

No Permanent Friends or Enemies: Tracking Relationships between Nations from News

Xiaochuang Han

College of Computing
Georgia Institute of Technology
Atlanta, GA
xc@gatech.edu

Eunsol Choi

Paul G. Allen School of CSE
University of Washington
Seattle, WA
eunsol@cs.washington.edu

Chenhao Tan

Dept. of Computer Science
University of Colorado Boulder
Boulder, CO
chenhao@chenhaot.com

Abstract

Understanding the dynamics of international politics is important yet challenging for civilians. In this work, we explore unsupervised neural models to infer relations between nations from news articles. We extend existing models by incorporating shallow linguistics information and propose a new automatic evaluation metric that aligns relationship dynamics with manually annotated key events. As understanding international relations requires carefully analyzing complex relationships, we conduct in-person human evaluations with three groups of participants. Overall, humans prefer the outputs of our model and give insightful feedback that suggests future directions for human-centered models. Furthermore, our model reveals interesting regional differences in news coverage. For instance, with respect to US-China relations, Singaporean media focus more on “strengthening” and “purchasing”, while US media focus more on “criticizing” and “denouncing”.

1 Introduction

In the context of growing globalization (Baylis et al., 2017), understanding complex international relations is increasingly relevant to our daily life. Yet this is a challenging task due to the inherently dynamic nature of international relations. As Kissinger famously said, “America has no permanent friends or enemies, only interests.” Staying informed becomes even harder in the continuous streams of information from news outlets and social media.

This very availability of such information, however, opens up exciting opportunities for natural language processing to support individuals in understanding international relations. Supervised extraction has been incredibly useful at identifying pre-defined relations and events (Dodington et al., 2004; Mintz et al., 2009) but fails to

capture emerging or complex information needs. Topic models and neural models have been proposed to explore relations between entities without supervision (O’Connor et al., 2013; Chaney et al., 2016; Iyyer et al., 2016). In particular, Iyyer et al. (2016) introduces an unsupervised neural model for tracking relations between fictional characters, and this approach outperforms baselines from topic models and hidden Markov models. In this work, we incorporate linguistic insights into this model to track relation dynamics between nations from news articles.

Our model reconstructs textual information in the embedding space using relation embeddings, as proposed in Iyyer et al. (2016). We integrate simple yet effective linguistic insights: verbal predicates often describe the relationship between entities,¹ while nouns and proper nouns provide the context of this relationship. For example, in “U.S. denounces Russia for its interference in the 2016 election”, *denounce* describes the relation, and *election* and *interference* provide the context. We show that this intuition leads the model to discover relation descriptors that are easier to interpret and less noisy.

Evaluating these exploratory models for subjective tasks poses a challenge as there are no gold labels. Along with the model, we propose new approaches for evaluation. We introduce a quantitative metric which aligns pre-annotated key events with the temporal trends of relationships produced by the models. Since this task requires careful analysis of complex international relations, we conduct in-person user studies with NLP researchers and undergraduate students recruited from political science and linguistics courses. Both quantitative evaluation and human evaluation indicate that our model better rep-

¹We use entities, countries, and nations interchangeably in this work.

resents the dynamic relationships between nations than the prior model (Iyyer et al., 2016): 75.9% of participants preferred our model for finding natural language words describing international relations and 85.5% preferred temporal trends generated by our model.

Finally, we qualitatively explore the context of relations provided by an attention-based mechanism and demonstrate a practical application of our model by studying regional differences in news coverage of relationships between two countries. We conclude with discussions on future directions for buildings models that can support individuals in navigating a large collection of news articles. Our code is available at <https://github.com/BoulderDS/LARN>.

2 Data

We start by introducing our dataset of news articles, the shallow linguistic information that we extract, and our annotation of key events.

News article collection. Our dataset is derived from the NOW corpus, the largest corpus of news articles that is available in full-text format.² The NOW corpus collects the news articles that are available on Google News in 20 English-speaking countries and thus include news articles that span a wide variety of topics, ranging from politics, to sports, to celebrity news from 23K media outlets. In this work, we consider the news articles in recent years, i.e., from January 2016 to June 2018, to facilitate human evaluation.

We consider 12 nations (U.S., Russia, China, UK, Germany, Canada, France, India, Japan, Iran, Israel, and Syria) and the 66 nation pairs between them. To identify mentions of each nation in news articles, we manually construct a set of aliases for each nation to cover common abbreviations and the names of political leaders (e.g., Trump, Putin). On average, each nation has 3.5 aliases. We then use these aliases to find sentences that contain a pair of nations under consideration and obtain 1.2M sentences associated with 634K articles.

Adding Shallow Linguistic Information. To incorporate shallow linguistic knowledge, we process the news article collection for each nation pairs to extract (1) verbal predicates and (2) nouns

²The dataset can be obtained from https://www.corpusdata.org/now_corpus.asp.

Total # of months considered	30
Total # of articles	7.7M
Total # of articles with valid nation pairs	634K
Avg. # of articles per nation pair	9.6K
Avg. # of sentences per nation pair	18.2K
Avg. # of predicates per nation pair	21.6K
Avg. # of nouns & proper nouns per nation pair	201K
Avg. # of key events per nation pair (8 pairs)	4.9

Table 1: Data statistics.

and proper nouns from each sentence. Specifically, we use a dependency parser to detect verbal predicates and their subjects and objects, and only include predicates for which both subjects and objects were found. For sentences with such predicates, we find nouns and proper nouns using part of speech tags. All data processing was done in spaCy (Honnibal and Montani, 2017) and full details can be found in the appendix.

Key events annotation. The main goal of our model is to support the exploration of international relations, which is very challenging to evaluate. To derive quantitative evaluation measures and provide the basic context of international relations, we manually identify key events over the 30 months for 8 most frequently mentioned nation pairs (i.e. US-[China, Russia, UK, India, Canada, Japan, Syria], China-India) by reading through the top Google search results for each two countries and each month.³ We identified roughly five key events per nation pair. For example, one key event for US-China relation is Chairman Xi’s visit to the US in April 2017. A complete list of key events is shown in the appendix. Table 1 summarizes the data statistics.

3 Model

In this section, we formally introduce our model that builds on Relation Modeling Network (RMN) (Iyyer et al., 2016). Our main contribution is to integrate shallow linguistic information (i.e., verbal predicates and nouns/proper nouns) and identify the context of relations.

3.1 Overview

The intuition behind our model follows RMN, i.e., inferring relation embeddings by reconstruct-

³ This annotation is done by the first author and is thus inherently subjective. We add a robustness check in the appendix based on another set of independently annotated key events by the third author. Automatic evaluation shows a similar trend holds for the two sets of annotations.

ing textual information in the embedding space. Specifically, we learn a fixed set of relation embeddings and use a convex combination of these relation embeddings to reconstruct information in sentences that mention both entities.

Our main hypothesis is that relation information is often encoded in verbal predicates and we can obtain more interpretable and robust relations if we focus on predicates. For each pair of entities, we extract information from the predicates in sentences with both entities and reconstruct these predicates using shared relation embeddings. In addition, we use nouns to provide the context for the relations. We refer to our model as Linguistically Aware Relationship Networks (*LARN*).

We now formally define our problem. The input of our model is a collection of news articles. For each entity pair e_i, e_j , we obtain a set of articles, \mathcal{A}_{e_i, e_j} , containing at least one sentence mentioning both entities. We identify sentences where both e_i and e_j occur based on any alias associated with an entity (nation). We extract all the verbal predicates from these sentences in article $a \in \mathcal{A}_{e_i, e_j}$ as $\{p_1^{a, e_i, e_j}, \dots, p_N^{a, e_i, e_j}\}$, as well as the proper nouns or nouns as $\{n_1^{a, e_i, e_j}, \dots, n_M^{a, e_i, e_j}\}$. Preprocessing details could be found in the appendix. We then use GloVe embeddings to represent these words, i.e., $v_{p_k}^{a, e_i, e_j}$ for p_k^{a, e_i, e_j} and $v_{n_k}^{a, e_i, e_j}$ for n_k^{a, e_i, e_j} (Pennington et al., 2014). These word embeddings are static in the entire learning process.

Our model learns relation embeddings $\mathbf{R} \in \mathbb{R}^{K \times d}$, where K is a hyperparameter for the number of relations and d is the dimension of the relation embedding as well as the word embedding. Following Iyyer et al. (2016), we obtain a list of natural language descriptors for each relation using the nearest neighbors of the relation embedding within the 500 most common predicates.⁴ The model also provides (1) a probability distribution over relations for each article, (2) a probability distribution over nouns for each relation between each entity pair, and (3) an embedding for each entity.

Figure 1 describes the overall architecture. We will describe the construction of v_{label}^a in §3.2, a

⁴Refer to the appendix for a version that include all words. The relation descriptors from *RMN* required a manual filtering step in Iyyer et al. (2016), and become unintelligible without the 500 common words constraint. Note that our model produces intelligible descriptors even without the 500 most common predicates constraint.




Figure 1: Model overview of an article with respect to a pair of nations. Our model approximates the predicates from the input article v_{label}^a as a weighted sum of relation embeddings from \mathbf{R} . The blue embeddings are fixed and the green ones are trained. The red cell indicates our attention mechanism. The shaded area is our extension: we use only predicates as the label and provide the context of relations via nouns.

reconstruction of v_{label}^a through a weighted sum of relation embeddings in §3.3, and finally the learning in §3.4.

3.2 Modeling Text with Predicates

We compute the representation for each article to be reconstructed as the sum of bag-of-words embeddings. While Iyyer et al. (2016) considers all the words in a window in which both entities occur, we only consider predicates in the sentences where both entities show up:

$$v_{\text{label}}^{a, e_i, e_j} = \sum_{k=1}^N v_{p_k}^{a, e_i, e_j}.$$

$v_{\text{label}}^{a, e_i, e_j}$ depends on both the news article and the entity pair. In the rest of the paper, We omit e_i, e_j in the superscript here for simplicity, i.e., v_{label}^a .

3.3 Reconstructing Text with Relation Embeddings

We represent each article as a weighted sum of relation embeddings. Assuming d^a represents a weight vector over the K relation embeddings that

sums to 1 for an article a with respect to a pair of entities (e_i, e_j) , we obtain r^a as follows:

$$r^a = \mathbf{R}^\top d^a.$$

d^a can also be thought of as a distribution over relations. This distribution over relations depends on entity pair (v_e^a) , predicate information (v_p^a) , and noun information (v_n^a) . While *RMN* simply takes the average of all words in the sentence, *LARN* focuses on verbal predicates and nouns to capture our intuition that predicates describe the main relations, whereas nouns provide background information to explain the relations. We now describe these three components in details.

Representing predicates and entity pair. We follow *RMN* to construct embedding for each entity pair and for the predicates. The entity pair vector, v_e^a , simply adds the embedding of the two entities. The predicate vector, v_p^a , is equivalent to v_{label}^a except for word dropout during training, i.e., setting b_k^a to be 0 or 1 with a probability of 0.5.

$$v_e^a = v_{e_i} + v_{e_j},$$

$$v_p^a = \sum_{k=1}^N b_k^a \cdot v_{p_k}^a.$$

Representing context with nouns. To understand relations between an entity pair in a sentence, nouns should be considered in addition to predicates. For example, *tariff* is indicative of the relation between US and China in “Originally, Trump favoured the simple imposition of a tariff on products from selected countries, especially China and Mexico”, despite the seemingly positive predicate *favour*. As nouns are much more common than predicates (see Table 1) and not all of them are meaningful for understanding international relations, we employ a weighted sum of noun vectors. We use an attention mechanism (Conneau et al., 2017; Bahdanau et al., 2015) and consider each entity pair as a unique key to compute the attention weights, since the same noun can be interpreted differently across different entity pairs. To this end, we train an attention query embedding q_{e_i, e_j} for each entity pair separately. We further encode the temporal information by concatenating a one-hot vector t^a that indicates the month when the article was published with the noun representation $v_{n_k}^a$. This allows us

to capture the shifts in a word’s meaning over time.

$$h_{n_k}^a = \tanh(W_{\text{proj}} \cdot [v_{n_k}^a; t^a]),$$

$$\alpha_{n_k}^a = \frac{\exp(h_{n_k}^a \cdot q_{e_i, e_j})}{\sum_{k'=1}^M \exp(h_{n_{k'}}^a \cdot q_{e_i, e_j})},$$

$$v_n^a = \sum_{k=1}^M \alpha_{n_k}^a h_{n_k}^a.$$

Finally, we concatenate the three representations as the input to a feedforward network and pass to a softmax layer to create the weight vector d^a .

$$v_{\text{final}}^a = \text{ReLU}(W_{\text{cat}} \cdot [v_p^a; v_e^a; v_n^a]),$$

$$d^a = \text{Softmax}(W_{\text{final}} \cdot v_{\text{final}}^a).$$

This d^a is multiplied with the descriptor matrix \mathbf{R} to get the final representation r^a . Different from *RMN*, we do not consider temporal dependencies between time steps in our model because it is important to understand sudden shifts in international relations rather than assuming that the relations slowly evolve. We also found the temporal dependencies were not helpful empirically in our domain but rather computationally expensive.

3.4 Learning Objective and Summary

The reconstruction objective pushes r^a to resemble v_{label}^a . Our formulation is identical to *RMN*: the loss function consists of a contrastive max-margin loss term, J , and an auxiliary loss term, X , to encourage unique relation embeddings.

$$L(\theta) = J(\theta) + \lambda X(\theta),$$

$$J(\theta) = \sum_{a \in \mathcal{A}_{e_i, e_j}} \sum_{v_{\text{label}'}^a \in \mathcal{N}} \max(0, 1 - \frac{r_a \cdot v_{\text{label}}^a}{\|r_a\| \cdot \|v_{\text{label}}^a\|} + \frac{r_a \cdot v_{\text{label}'}^a}{\|r_a\| \cdot \|v_{\text{label}'}^a\|}),$$

$$X(\theta) = \|RR^\top - I\|,$$

where $v_{\text{label}'}^a$ is a randomly sampled negative example, \mathcal{N} is a collection of them, and λ is a hyperparameter for balancing two loss terms.

4 Evaluation

In this section, we compare our model to *RMN*. For both models, we fixed the number of descriptors to 30 following Iyyer et al. (2016). As tracking dynamic international relations requires careful analysis, we hosted onsite user studies for quality control and in-person feedback. We first describe the model outputs, and then present both quantitative and qualitative evaluation results.

<i>LARN</i> (Linguistically Aware Relationship Network)	Weight	<i>RMN</i> (Relationship Modeling Network)	Weight
denounce, undermine, condemn, punish, oppose	5.46%	seem, thing, though, too, especially	11.79%
leave, tell, ask, know, want	4.93%	range, information, supply, infrastructure, value	9.42%
differ, indicate, affect, regard, determine	4.64%	theresa, emmanuel, yemen, poland, lebanon	8.24%
strengthen, enhance, improve, develop, boost	4.56%	negotiate, establish, impose, propose, negotiation	7.92%
hit, cut, end, kick, beat	4.05%	terrorism, militant, terror, terrorist, condemn	7.54%
buy, manufacture, use, brand, purchase	2.59%	propose, woman, represent, indians, spend	1.69e-09
win, defeat, beat, title, match	2.59%	de, defense, kremlin, cite, pressure	1.62e-09
receive, express, praise, send, acknowledge	2.58%	car, de, bomb, little, note	5.39e-10
offer, provide, deliver, feature, guarantee	2.46%	cost, cite, compare, co, m	2.94e-10
launch, announce, unveil, release, celebrate	2.36%	aid, hand, small, public, round	2.21e-10

Table 2: Relation descriptors of the most/least frequent five relations by *LARN* and *RMN* and their average weights in news articles. The relations generated from *LARN* are more semantically meaningful.

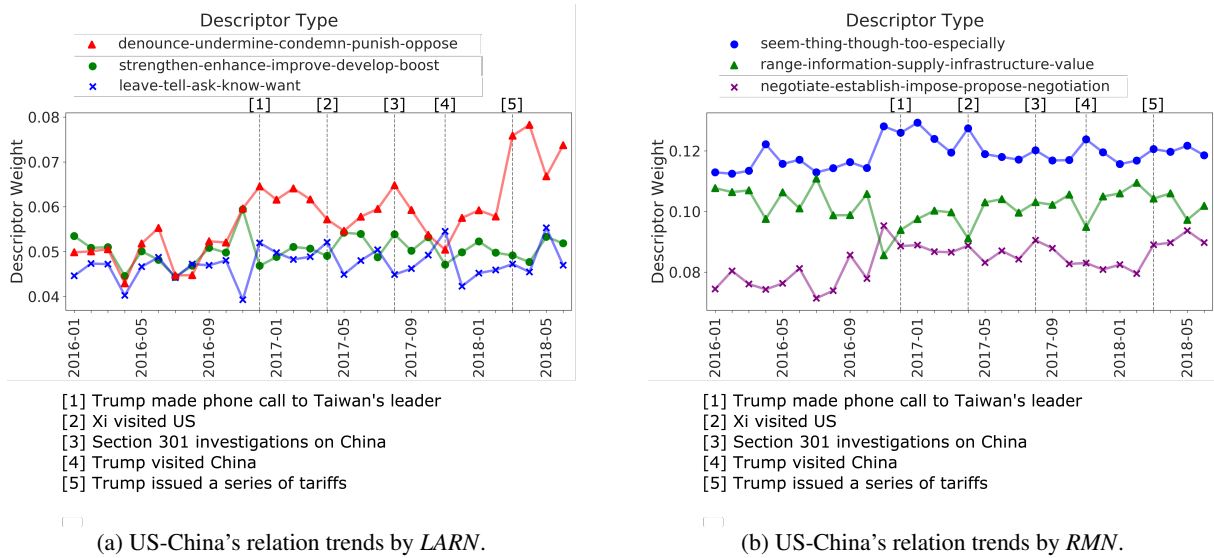


Figure 2: Temporal trends of top three relations between US and China based on *LARN* (Figure 2a), in comparison with results from *RMN* (Figure 2b). We highlight key events during this time period under the x-axes. Similar visualizations are also used in human evaluations. More figures are provided in the appendix.

4.1 Understanding Model Outputs

Given a set of time-stamped news articles and a list of nations of interest, both models provide a set of relation descriptors, where each one **defines** a type of relation and a temporal trend analysis of these descriptors that shows how the relation **evolves**.

Relation descriptors. Table 2 shows the top five and bottom five descriptors from *LARN* and *RMN* sorted by the average weights over all news articles related to the most frequently mentioned eight nation pairs.⁵ By using predicates to describe relations, our descriptors seem to contain more semantically meaningful words. For instance, the top relation in *RMN* consists of exclusively non-content words. Another interesting advantage of our model is that the five relations with the lowest

⁵We focus on the top eight nation pairs to be consistent with human evaluation.

weight have much higher weights in *LARN* than in *RMN*. This suggests that *RMN* tends to generate “useless” relations that do not show up in the data, while even bottom relations in *LARN* remain useful for describing the data.

Temporal trends. We visualize the temporal trends of the most prominent relations between nation pairs. We further provide our annotated key events as the context to interpret these temporal trends. Figure 2 gives an example for US and China. The top three relations based on *LARN* are “denounce”, “strengthen”, and “leave”,⁶ while the top three based on *RMN* are “seem”, “range”, and “negotiate”. We find our model generally aligns better with the key events: for instance, the “denounce” relation peaked around the time that Trump started issuing a series of tariffs based on

⁶For space reasons, we only include the top word in the relation descriptor.




Figure 3: Change rate at key event months vs. other months. Our model demonstrates a clearer relative difference between key event months and other months (66.9%) than *RMN* (50.1%). For the random baseline, we randomly choose the same number of key events and we observe no differences as expected.

our model (Figure 2a), while there do not exist similar fluctuations in Figure 2b based on *RMN*.

4.2 Quantitative Evaluation

We leverage our manually annotated key events to develop a novel automatic metric to evaluate how well the temporal trends from our model are aligned with key events. We define a change rate at each month Δ_t as the weighted average of changes in relation weights in the top three relations:

$$\Delta_t = \sum_{i \in \text{top 3 relations}} w_{t,i} * \frac{|d_{t,i} - d_{t_{prev},i}|}{d_{t_{prev},i}},$$

where $d_{t,i}$ is the average weight for the i -th relation in all the articles published at t , $d_{t_{prev},i}$ is the average weight before t in a window of W preceding months ($\frac{1}{W} \sum_{w=1}^W d_{(t-w),i}$), and weight $w_{t,i}$ is the normalized weight for top three descriptors ($\frac{d_{t,i}}{\sum_{j \in \text{top 3 relations}} d_{t,j}}$). Our results are robust to the choices of W , and we set it to 6 for presentation.

We expect change rates to be greater when significant events happen in international politics. Figure 3 compares the change rate at months where key events occurred with other months for eight nation pairs for which we annotated key events. Both models present more abrupt changes when key events occurred. Unlike *RMN*, our model does not have temporal dependencies between relation distributions over time,⁷ and thus has a higher discontinuity in general. However, even in relative terms, our model fluctuated more substantially than *RMN* when key events occurred.

⁷Please refer to Iyyer et al. (2016)’s Section 3.2.3 for more details on how they incorporated previous time steps.

	Model		P Value
	<i>LARN</i>	<i>RMN</i>	
Descriptor	22	7	8.1e-3
Temporal Trend	106	18	2.3e-16
Nation Pair Matching	38.0 %	45.2 %	0.33

Table 3: Human evaluation results: The first two rows represents the number of votes, while the last row represents % of nation pairs matched correctly with its temporal trend. We did a two-tail binomial test for the relation descriptor and temporal trend evaluation and an independent t-test for the nation pair matching evaluation. We show the p-values in the last column.

We also did a robustness check with another set of independently annotated key events and the results can be found in the appendix. This measure captures whether the model can detect the change points, but does not measure whether the model correctly captures the **semantics** of the key events, i.e., did a negative relation increase after hostile events such as war? To this end, we performed human evaluations.

4.3 Human Evaluation

We hosted three human evaluations with participants from different demographics: undergraduate students from political science classes, graduate students from a computer science department (mostly in NLP), and undergraduate students taking a linguistics class. The total number of participants was 29, roughly equally divided among the three groups. The participants were shown outputs from *RMN* and *LARN*, and asked to choose the output that better aligns with their intuitions. Each participant answered about 10 questions and provided justification for their answers to each question, taking roughly 30 minutes to an hour. Table 3 summarizes the results of our human evaluations.

Relation descriptor evaluation. The participants were shown a list of top five descriptors (as in Table 2) from two models, and prompted to select a set which adequately covers possible relationships that can occur between countries. 75.9% of participants preferred our model.

Temporal trends evaluation. We showed temporal trends between nation pairs annotated with key events, one from *RMN* and the other from *LARN* (as in Figure 2). We asked them to evaluate whether the temporal trends accurately reflect the dynamics in nation-to-nation relationships. Each participant evaluated four randomly chosen nation

pairs. The temporal trend from our model was preferred more frequently (85.5% of total responses).

Nation pair matching. We designed a novel task where we showed the participants the other four temporal trends without annotated key events and asked them to match each trend with the corresponding nation pair from the four candidates, based on their world knowledge about nation-nation relations. Each participant did the matching twice, once for *RMN* and once for *LARN*. The participants found correct temporal trends for 45.2% of entity pairs for *RMN*, and 38.0% of entity pairs for *LARN*, when random pairing would yield 25.0%. The difference between two models here is not statistically significant. Most participants found this task very challenging, as they did not know much about the relationship between certain entity pairs (e.g., a participant said “As an American, there’s no way to know the relation between China and India.”). Even political science students do not perform better than the other two groups.

Discussion. Overall the output from our model is preferred by the participants. We found that political science students paid more attention to detail, took a longer time to finish, and were more ambivalent between the performance of the two models. For example, for temporal trends, they preferred our model for 71.4% of the examples, compared to other groups which preferred our model for 90% of the examples. They also preferred *RMN*’s relation descriptors slightly (42.9% selected our model) and commented that a few concepts from *RMN*, like “infrastructure, supply, and value”, are more concrete (e.g., a participant said “I chose the left one (*RMN*), because it is easy to determine and remember the positive of items such as infrastructure, value, and supply. Those have more positive undertones, while it is easy to gauge negative sentiments with ‘terrorism,’ ‘condemn’ and those.”). As *LARN* encodes the background context specific to each nation pair separately, our relation descriptors do not contain such “concrete” concepts. In the next section, we will discuss **contexts** for relations, where these concepts appear in *LARN*.

5 Further Exploration

We present additional qualitative results to showcase the applications of our model. First, we ex-

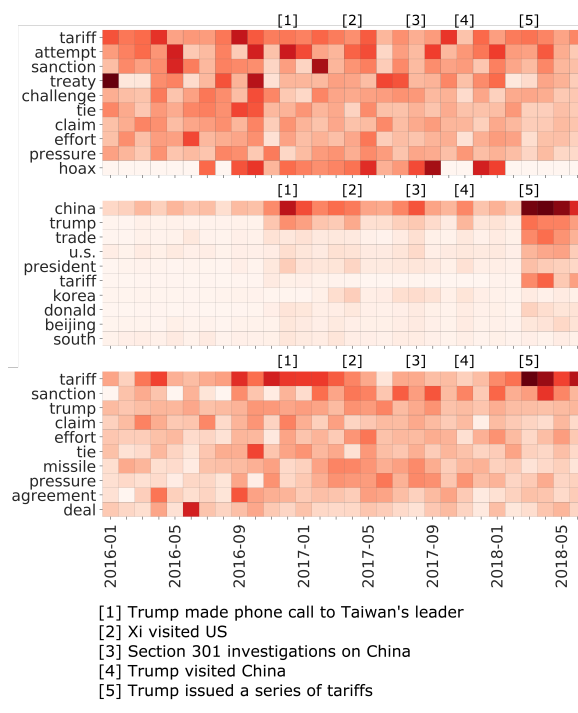


Figure 4: Top contextual words for US-China’s “denounce” relation derived from three approaches. *Top*: Showing word’s avg. attention score in all “denounce” articles each month. *Middle*: Showing word’s frequency in all “denounce” articles each month. *Bottom*: Showing word’s average attention score multiplied by $\log(\text{appearance})$ in each month. All figures are normalized by the global maximum score in the figure.

amine the context (nouns) associated with a relation between two nations based on an attention-based mechanism, which *RMN* does not handle. Second, we perform an in-depth analysis to show how our model can reveal regional differences in news coverage on the same topic.

5.1 Context for Relations

To help users better understand the inferred relations, we offer specific contexts that relate to the inferred relations based on the attention-based mechanism introduced in §3. For each relation, we find articles that place the most weight in that relation, and rank the nouns and proper nouns in those articles by their average attention score (i.e., $\alpha_{n_i}^a$ in §3.3).

The top part of Figure 4 shows the top 10 words for the “denounce” relation in Figure 2a, the temporal trend for US-China relations. Since *RMN* does not support such mechanism, we show the most frequent nouns that occurred in the documents that mention both entities as a baseline

(middle part).⁸ We find that the attended nouns from our model are more informative than frequent nouns: “tariff” is the most attended word; words such as “sanction”, “treaty”, and “pressure” also show up, while the frequency baseline centers around words like “China” and “president”. As we noticed that the frequency baseline can capture alignment with key events, we incorporate attention score and frequency in the bottom part of Figure 4. This augmented version captures informative words (e.g., “tariff”, “sanction”, and “missile”)⁹ and closely aligns with the key events.

5.2 Regional Differences in News Coverage

To further demonstrate the utility of our model, we explore regional differences in news coverage, as “it is possible to build real knowledge by comparing perspectives from different social contexts.”¹⁰ This also relates to the longstanding literature on framing in news coverage, i.e., “selecting some aspects of a perceived reality and make them more salient to promote problem definition/interpretation” (Entman, 1993; Chong and Druckman, 2007).

We picked two countries, Singapore and US, to study US-China relations.¹¹ Using country source of media outlets in the NOW corpus, we found 10K articles from Singaporean media and 5.7K articles from US media on US and China.

Table 4 shows the top five relations sorted by their absolute weight differences between US media outlets and Singaporean media outlets. Singaporean media more frequently use “positive” descriptors such as “strengthen” and “purchase”, whereas US media report negative relations such as “denounce” and “criticize” more frequently. Table 5 shows two example sentences from articles with the most weight in the “denounce” relation. Even though two media sources are focusing on events leading to the same type of relation, a reader who mainly consume news articles in Singapore would get a clearly different impression of US-China relations from those who read US news.

⁸We had a comparison between the top figure and the middle figure in the human evaluation, but we found an error in visualization and thus focus on qualitative comparisons.

⁹“missile” points us to another event related to US deploying missile defense system in South Korea, which also impacts US-China relations.

¹⁰<https://www.publicbooks.org/why-an-age-of-machine-learning-needs-the-humanities/>

¹¹China is not an English speaking country and is thus not in the NOW corpus. Singapore contained the most articles containing both US and China.

Relation descriptor	US	Singapore
strengthen, enhance, improve, develop, boost	4.62%	5.42%
purchase, sell, pay, buy, cost	3.18%	3.86%
denounce, undermine, condemn, punish, oppose	6.21%	5.66%
wag, criticise, accuse, tell, criticize	3.81%	3.29%
flee, destroy, invade, unmake, expel	3.70%	3.32%

Table 4: Top five relations between US and China sorted by absolute weight differences between US media outlets and Singaporean media outlets.

US media: “President Donald Trump is preparing to *impose* a *package* of \$60 billion in annual *tariffs* against Chinese products, following through on a longtime *threat* that he says will *punish* China for intellectual property theft and *create* more American jobs.”

Singaporean media: It does not look like just a *trade war*, but rather the US is trying to *bully* China and the rest of the world in order to make China *concede* economic resources and development *opportunities* to the US and make the US forever big and strong.

Table 5: Example sentences in top-scoring articles on US-China’s “denounce” relation in March 2018. We italicize key predicates (red color) and three most attended nouns (cyan color).

6 Related Work

Prior work (Chambers et al., 2015; Choi et al., 2016; Rashkin et al., 2017) studied entity-entity relations in terms of positive and negative sentiments between them. Similarly, literature on relation extraction (Riedel et al., 2010; Gardner and Krishnamurthy, 2017; Elson et al., 2010; Srivastava et al., 2016) focused on pre-defined relations between a pair of entities in the database schema. In comparison, our work discovers descriptors for relations between entity pairs instead of finding entity pairs matching pre-defined relation schema.

Topic modeling has been an important method to grasp important concepts from a large collection of documents in an unsupervised fashion (Blei et al., 2003; Das et al., 2015; Chang et al., 2009; Schein et al., 2015). Similar to our work, O’Connor et al. (2013) incorporates linguistic insights with topic models to identify event classes and detect conflicts. Our work additionally models

the context of relations through nouns and focuses on exploring the potential of neural models.

Most relevant to our work is Iyyer et al. (2016), which suggests *RMN* better capture dynamic relationships in literature than hidden Markov model (Gruber et al., 2007) and LDA (Blei et al., 2003). Recent work extended and applied *RMN* to other settings such as studying user roles in on-line communities (Wang et al., 2016; Frermann and Szarvas, 2017). Notably, Chaturvedi et al. (2017) suggests HMM with shallow linguistic features (i.e., frame net parses) and global constraints can outperform *RMN* for modeling relations in literature. In this work, we incorporate linguistic insights with *RMN* and apply it to news domain.

Last but not least, researchers have studied the dynamics of media coverage from a wide range of perspectives, ranging from framing (Card et al., 2015; Field et al., 2018), to relationship between ideas (Tan et al., 2017), to quotes of politicians (Niculae et al., 2015; Tan et al., 2018; Leskovec et al., 2009). There is also significant effort for building event databases in political science (Leetaru and Schrodt, 2013), and assisting journalists with tools (Handler and O’Connor, 2017), and dating historical text (Niculae et al., 2014).

7 Conclusion

We investigate the promise of unsupervised neural models for automatically inferring relations between nations. We find that incorporating shallow linguistic information is a simple yet effective strategy for deriving robust and interpretable relations. We develop a novel quantitative evaluation metric for understanding international relations and in-depth human evaluation also confirms the effectiveness of our model. We further show that our models can provide the background of relations using attention score and reveal regional differences for future studies on media framing.

Meanwhile, our work suggests important future directions for using NLP technologies to support individuals in navigating a large collection of news articles. Our participants often find it challenging to infer information simply from temporal dynamics of our inferred relations based on natural language descriptors. It is thus important to incorporate human cognitive preferences in developing such models and provide narratives beyond words such as key events. Furthermore, different popu-

lations pay attention to different parts of information. We need to understand the diversity when developing NLP technologies for end users and provide helpful personalized hints to lower the barrier of benefiting from model outputs.

Acknowledgments

We thank the anonymous reviewers, Dallas Card, and Mohit Iyyer for helpful feedback and discussion. We thank all the participants in our human evaluation from LING 3813 at Georgia Tech by Lelia Glass, students at X-lab at University of Washington, and undergraduate students at University of Colorado Boulder. We also thank Andy Baker and Sven Steinmo for recruiting students in their classes. Choi is supported by a Facebook Fellowship.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- John Baylis, Steve Smith, and Patricia Owens. 2017. *The globalization of world politics: an introduction to international relations*. Oxford University Press.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.
- Dallas Card, Amber E. Boydston, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. The media frames corpus: Annotations of frames across issues. In *ACL*.
- Nathanael Chambers, Victor Bowen, Ethan Genco, Xisen Tian, Eric Young, Ganesh Hariharan, and Eugene Yang. 2015. Identifying political sentiment between nation states with social media. In *EMNLP*.
- Allison June-Barlow Chaney, Hanna M. Wallach, Matthew Connelly, and David M. Blei. 2016. Detecting and characterizing events. In *EMNLP*.
- Jonathan Chang, Jordan L. Boyd-Graber, and David M. Blei. 2009. Connections between the lines: augmenting social networks with text. In *KDD*.
- Snigdha Chaturvedi, Mohit Iyyer, and Hal Daumé. 2017. Unsupervised learning of evolving relationships between literary characters. In *AAAI*.
- Eunsol Choi, Hannah Rashkin, Luke S. Zettlemoyer, and Yejin Choi. 2016. Document-level sentiment inference with social, faction, and discourse context. In *ACL*.

- Dennis Chong and James N Druckman. 2007. Framing theory. *Annual Review of Political Science*, 10:103–126.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *EMNLP*.
- Rajarshi Das, Manzil Zaheer, and Chris Dyer. 2015. Gaussian lda for topic models with word embeddings. In *ACL*.
- George R Doddington, Alexis Mitchell, Mark A Przybocki, Lance A Ramshaw, Stephanie Strassel, and Ralph M Weischedel. 2004. The automatic content extraction (ace) program-tasks, data, and evaluation. In *LREC*.
- David K. Elson, Nicholas Dames, and Kathleen McKeown. 2010. Extracting social networks from literary fiction. In *ACL*.
- Robert M Entman. 1993. Framing: Toward clarification of a fractured paradigm. *Journal of communication*, 43(4):51–58.
- Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. Bradford Books.
- Anjalie Field, Doron Kliger, Shuly Wintner, Jennifer Pan, Dan Jurafsky, and Yulia Tsvetkov. 2018. Framing and agenda-setting in russian news: a computational analysis of intricate political strategies. In *EMNLP*.
- Lea Frermann and György Szarvas. 2017. Inducing semantic micro-clusters from deep multi-view representations of novels. In *EMNLP*.
- Matt Gardner and Jayant Krishnamurthy. 2017. Open-vocabulary semantic parsing with both distributional statistics and formal knowledge. In *AAAI*.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*.
- Amit Gruber, Yair Weiss, and Michal Rosen-Zvi. 2007. Hidden topic markov models. In *AISTATS*.
- Abram Handler and Brendan T. O’Connor. 2017. Rookie: A unique approach for exploring news archives. In *Workshop on Data Science + Journalism at KDD*.
- Matthew Honnibal and Ines Montani. 2017. spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing.
- Mohit Iyyer, Anupam Guha, Snigdha Chaturvedi, Jordan Boyd-Graber, and Hal Daumé III. 2016. Feuding families and former friends: Unsupervised learning for dynamic fictional relationships. In *NAACL*.
- Kalev Leetaru and Philip A Schrodt. 2013. Gdelt: Global data on events, location, and tone, 1979–2012. In *ISA annual convention*.
- Jure Leskovec, Lars Backstrom, and Jon Kleinberg. 2009. Meme-tracking and the dynamics of the news cycle. In *KDD*.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *ACL*.
- Vlad Niculae, Caroline Suen, Justine Zhang, Cristian Danescu-Niculescu-Mizil, and Jure Leskovec. 2015. Quotus: The structure of political media coverage as revealed by quoting patterns. In *WWW*.
- Vlad Niculae, Marcos Zampieri, Liviu P. Dinu, and Alina Maria Ciobanu. 2014. Temporal text ranking and automatic dating of texts. In *EACL*.
- Brendan O’Connor, Brandon M Stewart, and Noah A Smith. 2013. Learning to extract international relations from political context. In *ACL*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*.
- Hannah Rashkin, Eric Bell, Yejin Choi, and Svitlana Volkova. 2017. Multilingual connotation frames: A case study on social media for targeted sentiment analysis and forecast. In *ACL*.
- Sebastian Riedel, Limin Yao, and Andrew McCallum. 2010. Modeling relations and their mentions without labeled text. In *ECML*.
- Aaron Schein, John Paisley, David M Blei, and Hanna Wallach. 2015. Bayesian poisson tensor factorization for inferring multilateral relations from sparse dyadic event counts. In *KDD*.
- Shashank Srivastava, Snigdha Chaturvedi, and Tom M. Mitchell. 2016. Inferring interpersonal relations in narrative summaries. In *AAAI*.
- Chenhao Tan, Dallas Card, and Noah A. Smith. 2017. Friendships, rivalries, and trysts: Characterizing relations between ideas in texts. In *ACL*.
- Chenhao Tan, Hao Peng, and Noah A Smith. 2018. ”you are no jack kennedy”: On media selection of highlights from presidential debates. In *WWW*.
- Alex Wang, William L. Hamilton, and Jure Leskovec. 2016. Learning linguistic descriptors of user roles in online communities. In *NLP+CSS@EMNLP*.

A Appendix

Data: Aliases and Key Events

We use a dictionary of nation aliases to find relevant sentences in the corpus that mentioned at least one nation pair.¹² Table A1 shows the full alias dictionary we used. Table A5 shows the complete list of key events used in this paper.

Nation	Aliases
U.S.	U.S., US, USA, Trump, Obama
China.	China, Chinese, Xi
Syria	Syria, Syrian, Assad
France	France, French, Macron, Hollande
Germany.	Germany, German, Merkel
Canada	Canada, Canadian, Trudeau
Russia	Russia, Russian, Putin
U.K.	U.K., UK, British, Britain, Cameron
India	India, Indian, Modi
Japan	Japan, Japanese, Abe
Iran	Iran, Iranian, Khamenei, Rouhani
Israel	Israel, Israeli, Netanyahu

Table A1: Aliases used for detecting nation mentions.

Preprocessing

We use spaCy to preprocess all sentences (Honni-bal and Montani, 2017). We then use a rule based extractor¹³ to get verbal predicates with detectable subjects and objects. When verbal predicate has a negation, we take its antonym in WordNet (Fellbaum, 1998) if it exists, ignored otherwise. We also extract nouns and proper nouns in sentences using spaCy’s part-of-speech tagging.

We follow the preprocessing steps in Iyyer et al. (2016) for *RMN*. The original *RMN*, which used a fiction dataset, removes the 500 most frequently occurring words and words that occur in fewer than 100 books. We also remove the 500 most frequent words for the input to *RMN*. However, since our news dataset doesn’t have the notion of “books”, we remove 5000 least common words (out of a 200K vocabulary) for *RMN*. Note that this is not done for our model *LARN*.

Implementation Details

We show all hyperparameters used in our study in Table A2. We generally use the same hyperparameter values as *RMN*, except a few model-

¹²We did not use coreference resolution to avoid noisy linking.

¹³https://github.com/NSchradling/intro-spacy-nlp/blob/master/subject_object_extraction.py

Hyperparameter	Value
v_{label}, v_p, v_n dimension	300
p_{drop}	0.5
v_e dimension	50
h_n (<i>LARN</i>) dimension	300 + 30
v_{final} dimension	300
recurrent enforcement α (<i>RMN</i>)	0.5
number of relations K	30
number of training epochs	15
learning rate	1e-3
orthogonal penalty λ	1e-1
number of negative samples	15
batch size	256

Table A2: Hyperparameters.

specific hyperparameters in Table A2. We also use Xavier initialization (Glorot and Bengio, 2010) for all trainable layers as in the original *RMN*.

Relation Descriptors Postprocessing

In addition to finding the nearest neighboring words to the relation embeddings, the original *RMN* also requires a manual filtering process to achieve an interpretable definition of the relation descriptor. To prevent introducing bias in this manual filtering step, we define each relation descriptor directly as its five nearest neighboring words in the input vocabulary. However, since the full vocabulary would contain many uncommon words which could hinder the interpretation of relation descriptors, we limit both models to choose descriptor words from the most frequently occurring 500 words in their own processed input vocabulary (i.e., verb predicates only for *LARN* and all words for *RMN*).

In Table A3, we present the descriptor set by *RMN* and *LARN* without the most common 500 words constraint on descriptor word selection. We find that the descriptor set of our model is robust to this change and does not suffer from a lowered interpretability as much as *RMN*.

Additional Temporal Relation Trends: US-Russia, US-India, US-Syria

See Figure A2, Figure A3, and Figure A4 for three more examples of temporal relation trends between nation pairs.

Additional Influential Background Words:US-Russia, US-India, US-Syria

In Figure A5, Figure A6, and Figure A7, we further show the corresponding attended background words for the top relation descriptor in each of the

<i>LARN</i> (Linguistically aware relationship network)	Weight	<i>RMN</i> (Relationship modeling network)	Weight
denounce, undermine, condemn, decry, legitimize	5.46%	always, something, seem, sense, thought	11.79%
leave,tell,forget,ask,know	4.93%	components, utilization, integrated, optimized, component	9.42%
differ, indicate, extent, affect, imply	4.64%	morroco, theresa, miriam, emmanuel, charlene	8.24%
strengthen, enhance, improve, develop, maximize	4.56%	mandate, authorize, assurances, rescind, ratify	7.92%
hit, fell, slump, edge, cut	4.05%	provocation, bloodshed, reprisals, incursion, hostilities	7.54%
buy, cheap, manufacture, use, shop	2.59%	una, esa, harridan, chicas, hija	1.69e-09
win, clinch, defeat, championship, champion	2.59%	gest, epoque, rol, chiat, tripper	1.62e-09
receive, bestow, redeem, entitle, outpour	2.58%	pharmacologically, offred, condemnable, chapelle, benassi	5.39e-10
offer, provide, deliver, ideal, cater	2.46%	ricoh, powe, mw, sy, hamill	2.94e-10
launch, relaunch, announce, debut, unveil	2.36%	tunstall, booksellers, seiko, pring, reflation	2.21e-10

Table A3: Relation descriptors of the most/least frequent five relations by *RMN* and *LARN* and their average weights in news articles (without constraint on most common words).

three nation pair relations shown in the previous section.

Robustness Check: Another Set of Key Events Annotations

Apart from the key events annotated by the first author, which were used in the change rate analysis in the main paper, the third author also did an independent annotation of the key events as a robustness check. The details of this annotation could found in Table A6. Figure A1 shows the change rate evaluation with respect to this annotation. A similar trend holds for both sets of annotations.




Figure A1: Change rate at key event months vs. other months using an additional set of key events annotations. Our model demonstrates a clearer relative difference between key event months and other months (41.74%) than *RMN* (36.75%). For the random baseline, we randomly choose the same number of key events and we observe no differences as expected.

A Simpler Baseline: Term-frequency Trend of Verbal Predicates


In addition to the comparison with *RMN*, we also checked a simpler baseline by replacing the descriptor weights with term-frequency values of each verbal predicates for each document. Specifically, the term-frequency of a predicate p in document d is $\frac{\text{count}(p)}{\text{length}(d)}$. Note that *document* refers to all sentences where a nation pair is mentioned. We show some temporal trend examples in Figure A8 and also the change rate evaluation result in Table A4. Note that the scales of these results are different from *LARN* and *RMN*, since the predicate words are much more sparse than the 30 descriptors used in *LARN* and *RMN*.

	Change rate		Δ
	key months	other months	
<i>LARN</i>	12.44 %	7.45 %	66.91 %
<i>RMN</i>	6.81 %	4.53 %	50.07 %
Predicate TF	69.86 %	62.36 %	12.02 %

Table A4: Change rate at key event months vs. other months with the predicate term-frequency baseline, compared to *LARN* and *RMN*.


Nation pair	Key events
U.S.-China	(2016-12) Trump made phone call to Taiwan's leader; (2017-04) Xi visited U.S.; (2017-08) Section 301 investigations on China; (2017-11) Trump visited China; (2018-03) Trump started issuing a series of tariffs.
U.S.-Russia	(2016-10) U.S. officially accused Russia's hacking; (2017-04) Syria airstrike; (2017-07) Trump and Putin's first meeting; (2017-11) Trump and Putin's meeting at APEC; (2018-02) Dozens of Russians killed by U.S.-backed Syria attack.
U.S.-Syria	(2016-10) U.S. suspended Syria ceasefire talk; (2017-04) Khan Shaykhun chemical attack and Shayrat missile strike; (2017-10) ISIS 'capital' captured; (2017-11) ISIS's defeat and aftermath; (2018-02) Battle of Khasham.
U.S.-U.K.	(2016-06) U.K. Brexit vote; (2017-01) Theresa May visited U.S.; (2017-12) Trump set a controversial visit to U.K.
U.S.-Canada	(2017-01) Trump said Nafta renegotiation to be started; (2017-04) Trump imposed tariff on Canadian lumber; (2018-06) Canada fought back with retaliatory tariff on U.S. products.
U.S.-India	(2016-06) Modi visited U.S. and met Obama; (2016-12) Trump made a complimentary phone call to Pakistan; (2017-06) Modi visited U.S. and met Trump; (2017-10) U.S. Secretary Of State Tillerson visited India; (2018-06) India imposed retaliatory tariffs on U.S.
U.S.-Japan	(2016-05) Obama gave memorial speech at Hiroshima with Japanese PM Abe; (2016-12) Abe visited Pearl Harbor; (2017-02) Abe visited Washington and met Trump; (2017-11) Trump visited Japan and met Abe; (2018-03) Trump accepted North Korea's invitation for direct nuclear talks.
China-India	(2016-04) Minister of Foreign Affairs meeting; (2016-11) China and India's joint military drill; (2017-02) India to develop a new missile; (2017-05) India refused to attend Belt and Road Summit; (2017-06) Doklam border standoff started; (2017-08) Doklam border standoff ended; (2017-11) China and India's WMCC meeting; (2018-03) China and India to boost trade.

Table A5: Key events annotations and references.



- [1] U.S. officially accused Russia's hacking
- [2] Syria airstrike
- [3] Trump and Putin's first meeting
- [4] Trump and Putin's meeting at APEC
- [5] Dozens of Russians killed by U.S.-backed Syria attack

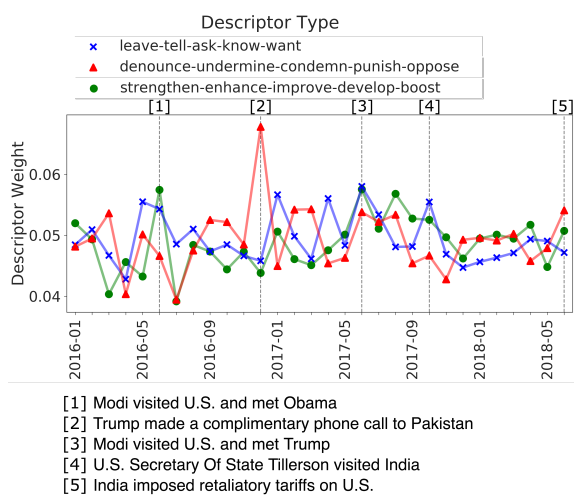
(a) US-Russia's relation trends by *LARN*.



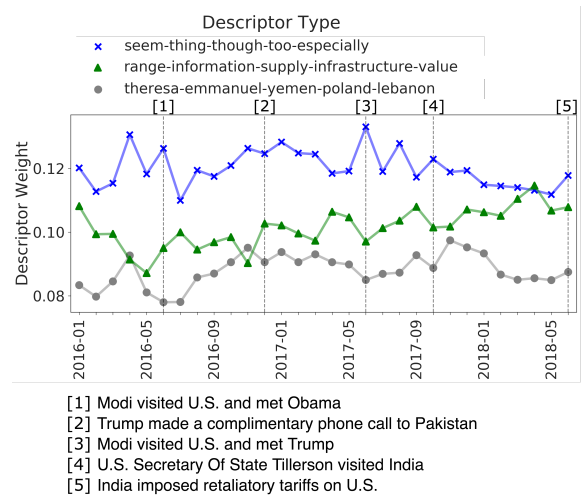
- [1] U.S. officially accused Russia's hacking
- [2] Syria airstrike
- [3] Trump and Putin's first meeting
- [4] Trump and Putin's meeting at APEC
- [5] Dozens of Russians killed by U.S.-backed Syria attack

(b) US-Russia's relation trends by *RMN*.

Figure A2: Temporal trends of top three relations between U.S. and Russia based on *LARN* (Figure A2a), in comparison with results from *RMN* (Figure A2b).




(a) US-India's relation trends by *LARN*.




(b) US-India's relation trends by *RMN*.

Figure A3: Temporal trends of top three relations between U.S. and India based on *LARN* (Figure A3a), in comparison with results from *RMN* (Figure A3b).



(a) US-Syria's relation trends by *LARN*.



(b) US-Syria's relation trends by *RMN*.

Figure A4: Temporal trends of top three relations between U.S. and Syria based on *LARN* (Figure A4a), in comparison with results from *RMN* (Figure A4b).




Figure A5: Influential words for US-Russia’s “denounce” relation derived from three approaches. *Top*: Showing word’s average attention score in all “denounce” articles each month. *Middle*: Showing word’s appearance frequency in all “denounce” articles each month. *Bottom*: Showing word’s average attention score multiplied by $\log(\text{appearance})$ in each month. All figures are normalized by the global maximum score in the figure.




Figure A6: Influential words for US-India’s “leave” relation derived from three approaches. *Top*: Showing word’s average attention score in all “leave” articles each month. *Middle*: Showing word’s appearance frequency in all “leave” articles each month. *Bottom*: Showing word’s average attention score multiplied by $\log(\text{appearance})$ in each month. All figures are normalized by the global maximum score in the figure.

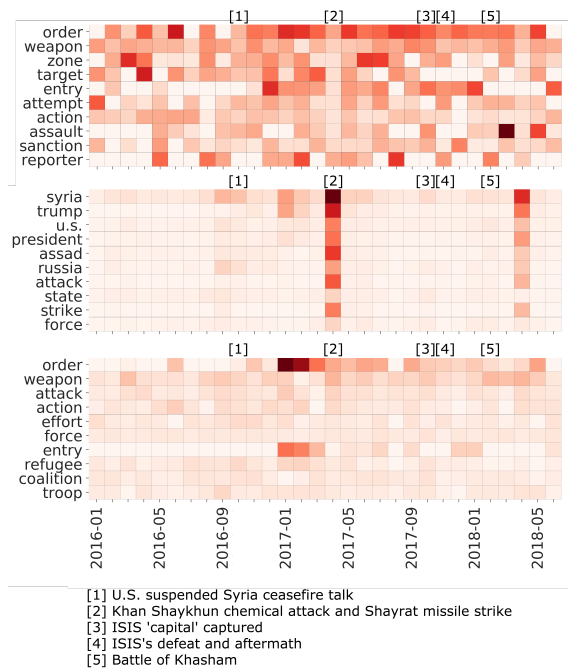
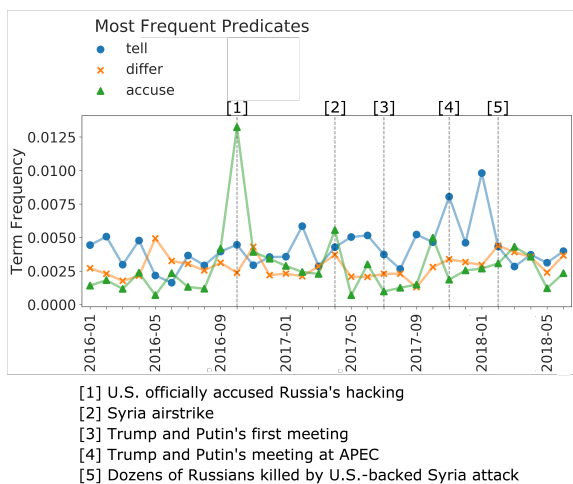



Figure A7: Influential words for US-Syria's "denounce" relation derived from three approaches. *Top*: Showing word's average attention score in all "denounce" articles each month. *Middle*: Showing word's appearance frequency in all "denounce" articles each month. *Bottom*: Showing word's average attention score multiplied by $\log(\text{appearance})$ in each month. All figures are normalized by the global maximum score in the figure.



(a) US-Russia's relation



(b) US-India's relation

Figure A8: Temporal trend examples by the predicate term-frequency baseline.

Nation pair	Key events
U.S.-China	(2016-02) China Sends Missiles to Contested South China Sea Island; (2016-06) Xi urges caution over THAAD deployment in South Korea; (2016-12) Trump-Tsai call; (2017-02) Trump affirms One China policy; (2017-04) Xi visits US; (2017-11) Trump visited Beijing; (2018-03) tariff targets China; (2018-04) China retaliates.
U.S.-Russia	(2016-02) U.S.-Russian deal on a “cessation of hostilities”; (2016-09) Russia and the U.S. announce joint peace plan for Syria; (2016-10) Russian interference with election; (2017-01) The U.S. Office of the Director of National Intelligence releases a declassified version of the Intelligence Community’s assessment that “Putin ordered an influence campaign in 2016 aimed at the U.S. presidential election”; (2017-04) US airstrike Syria; (2017-05) Comey fired and Trump met Russian ambassadors; (2017-07) Trump met Putin at G20 summit; (2018-04) Missile strikes against Syria.
U.S.-Syria	(2016-02) US and Russia agree to enforce new Syria ceasefire; (2016-09) Another ceasefire agreement broke in the same month; (2016-11) Senior Chief Petty Officer Scott Cooper Dayton, 42, of Woodbridge, Virginia, is killed in an improvised explosive device blast near the northern Syrian town of Ayn Issa, becoming the first American casualty in combat in the fight against IS in Syria.; (2017-01) Trump orders ban on Syrian refugees; (2017-04) Trump orders strikes against Syria in response to chemical attack; (2017-07) Another ceasefire agreement that involved united states; (2018-03) U.S. Master Sgt. Jonathan J. Dunbar is killed by a roadside bomb attack in Syria alongside a British serviceman; (2018-04) The US, The U.K. and France carried out missile strikes against Assad’s compounds in response to the Douma chemical attack.
U.S.-U.K.	(2016-04) Obama and Cameron joint news conference; (2017-01) May visit Trump; (2017-02) Parliament debate trump visit; (2017-09) May visit US for UN general assembly; (2018-04) Missile strikes in Syria.
U.S.-Canada	(2016-03) Trudeau visited the White House for an official visit and state dinner on March 10, 2016; (2017-02) Prime Minister Trudeau and President Trump formally met for the first time at the White House on February 13, 2017; (2017-06) The Trudeau government announced that Canada would continue to support coalition operations; (2018-03) President Donald J. Trump announced he will apply across-the-board tariffs, or import taxes, on steel and aluminum; (2018-05) Canada, Mexico, and the EU became subject to the steel and aluminium tariffs later in an announcement on May 31, 2018; (2018-06) Trump comments at G7 summit.
U.S.-India	(2016-02) The Obama administration notified the US Congress that it intended to provide Pakistan eight nuclear-capable F-16 fighters and assorted military goods; (2016-06) Modi visits Obama; (2016-08) U.S., India sign military logistics agreement; (2017-06) Modi visits Trump; (2017-09) Mattis visits India; (2017-10) Secretary of State Rex Tillerson visits India; (2018-03) The India-US 2 + 2 Dialogue postponed.
U.S.-Japan	(2016-05) Abe meets Obama; (2016-07) Abe, U.S. commander agree to carry out defense guidelines in steady manner; (2016-11) Abe meets Trump in New York; (2016-12) U.S. Returns Some Okinawa Land to Japan, but Unease Endures; Abe visits Pearl Harbor; (2017-01) Trump abandons TPP; (2017-02) Abe and Trump presented a united front on dealing with Pyongyang’s nuclear weapon test and multiple missile launches; (2017-04) Japan, U.S. hold missile defense drill in Sea of Japan; (2017-09) Japan, U.S., India vow to work together on strategic port development as China flexes clout; (2017-11) Trump visits Japan; (2018-02) Pence visits Japan; (2018-04) Abe visits US.
China-India	(2016-06) Tashkent on the sidelines of the Shanghai Cooperation Organization Summit; (2016-09) G20 and East Asia Summit; (2016-10) Modi meets Chinese president Xi Jinping on the sidelines of the Goa BRICS Summit; (2017-05) Indias decision to boycott the Belt and Road Initiative (BRI) summit held in Beijing in May; (2017-06) Chinese troops with construction vehicles and road-building equipment began extending an existing road southward in Doklam, a territory which is claimed by both China as well as India’s ally Bhutan.; (2017-08) China and India reached a consensus to put an end to the border stand-off; (2017-09) Xi Jinping Meets with Prime Minister Narendra Modi of India; (2017-11) India joins QUAD; (2018-03) China-India border affairs meeting held in New Delhi; (2018-06) Xi meets Modi.

Table A6: Extra key events annotations for robustness check.