

Cross-national analysis of global security discourse using word embeddings

Takuto Sakamoto^{*†}

Graduate School of Arts and Sciences, The University of Tokyo

Abstract

Recent methodological advances in quantitative text analysis enable a not only more rigorous, but also more systematic, investigation of the discursive dynamics that are unfolding among various international actors, in comparison with preceding approaches. This study conducts such an investigation in the context of policy deliberations on international peace and security, by performing cutting-edge text analysis on the entire body of meeting records of the Security Council of the United Nations for the past quarter of a century (1994-2019). Focusing on one of the most consequential notions for the council's policy-making, "threat to the peace," this study employs an unsupervised machine learning model, broadly termed "word embedding," to analyze how this notion has been discussed by relevant members of the council, especially its five permanent members, during the period under investigation. Word embedding estimates the semantic relationships among all words found in a given set of documents, by assigning a high-dimensional vector representation to each of these words. In combination with other natural language processing techniques and simple statistical operations, it reveals persistent patterns of cross-national convergence and divergence in security discourse. While these findings are mostly consistent with existing empirical observations on council politics, some are not, including, most notably, a considerable degree of correlation in how to conceive international threats found between the close-knit Western allies in the council and their Russian counterpart. The study also tracks changing conceptions of threats over time, unveiling distinctive aspects of international relations in each period that have strongly affected the evolution of these conceptions.

Key words: United Nations, international security, text analysis, word embedding, threat conception

1 Introduction

International relations contain a considerable amount of argumentation and persuasion. States and other actors, through the voices of their respective delegates, argue and counter-

^{*}Corresponding author.

[†]E-mail: sakamoto@hsp.c.u-tokyo.ac.jp

argue for specific policy positions or particular normative ideas, in order to influence the perceptions, the interests, and eventually the policy choices of their counterparts [1]. These discursive dynamics happen in various institutional settings (e.g., bilateral or multilateral) as well as different issue areas (e.g., security, trade, finance, and the environment), and include those involving hard decisions concerning the use of force, where the consideration of power and interests often seems to overwhelm the reason of rational arguments [2, 3].

The Security Council of the United Nations (UN) is a notable example in this regard. According to the UN Charter, the council has primary responsibility for the maintenance of international peace and security (Article 24). As such, its members hold different forms (public or closed, formal or informal) of consultations and deliberations almost on a daily basis, in order to address a wide array of issues concerning global security, from humanitarian crises to transnational terrorism, to nuclear proliferation. These consultations and deliberations have produced conflicting interpretations. Many observers, pointing to high-profile "hard cases" such as Kosovo in 1999, Iraq in 2003, and more recently Syria, draw highly confrontational pictures from these interactions among council members, especially its five permanent members (P5) [4–7]. Others, however, are more optimistic. Citing cases of great-power cooperation in areas such as counter-terrorism and peacekeeping, they see more integrative force in the council's discursive dynamics, which might bring its members some semblance of community, or even corporate identity [3, 8, 9]. Given the accumulation of intense deliberations on such a broad range of policy agendas over so many years, however, neither strand of assessment seems to represent the whole picture of council politics. In other words, the understanding and evaluation of the international relations unfolding at the Security Council have been severely hampered by the lack of systematic investigations of its discursive dynamics over time.

The present study addresses this gap by proposing and implementing a novel methodological framework for capturing discursive dynamics that involve multiple international actors and have taken place for many years. The study benefits from the recent advancement in large-scale text analysis using computational methodologies [10, 11]. In particular, it employs "word embedding" to construct measurements for quantifying cross-national differences in actors' understanding of a given notion in international relations (e.g., "peace," "threat," "sovereignty") and for tracking the evolution of such understanding over time. Word embedding is a relatively recent product from the blossoming field of computational linguistics [12–14]. It estimates numerical vector representations of words appearing in a given set of documents, or "corpus," based on the relations of co-occurrence among these words. This enables one to represent a semantic relationship between any combination of words (e.g., "threat" and "terrorism"), with some quantified measure of distance or similarity between the corresponding vectors. Such a measure, then, permits the rigorous examination of variations in conceptual or normative understanding concerning a given aspect of international relations, across actors and across time periods. In addition to word embedding, this study employs other natural language processing (NLP) tools, including part-of-speech (POS) tagging and phrase detection, in order to make the above analysis more fine-grained and more context-sensitive.

This study applies these analytical tools to the meeting records (*proces-verbaux*, or

P.V.) of the UN Security Council, which contain the transcripts of all the public statements delivered by various delegates and officials at the official meetings that were held openly¹. As of March 31, 2020, the council had convened 8,745 such meetings, starting from its first session on January 17, 1946. Among these, the study deals with more than 5,100 meetings, held between January 1994 and December 2019, covering a sizable part of the post-Cold War council politics. The records of these meetings are publicly available online, and many of them have actually been used as descriptive or historical materials for preceding works [3, 6, 9, 15, 16]. Their usage, however, remains partial. In contrast, this study treats thousands of meeting records as a single dataset, and conducts systematic text analysis on the entire body of the public deliberations recorded during the past quarter of a century.

Although one can address many interesting research questions using this large dataset, the present study focuses on one particular notion: "threat to the peace." Being "the charter's key threshold for council action" [17], this notion is highly consequential in policy debates in the council in that, if some issue or entity is determined to constitute a "threat to the peace" (Article 39), then the council is legally justified to invoke various collective enforcement measures stipulated in Chapter VII of the UN Charter, including military action (Article 42), in order to remove the designated threat. The problem here is that the charter itself does not define what constitutes a "threat to the peace," and instead leaves the determination of the specific contents of this notion entirely to the collective discretion of council members [18]. Accordingly, council deliberations involving the notion of threat have obvious implications for global security. This study thus attempts to explore how this notion has evolved over time, and how differently the council members have conceived the notion, by systematically investigating its semantic content, in relation to other concepts appearing in council deliberations, across actors and periods.

This manuscript is organized as follows. The next section briefly discusses the related literature to further illuminate the study's substantive and methodological contributions. Section 3 describes the specific procedures taken to obtain, preprocess, and analyze the main dataset, that is, the meeting records of the UN Security Council. The section also gives accounts of how to exploit word embeddings to quantitatively characterize an actor's threat conception, and then to compare such conceptions among different actors and over different time periods. Section 4 reports the results of all of these measurements and analyses, and the following section discusses their implications in relation to existing knowledge on council politics. The last section points out the remaining challenges for this study, and suggests several fruitful directions that further study could take. The online appendix provides further descriptions of the study's methodological aspects. It also reports on additional measurements and analyses that were conducted to examine the robustness of the results discussed in the main text.

¹Sometimes, an official meeting is held in a closed manner, leaving only a brief summary of discussion contained in the communiqué issued after the meeting. The following analyses exclude these closed meetings from consideration, and deal only with official meetings that are open to the public.

2 Related Literature

Almost any textbook on the UN mentions "the expansion of what the Council considers a threat to international peace and security" [18], especially after the end of Cold War [19–21]. For example, numerous observers have noted the increasing involvement of the council in what had been considered essentially as internal affairs, such as civil wars and humanitarian crises [2, 22]. This observation has often been accompanied by another typical note of its redoubled efforts for countering terrorism and containing proliferation, following the September 11 attacks on the United States [3, 5, 7]. To this expanding list of potential threats for the council's consideration, some might also add other pressing issues such as HIV/AIDS and the illicit flow of small arms and light weapons, which often leads to the situation that critics describe as a "theme park" of nontraditional security threats [23]. Although such enumeration conveys a general sense of how extensive the organ's roles have become during the past several decades, it does not help much to address the more analytically and practically important aspects of the council's policy making. Those aspects largely concern how each of council members conceives various issues and entities on its own scale of threat, and how different these conceptions are. Answering these questions has obvious importance, as it reveals much about the viability and the possible direction of the concerted action, especially under Chapter VII of the charter, that the council might take at a given moment.

The large accumulation of studies on the council and the UN at large certainly offers rich materials for examining these crucial aspects of council politics. It is, however, not necessarily helpful for doing so in systematic and rigorous fashions. As already mentioned, many studies have focused on relatively high-profile showdowns involving the P5 [5–7, 24, 25]. The in-depth descriptions and interpretations that this body of work provides are highly informative, as far as one's aim is, for example, to know specific aspects of divergence in the threat perceptions of the United States and the Russian Federation at the time of the Kosovo crisis in March 1999 [3, 8]. However, these particular events taking place at discrete moments of the history, no matter how shocking and consequential they are, concern only a part of the enormous amount of business that the council has continuously conducted over many years. This necessitates, in addition to the in-depth case studies that are common in the literature, a more holistic approach to the council's activities to place the former's insights in a proper context.

In this respect, some other studies might be more relevant as they indeed take such an approach [26, 27]. Specifically, there are notable studies on the council's major policy output, namely, council resolutions [28, 29]. Through analysis of various features of resolutions and their adoption (agenda, region to address, invocation of Chapter VII, voting pattern, etc.) over a long period of time, these studies have systematically revealed, for example, the changing pattern of policy agendas that the council has addressed over time, as well as the changing frequency of, and reasons for, vetoes cast by the P5 against draft resolutions. This body of work, however, is not sufficiently fine-grained to uncover the deeper processes at work when the council produces its policy output, that is, the underlying divergence and convergence in actors' normative and conceptual understanding of the world, including those concerning security threats.

The situation described so far is somewhat unfortunate, given the wide availability of data that could be utilized, with proper analytical tools, to overcome the challenges mentioned above. The UN has made publicly available a huge body of textual documents, which record traces of international interactions that have been taking place in this multi-lateral body during its 75 years of existence². Regarding the Security Council, these documents include its main policy outputs, such as resolutions and presidential statements, annual reports of its activities submitted to the General Assembly, and the speech records (P.V.) and communiqués of its official meetings. The systematic use of these and other UN documents has still been limited, however, with several notable exceptions [30–32]. This study leverages one series of council documents, P.V. ("S/PV" in the UN system of document symbols), to investigate the threat conceptions of different international actors, in a both systematic and fine-grained fashion.

In the wider field of political science, considerable progress has been made with the "text as data" approach [11, 33, 34], with many studies measuring various partisan and/or ideological differences by extensively using political texts such as speech transcripts of legislators (e.g., *Congressional Record* for the US Congress) [35–42]. Somewhat belatedly, International Relations (IR) has also started to embrace this development. Besides several outstanding works [43–46], it is noteworthy that some researchers have also conducted advanced text analysis on security documents, including council speech transcripts [47–50]. Using such quantitative tools as topic models (e.g., Latent Dirichlet Allocation; LDA) and spatial models, these studies have largely focused on the structural aspects of security discourse, such as framing, often in a particular issue area (e.g., Afghan conflict, Syrian conflict, or terrorism). Those tools, however, do not necessarily reveal the detailed semantics of a particular concept (e.g., threat) in a given discursive context, and its variation across multiple actors and periods, all of which are the main focus of this study.

Word embedding, the tool employed here, is highly useful in uncovering exactly these aspects of discourse [13, 14]. Though still underutilized in IR, it has been extensively applied in other scientific fields, including sociology. Specifically, word embedding has been effective in successfully recovering known human biases, racial and ethnic stereotypes, and conceptual associations from a broad array of texts, written in diverse social and historical settings [51–54]. The methodological framework described below owes much to these preceding applications.

3 Data and Methods³

3.1 Text Data and its Acquisition

The English transcripts of council meetings (S/PV) were automatically (i.e., by using a web-crawler) obtained from the Official Document System (ODS) of the UN, via hyperlinks embedded in the summary tables provided on the website of the Dag Hammarskjöld

²These documents are available from the Official Document System (ODS) of the United Nations (<https://undocs.org>; last accessed on March 25, 2020).

³Unless otherwise reported, the procedures described in this section were implemented using the Python programming language (version 3.7) and its numerous extensions (modules).

Library⁴. These tables, each row of which corresponds to each official council meeting, are highly useful, as they provide additional pieces of information on each meeting, including the meeting's date, topics under discussion, links to related documents (e.g., P.V., communiqué, or resolution), and the voting outcome (if a vote was taken on a draft resolution). At the time of data acquisition, these tables were available for the meetings from January 1994 onward⁵. This effectively set the temporal boundaries of the present study. That is, the study covers the official council meetings from the 3,326th (held on January 6, 1994; the first meeting of the year) to the 8,697th meeting (on December 20, 2019; the last meeting of the year). Among these 5,372 meetings, 238 were closed to the public, leaving only brief communiqués without any detailed records of deliberations. These closed meetings were excluded from the following analysis.

The proceedings of each meeting are recorded in a single document file (PDF format; for a large number of records, Microsoft Word documents are also available). Each file contains the transcripts of speeches delivered by speakers during the corresponding meeting, the names of these speakers, and the countries or agencies they represented, along with other information such as the agenda and the date of the meeting. Although these files themselves do not constitute a structured database, exploiting several formatting conventions of council meeting records enables one to locate and extract relevant pieces of information, while discarding irrelevant ones (e.g., page numbers, document symbols, unnecessary punctuation, etc.). For example, a typical transcript of a speech begins as follows: "Mr. Nebenzia (Russian Federation) (*spoke in Russian*): The Russian Federation will vote against the humanitarian troika's draft resolution S/2019/961, on the cross-border assistance mechanism in Syria..." (S/PV.8697 (December 20, 2019)). By making use of honorific titles ("Mr." in the above quote) and formatting symbols (the parentheses and the colon) as data locators, one can isolate the surname ("Nebenzia") and the country ("Russian Federation") of the speaker, and the text of the speech he delivered ("Russian Federation will vote against..."). More detailed information on the text extraction process is given in Appendix A. The Python code used for this procedure is available from GitHub⁶.

One aspect of council deliberations that is potentially confounding for their systematic treatment is that a sizable portion of the recorded speeches are actually procedural statements given by the council's monthly-rotated presidents (referred to as "The President" in the proceedings). These statements, for instance, include opening and closing remarks, vote announcements, and ceremonial greetings. This aspect has largely been overlooked in existing quantitative research on council meeting records [32]. Meanwhile, a president can also give a non-procedural policy statement as a representative of his or her country. Such a statement typically begins with a sentence like "I shall now make a statement in

⁴<http://research.un.org/en/docs/sc/quick/meetings/> (last accessed on March 25, 2020)

⁵Summary tables later became available for all the meetings before 1994 as well. This is a welcome development. However, several considerations precluded rerunning the entire process of preprocessing and embedding yet again for all the meeting records from 1946 at this stage. One reason is that the formatting convention of S/PV documents was not established well for older meeting records. Extracting relevant text (speeches) and related information (the names and the countries/agencies of speakers) from these records in a consistent way is thus expected to be considerably arduous and time-consuming.

⁶<https://github.com/takutos/undocs>

my national capacity" or "I shall now make a statement in my capacity as representative of the United States." This formalism gives another useful convention of council deliberations, which was exploited here to screen out non-policy procedural statements. With this additional procedure, a total of 48,887 (supposedly substantive) speeches, which consist of 40,489,764 words, were obtained (82,136 speeches consisting of 43,137,749 words including the discarded procedural statements).

Appendix B provides several illustrations concerning how the extracted speeches are distributed across different actors and across different years. They confirm the general tendencies of council meetings already pointed out in prior work [32]. One such tendency is the steadily expanding volume of council debates, with an increasing number of speakers appearing in the council chamber over years. This tendency has actually been accelerating over the past five years. Another notable tendency is the sizable, but not necessarily overwhelming, volume of speeches contributed by P5 delegates over the past quarter of a century. In fact, the combined number of P5 speeches is 9,843, barely over 20% of the total. The most vocal countries in terms of the number of speeches are: Russian Federation (2125), United States (2002), France (1967), China (1912), United Kingdom (1837), Japan (899), Germany (756), South Africa (673), and Brazil (631)⁷.

3.2 Text Preprocessing⁸

A series of preprocessing measures were applied to the obtained text data [10, 11]. These measures include the removal of punctuation and other non-alphanumeric symbols, converting uppercase characters to lowercase ones, converting the spelling of English words (from British to American English spelling)⁹, and removing "stop words." Usually comprising articles (e.g., "a," "the"), pronouns ("I," "she"), prepositions ("of," "in") and other function words, stop words are a set of words that are considered too common to be sufficiently informative and discriminating for use in a given analysis. This study used the standard set of English stop words (179 in total) provided by the Python module (NLTK) employed here.

In addition to these fairly common and simple procedures, two statistical techniques for natural language processing were also used. One was part-of-speech (POS) tagging, which employs a statistical model and a dictionary to estimate a POS (e.g., "noun," "verb," "adjective") for each word appearing in a given text. This information was needed in the later analysis to sort out various objects, named (proper nouns) as well as unnamed (common nouns), from the other words mentioned in the extracted speeches.

The other procedure was phrase detection. Using a statistical algorithm¹⁰, this procedure combines frequently collocated words into a single phrase. For example, unsur-

⁷Another body, the UN Secretariat, would likely appear somewhere in this list. It is not listed here because many of its officials (except the Secretary-Generals), whose affiliations, unlike those of state representatives, are not annotated next to their surnames in meeting records, have not yet been identified.

⁸Unless otherwise reported, the preprocessing procedures described below were implemented by the widely-used Python module named NLTK (Natural Language Toolkit).

⁹This procedure made use of an online spelling list (<http://www.tysto.com/uk-us-spelling-list.html>; last accessed on April 2, 2020).

¹⁰This study used the phrase detection algorithm provided by Gensim, a Python-based library for advanced text processing and analysis.

prisingly in the context of council deliberations, the word "Security" ("security:NNP" in the processed corpus, where the POS-tag, "NNP," stands for a singular proper noun) is often followed by the word "Council" ("council:NNP"). Therefore, the phrase detection algorithm would likely combine these two words ("security:NNP_council:NNP"). This study iterated this procedure twice, so that up to four words could be combined into a phrase (e.g., "international:NNP_atomic:NNP_energy:NNP_agency:NNP"). Moreover, it applied the algorithm only to fairly frequently occurring words and phrases, or more specifically, those that have appeared in deliberations 1,000 times or more during the study period. The phrase detection was expected to partially mitigate the brutal effects of the so-called "bag-of-words" approach required in the following analysis, which, in its representation of a text, largely disregards the syntax structure of the text.

With the above preprocessing measures, the speech texts consist of a total of 19,812,732 POS-tagged tokens (words and phrases), with 133,128 unique tokens. The entire set of these speeches, or the corpus, will be available from the Harvard Dataverse. The Python code used for its derivation is available on GitHub¹¹.

3.3 Word Embedding

Word embedding refers to a class of unsupervised machine learning models that learn high-dimensional vector representations of words (or POS-tagged tokens in the present context) in a given corpus. Each word is "embedded," or assigned a series of numbers denoting its coordinate, in vector space, such that two words that tend to co-occur with a similar set of words in text are likely to have a smaller distance in that space. While there are several popular methods for accomplishing this job, such as "word2vec" algorithms [13], this study employed GloVe (Global Vectors for Word Representation), an embedding algorithm that has been proven to effectively perform various linguistic tasks (e.g., word analogy) on a wide variety of corpora of different sizes [14, 51–53]. GloVe efficiently computes a word-word co-occurrence matrix from an input corpus, and then applies a log-bilinear regression model to this matrix to estimate the numerical elements of a vector denoting each word in the corpus [14]¹².

The derived vector space essentially comprises semantic relations among the various concepts appearing in a given set of documents. These relations are relative to the corpus concerned, and thus should depend on the context (time, author, subject, etc.) in which the corpus was generated. Therefore, in addition to running the embedding algorithm for the entire set of council speeches, this study also divided the corpus in different ways, according to countries (e.g., a set of speeches delivered by U.S. delegates), time periods (e.g., speeches delivered during 2015-2019), or both (e.g., speeches delivered by U.S. delegates during 2015-2019). It then applied the algorithm to each of these subsets of the original corpus. As too much division would result in too small subsets, which would render reliable statistical estimation untenable, there was a certain limitation to this endeavor. Accordingly, most of the analyses reported below rest on word embeddings that

¹¹<https://github.com/takutos/undocs>

¹²The GloVe implementation code is available from the developers' website (<https://nlp.stanford.edu/projects/glove/>; last accessed on April 3, 2020).

were derived from relatively large sets of speeches, in particular, speeches contributed by countries with a sizable council presence (typically permanent members) as well as those spanning fairly long periods of time (e.g., five years).

GloVe works with several learning parameters. Among them, the number of dimensions of embedding space was set to be 100. Another important parameter, the size of the "context window," was set to be ten in both directions. That is, if a word is located up-to-ten words ahead or behind another word, then these two words are considered to have co-occurred. The word embeddings (vectors), derived (or "trained" in the machine-learning terminology) in this parameter setting, for all the POS-tag tokens in the entire corpus (more than 130,000 in total), will be available from the Harvard Dataverse. The above parameter values are fairly common, but are undeniably arbitrary. Appendix F reports the results obtained in different learning settings, which show no major difference from the results reported in the next section.

3.4 Threat Representation and Measurement

Each numerical entry of a 100-dimensional embedding vector has no intrinsic meaning. However, the quantitative relationship between the vectors induced by these numbers has a substantive interpretation. That is, the distance or the similarity between two embedding vectors quantifies the degree of semantic association between the corresponding words or tokens. A strong association implies that the two words concerned are identical except for some grammatical features (e.g., "threat" and "threats"), synonymous ("threat" and "menace"), otherwise semantically similar ("threat" and "danger"), or frequently mentioned in a similar context ("threat" and "terrorism" in council deliberations). The meaning of a given concept in a corpus can be revealed in a relative manner by probing its semantic associations with other concepts appearing in the same corpus [53]. The representation of threat conception described below leverages the relational understanding of meaning that word embedding allows to be formalized.

Specifically, let $w_i \in \mathbb{R}^d$ denote one of the trained d -dimensional embedding vectors ($d=100$ in the current setting) that represents a POS-tagged token i . Let also $i \in V_C$, where V_C is the set of all unique POS-tagged tokens contained in a corpus C . Then, for two tokens, i and j , in V_C , the following cosine similarity between the two corresponding vectors gives a widely-used measure of the semantic association between i and j . The product in the numerator is an inner or dot product of two vectors. The measure takes a value from -1.0 (diametrically opposite) to 1.0 (identical).

$$\text{similarity}_c(i, j) = \frac{w_i \cdot w_j}{\|w_i\| \|w_j\|} \quad (1)$$

Building on the above measure, this study represented the threat conception latent in a given set of council speeches as a series of similarity measures computed between threat-related tokens (e.g., "threat," "threats") and some other reference tokens used for characterizing that conception (e.g., "terrorism," "al-Qaeda," "HIV/AIDS"). Suppose that T_C denotes a subset of V_C , which consists of threat-related tokens. Suppose also that W_C denotes another subset of V_C , which contains reference tokens. Further, let $\tau \in \mathbb{R}^d$ represent the mean of the embedding vectors corresponding to the elements of T_C . Using these

notations, a measure of the threat conception ("threat measure") for corpus C , $threat(C)$, can be expressed as the following series of cosine similarities:

$$threat(C) = \left(\frac{\tau \cdot w_i}{\|\tau\| \|w_i\|} \mid i \in W_C \right) \quad (2)$$

$$\tau = \frac{\sum_{j \in T_C} w_j}{|T_C|}$$

In most of the analyses reported below, $T_C = \{"threat:NN," "threats:NNS"\}$, where the tags, "NN" and "NNS," denote a singular common noun and a plural common noun, respectively. "threat:NN" and "threats:NNS" are two of the most frequent tokens that contain a string "threat"¹³. W_C , on the other hand, was given 2,000 of the most frequent POS-tagged tokens from the whole corpus that contain noun words (both common and proper nouns)¹⁴. Nouns were used to isolate possible subjects and objects that could be associated with the notion of threat. Less frequent tokens were excluded, as the estimation of their embeddings could be highly unstable. Other compositions of these token sets (e.g., 1,000 tokens instead of 2,000, proper nouns instead of all nouns) were also tried. Again, as reported in Appendix G, the main results described below were largely resistant to these manipulations.

Cross-national and temporal comparisons of threat conceptions can be made by computing and comparing $threat(C)$ over different corpora or subsets of a corpus. Although there are numerous ways to do this (e.g., clustering, principal component analysis), the present study used the simplest techniques: pair-wise operations of threat measures, such as subtracting one from another and computing a Pearson correlation coefficient for two measures, both in reference to the common tokens W_C ¹⁵. For example, consider that two subsets of the council corpus, C_{US} and C_{RU} , represent distinct sets of the speeches delivered by representatives of the United States and the Russian Federation, respectively. The subtraction, $threat(C_{US}) - threat(C_{RU})$, performed for each element of W_C might uncover specific issues and entities on which the threat conceptions of the U.S. and Russia are particularly divergent. Similarly, the Pearson correlation calculated between $threat(C_{US})$ and $threat(C_{RU})$ gives a rough measure of the overall closeness of the two countries, in terms of the issues and entities associated with the threat to international peace and security.

In reporting the numerical values of the quantities defined above, their confidence intervals were computed by using bootstrapped samples of council speeches (for threat measures) and reference tokens (for correlation coefficients). Using the latter samples, Kolmogorov–Smirnov (KS) two-sample tests were also performed to evaluate differences

¹³In the entire corpus, "threat:NN" and "threats:NNS" appear 12,530 and 9,758 times, respectively. There are actually 57 other tokens that contain "threat," including, for instance, "threat:NN_international:JJ" (2,016 times) and "threat:NN_posed:VBN" (1,249 times). These other tokens were excluded from the analysis because of their relatively low frequencies.

¹⁴One of the least frequent among these 2,000 tokens is "gender-based:JJ_violence:NN," whose frequency is 1,030 times in the corpus. Unsurprisingly, the most frequent token is "council:NNP" (102,028 times), which is followed by "security:NNP_council:NNP" (88,465 times).

¹⁵In the case that one or both of the subsets which were compared lacked some elements of W_C , these tokens were skipped in pair-wise operations on the subsets concerned.

in correlation coefficients. Appendix C gives a detailed account of these statistical procedures.

4 Results

4.1 Plausibility of Word Embeddings

Before focusing on threat conception, it might be useful to gain a sense of overall plausibility about the derived word embeddings. Table 1 shows, for each of the tokens listed in the first column, the ten most strongly-associated tokens that contain noun words¹⁶. For the sake of readability, the POS-tags of the tokens in the second column were removed. The numbers in parentheses in the first column are the frequencies of the listed tokens, while those in the second are the cosine similarities, as defined in Equation (1), between the corresponding tokens in the two columns. The embedding vectors representing these tokens were trained using the entire set of council speeches, disregarding any cross-national and temporal differences. Therefore, the associations displayed in the table can be considered as reflecting the council's aggregated understandings of the listed notions over the past quarter of a century.

The table suggests a variety of ways in which two notions are associated with each other by their respective vector representations. For example, two tokens can just be semantically similar almost in any context ("security" and "stability," "threat" and "danger"), or interchangeable in the context of council deliberations (e.g., "security_council" and "council," "africa" and "continent"). In contrast, these tokens might be only weakly related to each other or even have opposing meanings in the general context (e.g., "somalia" and "coast," "peace" and "conflict," "security" and "threat"), but nevertheless form some common cluster of notions inside the council chamber. Despite the considerable degree of context dependence, however, most of the associations shown in the table would presumably seem natural to the eyes of observers of council politics, or broader international relations. The most notable in this regard are the tokens associated with "sovereignty:NN" (the last row of the table). These tokens constitute a familiar set of words and phrases (i.e., "territorial integrity," "sovereignty and territorial integrity," "independence," "non-interference in internal affairs," etc.) that typically accompany the notion of state sovereignty in its traditional conception.

4.2 Threat Measures

Table 1 lists a small number of tokens, such as "danger" and "terrorism," that are strongly associated with a single token, "threat:NN," at the most aggregate level of the corpus. In order to allow for more comprehensive and more fine-grained analyses, the threat measure

¹⁶Although the NLTK POS tagger employed here works effectively most of the time, it sometimes makes obvious misclassification. In the present context, for example, some adverbs that appear at the beginning of a sentence (thus its first character is capitalized in text), including "Furthermore," "Nevertheless," "Hence," "Today," "First," "Therefore," and so forth, are often classified as proper noun (given a "NNP" POS-tag) rather than as adverb ("RB"). As long as these cases could be manually detected, they were removed from relevant tables and figures for the purpose of avoiding confusion. They were not removed from measurement and calculation for avoiding arbitrariness.

TABLE 1
Token-Token Associations Computed from the Derived Embeddings

Example Token	Strongly Associated Tokens
security:NNP_council:NNP (88,465)	council (0.939), resolution (0.741), time (0.736), international_community (0.735), members_council (0.717), members_security_council (0.715), matter (0.714), action (0.698), council_members (0.696), united_nations (0.692)
peace:NN (51,846)	lasting_peace (0.806), peace_process (0.751), peace_security (0.738), sustainable_peace (0.734), peace_stability (0.726), reconciliation (0.714), stability (0.697), conflict (0.687), region (0.684), achievement (0.677)
security:NN (50,623)	development (0.752), stability (0.746), peace_security (0.709), threats (0.696), challenges (0.689), region (0.688), context (0.675), order (0.665), threat (0.663), security_situation (0.662)
united:NNP_nations:NNP (44,532)	field (0.734), organization (0.713), missions (0.678), united_nations_system (0.675), addition (0.671), peace-keeping (0.671), peacekeeping_operations (0.669), personnel (0.663), mission (0.662), context (0.659)
africa:NNP (18,893)	continent (0.893), african (0.757), african_countries (0.752), great_lakes_region (0.709), subregion (0.682), world (0.680), conflicts (0.675), region (0.672), peace_security (0.649), problems (0.646)
somalia:NNP (15,337)	somali (0.746), afghanistan (0.698), federal_government (0.685), horn (0.659), african_union_mission_somalia (0.631), piracy (0.628), amisom (0.627), coast (0.623), liberia (0.604), countries_region (0.603)
threat:NN (12,530)	threats (0.840), threat_international (0.789), danger (0.752), threat_posed (0.743), challenge (0.729), terrorism (0.724), international_peace_security (0.706), peace_security (0.685), phenomenon (0.668), security (0.663)
sovereignty:NN (5,621)	territorial_integrity (0.905), sovereignty_territorial_integrity (0.846), independence (0.736), internal_affairs (0.694), principle (0.683), unity (0.672), non-interference (0.658), territory (0.641), respect (0.640), state (0.628)

defined in Equation (2) was computed for different subsets of the corpus with the common token sets, T_C and W_C . Table 2 displays a tiny fraction of the output (more systematic results will be available from the Harvard Dataverse). The table consists of cosine similarities between the mean embedding vector for T_C ("threat:NN" and "threats:NNS") and vectors representing a small subset of W_C ("terrorism:NN," "proliferation:NN," ..., "sahel:NNP" and "iran:NNP"). Each row corresponds to one of six sets of council speeches, that is, the whole corpus (denoted as "All" in the first column) and its five subsets, each of which comprises all of the speeches contributed by one of the P5. Numbers in parentheses

in a cell indicate the 90% confidence interval (CI) of the corresponding cosine similarity, which was calculated from bootstrapped samples of the original speeches (see Appendix C).

Despite its lack of representativeness, the table is still revealing regarding the somewhat intricate nature of the threat conceptions of powerful council members. Overall, the P5's threat conceptions seem to be not radically different from each other at this level of aggregation. In particular, the P5 members almost equally strongly associate some tokens such as "terrorism" and "proliferation" with T_C without marked cross-national differences. Nevertheless, some members put more emphasis on some issues than the others do. For example, both "drug_trafficking" and "afghanistan" have far more salience in Russia's threat measure than in that of any other P5 member. Equally, France's geographic focus on "sahel" is obvious, as is the readiness of the U.S. to emphasize "iran" in its threat conception. It is also notable that "climate_change" has received markedly stronger emphasis in the threat measures for the West European members of the council (France and the U.K.). As will be briefly discussed in the next section, these distinctive features derived from the threat measures are largely consistent with existing empirical observations about the P5.

One can express threat measures in a more concise manner. Figure 1 displays, for each of the different subsets of the corpus, a varying degree of association of W_C tokens with T_C , in the form of a word cloud¹⁷. For the sake of visual clarity, only W_C tokens with proper nouns, that is, named entities, are displayed. This representation can be a highly useful tool for efficiently describing the threat conception of a given actor for a given period of time.

4.3 Cross-National Correlations

As explained in the previous section, computing the correlations between threat measures allows more extensive cross-national comparisons to be made. The heatmap in Figure 2 shows the degree of correlation between threat measures over the entire W_C , for each combination of the P5 members, across the whole study period. By a crude mathematical operation, the underlying correlation table has been arranged so that two countries that have mutually similar series of correlation coefficients tend to be located closer to each other in both axes of the table (rows and columns)¹⁸. This makes it easier to recognize the structural configuration of the threat conceptions of the council members.

Several observations can be made regarding this heatmap. Firstly, among the P5, the United States, the United Kingdom, and France, which are also members of the North Atlantic Treaty Organization (NATO), show a predictable tendency toward having closely similar threat conceptions. The pair-wise correlations of threat measures among them are, in descending order, 0.704 (90% CI: 0.682-0.725) between the U.S. and the U.K., 0.676 (0.653-0.699) between the U.K. and France, and 0.666 (0.640-0.691) between the U.S.

¹⁷These representations were produced with Python's wordcloud module.

¹⁸More specifically, the algorithm projects the correlation coefficients computed for each country onto the first eigenvector of the correlation matrix, and then sorts the countries concerned according to their projected values. Other algorithms such as agglomerative hierarchical clustering return mostly similar, but not necessarily same configurations.

TABLE 2
Aggregated Threat Measures for the P5 and the Whole Council, 1994-2019 (Excerpt)

	terrorism	proliferation	drug_trafficking	climate_change	al-qaida	afghanistan	sahel	iran
China	0.651 (0.613-0.681)	0.498 (0.400-0.584)	0.443 (0.383-0.553)	0.342 (0.262-0.388)	0.190 (0.070-0.307)	0.271 (0.231-0.311)	0.262 (0.183-0.362)	0.131 (0.054-0.220)
France	0.678 (0.634-0.716)	0.597 (0.503-0.657)	0.495 (0.429-0.576)	0.490 (0.396-0.554)	0.317 (0.247-0.408)	0.313 (0.233-0.380)	0.440 (0.364-0.503)	0.142 (0.106-0.183)
Russia	0.716 (0.681-0.746)	0.488 (0.463-0.527)	0.583 (0.497-0.670)	0.292 (0.115-0.414)	0.420 (0.384-0.452)	0.528 (0.497-0.576)	0.330 (0.256-0.398)	0.166 (0.053-0.233)
U.K.	0.706 (0.659-0.755)	0.584 (0.509-0.698)	0.342 (0.215-0.427)	0.480 (0.362-0.594)	0.326 (0.199-0.419)	0.333 (0.244-0.389)	0.319 (0.175-0.418)	0.214 (0.168-0.273)
U.S.	0.701 (0.640-0.734)	0.567 (0.525-0.609)	0.373 (0.276-0.460)	0.361 (0.244-0.450)	0.470 (0.399-0.521)	0.338 (0.300-0.390)	0.408 (0.339-0.451)	0.321 (0.274-0.357)
All	0.725 (0.711-0.744)	0.643 (0.619-0.679)	0.573 (0.549-0.598)	0.569 (0.546-0.587)	0.482 (0.453-0.508)	0.475 (0.448-0.503)	0.469 (0.436-0.502)	0.352 (0.323-0.376)



FIGURE 1: Word-Cloud Representations of Threat Measures




FIGURE 2: Correlations of Threat Measures among the P5, 1994-2019

and France. This order is statistically significant, on the basis of KS two-sample tests that compared the possible combinations of these correlations¹⁹. Secondly, perhaps less predictably, Russia is not particularly dissimilar to the close-knit Western members in terms of the threat measures defined here. Russia's correlations with these countries are 0.674 (0.651-0.696) with France, 0.651 (0.627-0.674) with the U.S., and 0.644 (0.619-0.668) with the U.K. These values are fairly comparable with the previous correlations. In particular, the correlation between Russia and France is significantly higher than that between the U.S. and France²⁰. Otherwise, the correlations are significantly lower than any of the other correlations among the NATO members. Thirdly, as is very clear from Figure 2, China is distinctively autonomous in its threat conception, in comparison with the other permanent members. The country whose threat measure is most correlated with China's is France (correlation: 0.576; CI: 0.548-0.604), then narrowly followed by Russia (0.574; 0.546-0.601)²¹. These correlations are, however, markedly lower than those among the other permanent members.

The correlations reported above were computed from the temporally aggregated sets of speeches. Appendix D provides additional correlation heatmaps that illuminate tempo-

¹⁹The KS statistics that were computed from bootstrapped samples (the size of each sample: 10,000) of threat measures, under the null hypothesis that two dyads of countries are indistinguishable in terms of their distributions of threat measure correlations, are: 0.816 for U.S.-U.K. correlation against U.K.-France correlation, 0.279 for U.K.-France against U.S.-France, and 0.816 for U.S.-U.K. against U.S.-France. The corresponding P-values are all negligible.

²⁰The corresponding KS statistic is 0.228. The P-value is again near zero.

²¹The difference between China's correlations with France and Russia is significant. The corresponding KS statistic is 0.060 with near zero P-value.

ral variations in the threat measure correlations among the P5. One can see considerably fluctuating dynamics from these figures. For example, in relative terms, in the period from 1994 to 1999, the U.S. and Russia had more similar threat conceptions (correlation: 0.274; CI: 0.233-0.314) than any other dyad, whereas, in the ensuing five years (2000-2004), France and Russia had the highest threat measure correlation (0.367; 0.332-0.401) between each other. For comparison, the correlations between the U.S. and the U.K. in the same periods were 0.238 (0.196-0.279) and 0.344 (0.306-0.382), respectively. Otherwise, however, deviations from the observations in the preceding paragraph are limited, in particular for recent periods such as 2010-2014 and 2015-2019. Furthermore, the distance of China's threat conception from those of the other members has been consistently conspicuous over all the periods.

Finally, including other council members in this analysis requires caution. Unlike the permanent members, the other ten members ("E10") of the council are elected for a two-year term, without immediate re-election. Moreover, half of them are replaced with newly elected members annually (Article 23 of the Charter). Without adequately taking account of these aspects of council membership, comparisons involving elected members could lead to misleading results, whereas imposing too many constraints on the input data (e.g., using a two-year subset of the corpus) could run the risk of making a shaky estimation.

With these caveats, Figure 3 depicts the correlation table for the threat measures of the ten states that served on the council in both 2018 and 2019, for which a relatively large number of speech records are available. These states are: Russia (430 speeches), the U.S. (379), the U.K. (378), China (377), France (371), Peru (371), Kuwait (357), Equatorial Guinea (354), Poland (354), and Côte d'Ivoire (345). While adding the elected members to the analysis obviously made the resultant patterns more complicated, it is still possible to recognize some structural features, including those already pointed out regarding the P5, such as the persistent difference between China's threat conceptions and those of the other permanent members. It is also notable that the five elected members seem to form a somewhat autonomous cluster, which is relatively distinct from the P5. Another noticeable feature is the considerable difference existing among the P5 in the correlations of their threat measures with those of the elected members. That is, China, the U.S., and, to a lesser degree, Russia have relatively tenuous threat measure correlations with the elected members, whereas France and, to a lesser degree, the U.K. have much stronger correlations.

4.4 Points of Divergence

Cross-national convergence or divergence of threat measures, such as the ones reported above, can be driven by different factors. Figure 4 gives just one illustration of this (see Appendix E for more examples). The figure was derived from another pair-wise operation, subtraction, of the threat measures of the U.S. and Russia for the entire period, namely, $threat(C_{US}) - threat(C_{RU})$. This operation yielded a series of differences in the cosine similarity between the two countries, over the whole elements of W_C . These differences were sorted in descending order, and then their top (distinctively U.S.) and bottom (distinctively Russian) twenties were extracted with the corresponding tokens for




FIGURE 3: Correlations of Threat Measures among Council Members, 2018-2019

the illustration.

As the figure depicts, relative to Russia, the U.S. tends to emphasize concepts such as "nations" (cosine-similarity difference: 0.578), "sexual_gender-based_violence" (0.433), "north_korea" (0.380), "bulgaria" (0.371), and "generations (0.364)" in conceiving the notion of threat. Similarly, "statement_delivered" (-0.465), "auspices" (-0.460), "non-aligned_movement" (-0.456), "consolidation" (-0.400), "mechanisms" (-0.378), and so forth have received more salience in Russia's threat conception, in comparison with that of the U.S. Among these, a sizable number of tokens, such as "generations," "state-ment_delivered," and "auspices," are perhaps not informative in the present context, as they are likely to indicate the diverging dispositions of the two countries (even the disposition of the council's translators in the case of Russia) towards the particular use of words and expressions in their statements involving the notion of threat. In other words, they are not likely to reflect any substantive policy differences. Presumably, these consequential differences are more directly captured by tokens such as "sexual_gender-based_violence," "north_korea," "hamas" (0.364), "afghan" (-0.337), and "drug" (-0.370). Although a substantial amount of prior knowledge and interpretation is needed to establish the relevance or irrelevance of a given concept, this type of analysis can be a useful screening device for detecting divisive issues and entities that might hinder the convergence of the threat conceptions of two actors.




FIGURE 4: Distinctive Threat-Associated Tokens for the U.S. (top) and Russia (bottom)

A similar analysis can be performed on two threat measures representing different time periods. For example, Figure 5 was yielded by $threat(C_{2015-2019}) - threat(C_{1994-1999})$, where $C_{2015-2019}$ and $C_{1994-1999}$ denote the aggregate speech sets for the periods of 2015-2019 and 1994-1999, respectively. While still containing noisy elements, the figure unveils particular issues and entities that constitute substantially changed aspects of the threat notions discussed in the council over the past quarter of a century. More specifically, among the tokens strongly associated with T_C in the late 1990s, "peace_building" (-0.686), "east_timor" (-0.613), "federal_republic_yugoslavia" (-0.381), "serb" (-0.316), "kuwait" (-0.307), etc. had lost such associations in the late 2010s. Instead, the latter period saw a surge of attention to the issues symbolized by such tokens as "transnational_organized_crime" (0.765), "sahel" (0.660), "fight_terrorism" (0.593), "sustaining_peace" (0.583), and "gulf_guine" (0.527). This semi-automated retrieval of threat-associated concepts over years can again offer an efficient means for capturing the changing facets of global security, as seen inside the council chamber. Appendix E provides further examples of this usage of threat measures.




FIGURE 5: Distinctive Threat-Associated Tokens for 2015-2019 (top) and 1994-1999 (bottom)

5 Discussion

While the measurements and analyses reported above are far from being exhaustive, they still contribute some important and useful insights to the literature on council politics. First of all, it should be noted that many of the results obtained so far confirm existing observations about the council. For example, it is not surprising to learn that drug-related and Afghan-related notions (e.g., "drug," "drug_trafficking," "afghan," "afghanistan") are among the most salient aspects of Russia's threat conceptions, given that the trafficking of heroin and opium produced in its southern neighbor, Afghanistan, is one of the security concerns that Russia has persistently raised in the council chamber over years [55]. One can also find abundant supporting evidence in preceding studies for other thematic (e.g., climate change, sexual violence) and geographic (e.g., Sahel, Iran, North Korea) focuses of council members, as well as for topics that have strong rallying effects on these members, such as the global terrorism [7, 19, 20, 24, 25]. Apparently, word embedding, on which all the above analyses rest, seems to quantify only those that are already known, or are intuitively clear about the council and its members.

On a few crucial points, however, word embedding has also revealed insights that deviate considerably from the dominant understandings of council politics. These in-

sights concern the structural configuration of council members, particularly the P5, as seen through their threat conceptions. Unlike the frequent depiction of the council as an arena for great-power competition, most typically, between its close-knit Western members (the U.S., the U.K., and France) on one hand and Russia and China on the other [5, 17, 19, 23], the above results support a more nuanced picture (Figure 2). According to the latter image, Russia has not particularly been far removed from its Western counterparts (whose relative cohesion is undisputed), definitely in comparison with China, and perhaps also in comparison with other elected members (Figure 3). As the previous analysis illustrates (Figure 4), the two sides can differ on numerous specific issues, but they still have, at least in relative terms, some common ground regarding what matters should be considered as common threats to global security, even if there is not total policy unity concerning how to handle these matters.

Moreover, the above results also indicate that Russia and China have not particularly been close to each other. Indeed, in comparison with the other permanent members, and even with other elected members, there is nothing conspicuous about the correlation of their conceptions of security threat. Rather, China is almost equally distanced from any other council member in this regard, which, despite its recent rise as a global power, is largely consistent with Communist China's historical role in the council as a relatively autonomous power, with a regional, rather than global, focus [17, 56]. Meanwhile, the lack of convergence of threat conceptions between Russia and China might not prevent the two countries from taking common action (e.g., exercising veto) against what they see as the U.S. and its allies' abuse of the council on some occasions, such as recent confrontations on Syria and Venezuela. Yet, it certainly does not help them to forge lasting relationships, including a durable alliance [57].

Furthermore, the pattern of cross-national divergence and convergence described so far has been largely stable over time. Although the specific issues and entities that characterize security threat at a given moment can vary considerably from time to time (see Figure 5), the configuration of council members' threat conceptions itself appears to have been fairly robust against these temporal changes (see Appendix D). As long as this configuration contains persistent elements of heterogeneity, it gives a strong caution against another strand of views on the council, namely, seeing too much unity, autonomy, or even agency in the working of this body [3, 8, 9]. On the other hand, it neither implies that the council has been losing a foundation for serving some of its crucial functions, whether the collective legitimization of the use of force, great-power concert, or other governance functions, such as deploying peacekeepers [58–60]. Despite frequent suggestions to the contrary, the durable nature of a certain degree of accord among powerful members of the council still allows cautious optimism for its continued relevance.

Lastly, the validity of these observations could be questioned on several grounds. As has already been discussed, the overall correlation of threat measures between two countries might be confounded to some extent by factors that do not necessarily reflect relevant policy differences. Alternatively, one could also suspect that the measurements employed here were too restrictive for the purpose of this study. That is, it might seem far-fetched that any meaningful insight about structural aspects of the council could be derived from investigating just one or two notions ("threat" and "threats"), no matter how

important these notions are. Appendix G reports on further analyses that partially address these concerns by manipulating the composition of T_C and W_C in different ways. Qualitatively, these manipulations, as long as they remained in a certain range, brought no radical change to the original results. For instance, even if the original elements of T_C , "threat:NN" and "threat:NNP," were replaced altogether with other (yet somewhat related) tokens such as "peace:NN" and "security:NN," the resultant configuration of council members remained almost indistinguishable from that described above. While still being limited in scope, these additional examinations suggest that the obtained insights are not just the product of contingent factors.

6 Conclusions

Having been almost regularly held for more than seven decades, the public meetings of the UN Security Council have become a highly institutionalized, perhaps even ritualized, place. As is well-known, a substantial part of the council's decisions are often made elsewhere. A large number of important decisions have actually been made literally behind closed doors, namely, in informal consultations, which often involve only P5 representatives [17, 60]. This does not mean, however, that open deliberations in council meetings are unimportant. Over time, these deliberations have left behind a massive amount of signals or traces that contain valuable information on council members and other participants, which illuminates their policy preferences, normative ideas, and conceptual understandings. By systematically recovering these signals with adequate processing tools, one can gain useful insights about the council and its politics.

Focusing on one of the most consequential notions for the council's policy making, namely, "threat to the peace," this study purports to be one of the first illustrations of such an endeavor. By employing word embedding and other advanced techniques for text analysis, the study proposes a novel way for measuring the threat conceptions of different actors, over different time periods. These measurements constitute a significant methodological advance with further potential applications, as they pioneer ways for rigorously operationalizing conceptual and normative factors in international relations, which have not necessarily been amenable to quantitative treatment. Furthermore, the study systematically conducted this measurement using a large body of texts, with considerable cross-national and temporal variations. The resultant patterns of convergence and divergence among different council members concerning their threat conceptions are intricate, yet relatively stable over time. These patterns imply that some modifications are required to the previous understanding of council politics.

There are still a number of remaining challenges, which will be addressed in future work. Among them, the following three are pressing. First, extending the time horizon of analysis far back beyond the post-Cold War period would make the whole endeavor obviously more dynamic, and probably more interesting. However, this would pose a formidable methodological challenge, because of the relative paucity of text data, along with the unsettled formatting conventions of meeting records for earlier periods. Second, there is still an enormous amount of untapped information in the derived embedding vec-

tor space. In focusing on several threat-related tokens and their relations with thousands of other tokens, the present study delved into only a limited area of this vast space. By leveraging more global relationships among the entire set of word embeddings, which could be uncovered, for example, by clustering algorithms, a more holistic picture of council members' configuration could be obtained. Third, beyond somewhat cursory plausibility checks, more systematic and rigorous validation is needed [11]. Unlike preceding sociological studies, which have taken pains to validate trained embedding vectors [52, 53], there are few corresponding historical surveys or interviews that allow a direct investigation of the embeddings derived from the statements of council delegates. For more indirect approaches, for example, it is still possible to compare the induced semantic associations in the council corpus with those derived from other corpora (e.g., national security documents). Such external validation would also be helpful for contextualizing the findings of this study, which were obtained exclusively from discursive dynamics inside the council chamber.

Acknowledgements

The author highly appreciates the research assistance provided by Tomoyuki Matsuoka. The author would also like to thank the organizers and attendants of the following events for the valuable occasions they offered him for presenting the study and obtaining feedback: the 2nd Annual Conference on Politics and Computational Social Science (PaCSS) at Georgetown University; the 2019 Annual Convention of the Japanese Association of International Relations (JAIR) in Omiya; European Research Center for Political Culture (ERCAM) Workshop at the University of Bucharest; and Social Fabrics Research Lab (FABLAB) Seminar on Big Data Analysis in International Relations at the West University of Timișoara. Finally, the author acknowledges the financial support (Grant Number: 18H03621) provided by the Japan Society for the Promotion of Science (JSPS).

References

- [1] Anatol Rapoport. *Fights, games, and debates*. University of Michigan Press, Ann Arbor, 1960.
- [2] Martha Finnemore. *The purpose of intervention: changing beliefs about the use of force*. Cornell University Press, Ithaca, 2003.
- [3] Ian Johnstone. *The power of deliberation: international law, politics and organizations*. Oxford University Press, Oxford ; New York, 2011.
- [4] Simon Chesterman, Ian Johnstone, David Malone, and Thomas M. Franck. *Law and practice of the United Nations : documents and commentary*. Oxford University Press, Oxford ; New York, NY, second edition. edition, 2016.
- [5] Brian Frederking. *The United States and the Security Council : collective security since the Cold War*. Routledge, London ; New York, 2007.

-
- [6] David Malone. *The international struggle over Iraq : politics in the UN Security Council 1980-2005*. Oxford University Press, Oxford ; New York, 2006.
 - [7] Ramesh Chandra Thakur. *The United Nations, peace and security: from collective security to the responsibility to protect*. Cambridge University Press, Cambridge ; New York, second edition. edition, 2017.
 - [8] Ian Johnstone. Security Council Deliberations: The Power of the Better Argument. *European Journal of International Law*, 14(3):437–480, 2003. doi: 10.1093/ejil/14.3.437.
 - [9] Ian Hurd. *After anarchy : legitimacy and power in the United Nations Security Council*. Princeton University Press, Princeton, N.J., 2007.
 - [10] James A. Evans and Pedro Aceves. Machine Translation: Mining Text for Social Theory. *Annual Review of Sociology*, 42(1):21–50, 2016. doi: doi:10.1146/annurev-soc-081715-074206.
 - [11] Justin Grimmer and Brandon M. Stewart. Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts. *Political Analysis*, 21: 267–297, 2013. doi: 10.1093/pan/mps028.
 - [12] Christopher D. Manning. Computational Linguistics and Deep Learning. *Computational Linguistics*, 41(4):701–707, 2015. doi: 10.1162/COLI_a_00239.
 - [13] Tomas Mikolov, Kai Chen, Greg S. Corrado, and Jeffrey Dean. Efficient Estimation of Word Representations in Vector Space, 2013.
 - [14] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. GloVe: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.
 - [15] Ray T. Donahue and Michael H. Prosser. *Diplomatic discourse : international conflict at the United Nations- addresses and analysis*. Ablex Pub. Corp., Greenwich, Conn., 1997.
 - [16] Natalie Florea Hudson. *Gender, human security and the United Nations : security language as a political framework for women*. Routledge, New York, 2010.
 - [17] David L. Bosco. *Five to rule them all : the UN Security Council and the making of the modern world*. Oxford University Press, Oxford ; New York, 2009.
 - [18] Adam Roberts, Dominik Zaum, and International Institute for Strategic Studies. *Selective security : war and the United Nations Security Council since 1945*. Routledge, Abingdon, 2008.
 - [19] Karen A. Mingst, Margaret P. Karns, and Alynna Lyon. *The United Nations in the 21st century*. Westview Press, Boulder, CO, fifth edition. edition, 2017.

-
- [20] Thomas G. Weiss, David P. Forsythe, Roger A. Coate, and Kelly-Kate Pease. *The United Nations and changing world politics*. Westview Press, Boulder, Colorado, eighth edition. edition, 2017.
 - [21] Thomas G. Weiss and Rorden Wilkinson. *International organization and global governance*. Routledge, Milton Park, Abingdon, Oxon ; New York, NY, second edition. edition, 2018.
 - [22] Martin Binder. *The United Nations and the politics of selective humanitarian intervention*. Springer Berlin Heidelberg, New York, NY, 2017.
 - [23] Sebastian von Einsiedel, David M. Malone, and Bruno Stagno Ugarte. Introduction. In Sebastian von Einsiedel, David M. Malone, and Bruno Stagno Ugarte, editors, *The UN Security Council in the twenty-first century*, pages 1–26, Boulder, Colorado, 2016. Lynne Rienner Publishers.
 - [24] Sebastian von Einsiedel, David M. Malone, and Bruno Stagno Ugarte. *The UN Security Council in the twenty-first century*. Lynne Rienner Publishers, Boulder, Colorado, 2016.
 - [25] Vaughan Lowe, Adam Roberts, Jennifer Welsh, and Dominik Zaum. *The United Nations Security Council and war : the evolution of thought and practice since 1945*. Oxford University Press, Oxford ; New York, 2008.
 - [26] Steve Chan. Power, Satisfaction and Popularity: A Poisson Analysis of UN Security Council Vetoes. *Cooperation and Conflict*, 38(4):339–359, 2003. doi: 10.1177/0010836703384001.
 - [27] Carla Monteleone. Coalition building in the UN Security Council. *International Relations*, 29(1):45–68, 2015. doi: 10.1177/0047117814552140.
 - [28] Patrik Johansson. The Humdrum Use of Ultimate Authority: Defining and Analysing Chapter VII Resolutions. *Nordic Journal of International Law*, 78(3): 309–342, 2009.
 - [29] Peter Wallensteen and Patrik Johansson. The UN Security Council: decisions and actions. In Sebastian von Einsiedel, David M. Malone, and Bruno Stagno Ugarte, editors, *The UN Security Council in the twenty-first century*, pages 27–54, Boulder, Colorado, 2015. Lynne Rienner Publishers.
 - [30] Alexander Baturo, Niheer Dasandi, and Slava J. Mikhaylov. Understanding state preferences with text as data: Introducing the UN General Debate corpus. *Research & Politics*, 4(2):2053168017712821, 2017. doi: 10.1177/2053168017712821.
 - [31] Lise Morjé Howard and Alexandra Stark. How Civil Wars End: The International System, Norms, and the Role of External Actors. *International Security*, 42(3): 127–171, 2018. doi: 10.1162/ISEC_a_00305.

-
- [32] Mirco Schönfeld, Steffen Eckhard, Ronny Patz, and Hilde van Meegdenburg. The UN Security Council debates 1995-2017. *arXiv*, 1906.10969:13, 2019.
 - [33] Christopher Lucas, Richard A. Nielsen, Margaret E. Roberts, Brandon M. Stewart, Alex Storer, and Dustin Tingley. Computer-Assisted Text Analysis for Comparative Politics. *Political Analysis*, 23(2):254–277, 2015. doi: 10.1093/pan/mpu019.
 - [34] John Wilkerson and Andreu Casas. Large-Scale Computerized Text Analysis in Political Science: Opportunities and Challenges. *Annual Review of Political Science*, 20(1):529–544, 2017. doi: 10.1146/annurev-polisci-052615-025542.
 - [35] Jeremy A. Frimer, Karl Aquino, Jochen E. Gebauer, Luke Zhu, and Harrison Oakes. A decline in prosocial language helps explain public disapproval of the US Congress. *Proceedings of the National Academy of Sciences*, 112(21):6591–6594, 2015. doi: 10.1073/pnas.1500355112.
 - [36] Justin Grimmer. A Bayesian Hierarchical Topic Model for Political Texts: Measuring Expressed Agendas in Senate Press Releases. *Political Analysis*, 18(1):1–35, 2010. doi: 10.1093/pan/mpp034.
 - [37] Justin Grimmer and Gary King. General purpose computer-assisted clustering and conceptualization. *Proceedings of the National Academy of Sciences*, 108(7):2643–2650, 2011. doi: 10.1073/pnas.1018067108.
 - [38] Michael Laver, Kenneth Benoit, and John Garry. Extracting Policy Positions from Political Texts Using Words as Data. *American Political Science Review*, 97(2): 311–331, 2003. doi: 10.1017/S0003055403000698.
 - [39] Burt L. Monroe, Michael P. Colaresi, and Kevin M. Quinn. Fightin’ Words: Lexical Feature Selection and Evaluation for Identifying the Content of Political Conflict. *Political Analysis*, 16(4):372–403, 2008. doi: 10.1093/pan/mpn018.
 - [40] Kevin M. Quinn, Burt L. Monroe, Michael Colaresi, Michael H. Crespín, and Dragomir R. Radev. How to Analyze Political Attention with Minimal Assumptions and Costs. *American Journal of Political Science*, 54(1):209–228, 2010. doi: 10.1111/j.1540-5907.2009.00427.x.
 - [41] Takuto Sakamoto and Hiroki Takikawa. Cross-National Measurement of Polarization in Political Discourse: Analyzing floor debate in the U.S. and the Japanese legislatures. *Proceedings of 2017 IEEE International Conference on Big Data*, abs/1711.02977:3022–3028, 2017.
 - [42] Hiroki Takikawa and Takuto Sakamoto. The moral–emotional foundations of political discourse: a comparative analysis of the speech records of the U.S. and the Japanese legislatures. *Quality & Quantity*, 2019. doi: 10.1007/s11135-019-00912-7.

-
- [43] Azusa Katagiri and Eric Min. The Credibility of Public and Private Signals: A Document-Based Approach. *American Political Science Review*, 113(1):156–172, 2019. doi: 10.1017/S0003055418000643.
 - [44] Gary King and Will Lowe. An Automated Information Extraction Tool For International Conflict Data with Performance as Good as Human Coders: A Rare Events Evaluation Design. *International Organization*, 57:617–642, 2003.
 - [45] Gary King, Jennifer Pan, and Margaret E. Roberts. How Censorship in China Allows Government Criticism but Silences Collective Expression. *American Political Science Review*, 107(2):326–343, 2013. doi: 10.1017/S0003055413000014.
 - [46] Gary King, Jennifer Pan, and Margaret E. Roberts. How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, Not Engaged Argument. *American Political Science Review*, 111(3):484–501, 2017. doi: 10.1017/S0003055417000144.
 - [47] Tabitha Bonilla and Justin Grimmer. Elevated threat levels and decreased expectations: How democracy handles terrorist threats. *Poetics*, 41(6):650–669, 2013. doi: <http://dx.doi.org/10.1016/j.poetic.2013.06.003>.
 - [48] Juraj Medzihorsky, Milos Popovic, and Erin K. Jenne. Rhetoric of civil conflict management: United Nations Security Council debates over the Syrian civil war. *Research & Politics*, 4(2):2053168017702982, 2017. doi: 10.1177/2053168017702982.
 - [49] John W. Mohr, Robin Wagner-Pacifici, Ronald L. Breiger, and Petko Bogdanov. Graphing the grammar of motives in National Security Strategies: Cultural interpretation, automated text analysis and the drama of global politics. *Poetics*, 41(6): 670–700, 2013. doi: <https://doi.org/10.1016/j.poetic.2013.08.003>.
 - [50] Mirco Schoenfeld, Steffen Eckhard, Ronny Patz, and Hilde van Meegdenburg. Discursive Landscapes and Unsupervised Topic Modeling in IR: A Validation of Text-As-Data Approaches through a New Corpus of UN Security Council Speeches on Afghanistan. *arXiv*, 1810.05572:38, 2018.
 - [51] Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334): 183–186, 2017. doi: 10.1126/science.aal4230.
 - [52] Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 2018. doi: 10.1073/pnas.1720347115.
 - [53] Austin C. Kozlowski, Matt Taddy, and James A. Evans. The Geometry of Culture: Analyzing the Meanings of Class through Word Embeddings. *American Sociological Review*, 84(5):905–949, 2019. doi: 10.1177/0003122419877135.

-
- [54] Vivek Kulkarni, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. Statistically Significant Detection of Linguistic Change. In and others, editor, *Proceedings of the 24th International Conference on World Wide Web, Florence, Italy*, pages 625–635. International World Wide Web Conferences Steering Committee, 2015.
 - [55] Dmitri Trenin. Russia in the Security Council. In Sebastian von Einsiedel, David M. Malone, and Bruno Stagno Ugarte, editors, *The UN Security Council in the twenty-first century*, pages 105–119, Boulder, Colorado, 2016. Lynne Rienner Publishers.
 - [56] Zhu Wenqi and Leng Xinyu. China in the Security Council. In Sebastian von Einsiedel, David M. Malone, and Bruno Stagno Ugarte, editors, *The UN Security Council in the twenty-first century*, pages 83–103, Boulder, Colorado, 2016. Lynne Rienner Publishers.
 - [57] Stephen M. Walt. *The origins of alliances*. Cornell University Press, Ithaca, 1987.
 - [58] David Bosco. Assessing the UN Security Council: A Concert Perspective. *Global Governance*, 20:545–561, 2014. doi: 10.1163/19426720-02004005.
 - [59] Inis L. Claude. Collective Legitimization as a Political Function of the United Nations. *International Organization*, 20(3):367–379, 1966. doi: 10.1017/S0020818300012832.
 - [60] Erik Voeten. The Political Origins of the UN Security Council’s Ability to Legitimize the Use of Force. *International Organization*, 59(3):527–557, 2005. doi: 10.1017/S0020818305050198.