAN, BETHANY ALBERTSON, AND DAVID G. RAND. 2014. "Structural Topic Models for Open-Ended Survey Responses." *American Journal of Political Science* 58 (4): 1064–82.
- SCHIEBER, CLARA EVE. 1921. "The Transformation of American Sentiment toward Germany, 1870–1914." *Journal of International Relations* 12 (1): 50–74.
- SIGNORINO, CURTIS S., AND JEFFREY M. RITTER. 1999. "Tau-b or Not Tau-b: Measuring Alliance Portfolio Similarity." *International Studies Quarterly* 43 (2): 115–44.
- SINGER, J. DAVID, STUART BREMER, AND JOHN STUCKEY. 1972. "Capability Distribution, Uncertainty, and Major Power War, 1820–1965." In *Peace, War, and Numbers*, edited by Bruce Russett, 19–48. Beverly Hills, CA: Sage.
- SNYDER, JACK. 1991. *Myths of Empire: Domestic Politics and International Ambition*. Ithaca, NY: Cornell University Press.
- TRUBOWITZ, PETER. 2011. *Politics and Strategy: Partisan Ambition and American Statecraft*. Princeton, NJ: Princeton University Press.
- TURNER, PETER D., AND MICHAEL L. LITTMAN. 2003. "Measuring Praise and Criticism: Inference of Semantic Orientation from Association." *ACM Transactions on Information Systems* 21 (4): 315–46.
- WALLANDER, CELESTE A., AND ROBERT O. KEOHANE. 1999. "Risk, Threat, and Security Institutions." In *Imperfect Unions: Security Institutions over Time and Space*, edited by Helga Haftendorn, Robert O. Keohane and Celeste A Wallander, 21–47. New York: Oxford University Press.
- WATANABE, KOHEI. 2018. "Newsmap: A Semi-Supervised Approach to Geographical News Classification." *Digital Journalism* 6 (3): 294–309.
- . 2020. "Latent Semantic Scaling: A Semi-Supervised Text Analysis Technique for New Domains and Languages." *Communication Methods and Measures*. <https://doi.org/10.1080/19312458.2020.1832976>.
- WOHLFORTH, WILLIAM C. 1993. *The Elusive Balance: Power and Perceptions during the Cold War*. Ithaca, NY: Cornell University Press.
- YARHI-MILO, KAREN. 2013. "In the Eye of the Beholder: How Leaders and Intelligence Communities Assess the Intentions of Adversaries." *International Security* 38 (1): 7–51.
- YOUNG, LORI, AND STUART SOROKA. 2012. "Affective News: The Automated Coding of Sentiment in Political Texts." *Political Communication* 29 (2): 205–31.

²⁷ *Newsmap* is currently available in the following languages: Arabic, Chinese (traditional and simplified), English, French, German, Hebrew, Italian, Japanese, Russian, and Spanish. Because *Latent Semantic Scaling* requires only a few seed words, it is easily adaptable to other languages. (The *LSS* algorithm is based on word embedding, which does not have language dependency.)