

Not all ANIMALs are equal: metaphorical framing through source domains and semantic frames

Anonymous ACL submission

Abstract

Metaphors are powerful framing devices, yet their source domains alone do not fully explain the specific associations they evoke. We argue that the interplay between source domains and semantic frames determines how metaphors shape understanding of complex issues, and present a computational framework that allows to derive salient discourse metaphors through their source domains and semantic frames. Applying this framework to climate change news, we uncover not only well-known source domains but also reveal nuanced frame-level associations that distinguish how the issue is portrayed. In analyzing immigration discourse across political ideologies, we demonstrate that liberals and conservatives systematically employ different semantic frames within the same source domains, with conservatives favoring frames emphasizing uncontrollability and liberals choosing neutral or more "victimizing" semantic frames. Our work bridges conceptual metaphor theory and linguistics, providing the first NLP approach for discovery of discourse metaphors and fine-grained analysis of differences in metaphorical framing.¹

1 Introduction

Metaphors help us understand and explain our world by transferring what we know about physical, tangible objects to more abstract, hard-to-define concepts and notions, or, as Lakoff and Johnson (2008) succinctly put it, "understanding and experiencing one kind of thing in terms of another". In particular, Lakoff and Johnson (2008) showed that **target** (more abstract) concepts are understood in terms of more concrete, physical **source** domains, such as LOVE IS BATTLE or THEORIES ARE BUILDINGS. Following their seminal work, *conceptual* metaphor theory (CMT) became a produc-

¹Code and data are available at <https://anonymous.4open.science/r/ConceptFrameMet-50F8/>. Pre-trained models will be released upon acceptance.

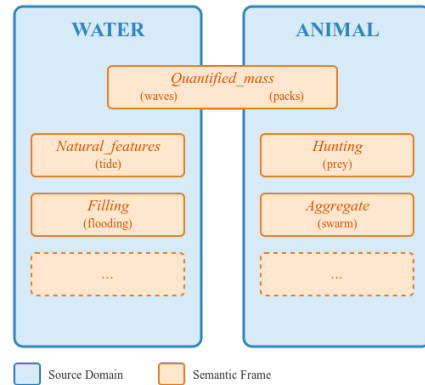


Figure 1: Interaction between semantic frames and domains

tive method to generalize from specific metaphorical expressions and arrive at a more abstract analysis of associations carried over from the source domain. In NLP, work on metaphorical understanding also focuses on mapping metaphors to their source domains (Shutova and Teufel, 2010; Mohler et al., 2016; Mendelsohn and Budak, 2025).

However, the source domain itself does not fully explain which associations are carried over. Consider the following metaphors, both deriving from the source domain of WATER, commonly used in immigration discourse²:

- (1) Illegal aliens continue to *flood* into our country, ruining our economy.
- (2) Maybe a high *tide* raises all boats? *Waves* of immigrants have always enriched us.

The first example compares immigration to *filling* with water, which must be controlled and stopped; the second treats it as a *mass of water* which, however, does not pose a risk of "overflowing" and can be considered a natural *feature of*

²Simplified examples from Mendelsohn and Budak (2025)'s dataset.

061 *landscape*, thus instilling a more positive, accept- 106
062 ing attitude towards immigration. 107

063 Such associations have been explained by lin- 108
064 guistic, constructionalist theories of metaphors 109
065 (Sullivan, 2025, 2013), positing that the associa- 110
066 tions derived from the source domain are related 111
067 to the *semantic frame* of a lexical item used in the 112
068 metaphor (Fillmore and Baker, 2001). In our ex- 113
069 ample, *flood* has a semantic frame of *Filling* which 114
070 emphasizes the movement of water and the nega- 115
071 tive result of such movement. On the other hand, 116
072 *wave* and *tide* have semantic frames of *Quantified*
073 *mass* and *Natural_features* with more neutral 117
074 associations. Semantic frames are generic struc- 118
075 tures that can be applied to many source domains. 119
076 For example, the *Quantified_mass* semantic frame 120
077 can also be applied to the ANIMAL source do- 121
078 main (*packs of immigrants*) or PRESSURE/PHYS- 122
079 ICAL_BURDEN source domain (*loads of illegals*). 123
080 Thus, it is the interplay of the source domain and 124
081 semantic frame that uniquely defines the associa- 125
082 tions of the metaphor, where the source domain 126
083 points to a cluster of associations, and the semantic 127
084 frame allows to pick out specific ones (Figure 1). 128

085 This nuanced understanding of associations is 129
086 particularly important when metaphors are being 130
087 used as *framing devices* to emphasize and high- 131
088 light a particular aspect of an issue or a debate, 132
089 i.e. to frame it (Entman, 1993).³ These discourse 133
090 metaphors or “relatively stable metaphorical pro- 134
091 jections that function as key framing devices within 135
092 a particular discourse over a certain period of time” 136
093 (Zinken et al., 2008; Nerlich and Jaspal, 2012) are 137
094 the focus of our study. Specifically, we show that 138
095 differences in metaphor use between political ide- 139
096 ologies — which has not been fully explained in 140
097 prior work (Mendelsohn, 2024; El Refaie, 2001) 141
098 — can be effectively analyzed through consistent, 142
099 prevalent choice of semantic frames that funnel 143
100 particular associations from the source domain. 144

101 To support our analysis of discourse metaphors, 145
102 we implement a framework that consists of two 146
103 components: a pre-trained language model that de- 147
104 tects metaphors and predicts their semantic frame 148
105 and source domain, and a statistical module that 149

³The term *frame* refers to (at least) two distinct concepts — semantic frames as encoded in FrameNet (Fillmore and Baker, 2001), and media frames, i.e. consistent emphasis of particular aspects of an issue to evoke specific associations in reader’s mind (Entman, 1993); see Otmakhova et al. (2024); Sullivan (2023) for the discussion of their relation to each other and distinctions. To avoid confusion, we use terms *semantic frames* and *media frames*.

106 uses log-likelihood ratio estimation (Rayson and 107
108 Garside, 2000) to highlight salient source domains 109
110 and semantic frames of metaphors within a particu- 111
112 lar discourse compared to a control corpus. 113

114 We apply our framework to analyze metaphor- 115
116 ical framing in two scenarios. First, we discover 117
118 discourse metaphors used to frame *climate change* 119
120 and not only discover well-known source domains 121
122 of climate change metaphors (such as WAR), but 123
124 also reveal more nuanced semantic frame associa- 124
125 tions consistently chosen within the source do- 125
126 mains. Second, we analyze the source domains 127
128 of WATER and ANIMAL in metaphors in immi- 128
129 gration discourse, comparing their use by liberals 129
130 and conservatives (Mendelsohn and Budak, 2025), 130
131 and show that these domains are heterogeneous, 131
132 and that tweets with different political leanings 132
133 routinely different semantic frames of metaphors 133
134 within the same domains to convey particular emo- 134
135 tions towards immigrants. 135

136 In sum, we are the first to connect source do- 136
137 mains and semantic frames in an NLP approach; 137
138 release a manually labeled corpus in the climate 138
139 domain; and apply our framework across two do- 139
140 mains showing it can be used to *discover* dis- 140
141 course metaphors within a particular topic, and 141
142 *highlight nuanced differences* in superficially simi- 142
143 lar metaphor usage. 143

2 Related work 134

Metaphor analysis in NLP Metaphors have at- 135
136 tracted significant attention, with an extensive 136
137 body of work focusing on annotating metaphors 137
138 and creating corpora (Group, 2007; Steen et al., 138
139 2010; Boisson et al., 2025), automatically detect- 139
140 ing them in text (Wang et al., 2025; Zhang and 140
141 Liu, 2022; Uduehi and Bunescu, 2024; Tian et al., 141
142 2024; Reimann and Scheffler, 2025), understand- 142
143 ing (Tong et al., 2024; Ye et al., 2025), genera- 143
144 tion (Joseph et al., 2023; Chakrabarty et al., 2021; 144
145 Veale, 2016), and interpretation, such as the ability 145
146 to handle them in inference and question answering 146
147 tasks (Liu et al., 2022; Sanchez-Bayona and Aggeri, 147
148 2025) (see Ge et al. (2023); Rai and Chakraverty 148
149 (2020) for an extensive survey). 149

150 Concept metaphor theory (CMT) (Lakoff and 150
151 Johnson, 2008) has been widely adopted as a tool 151
152 to improve detection, explanation, and generation 152
153 of metaphors (Jones, 1992; Stowe et al., 2021a; Ge 153
154 et al., 2022; Mao et al., 2023; Tian et al., 2025). 154
155 Some works also use semantic frames as a proxy, 155

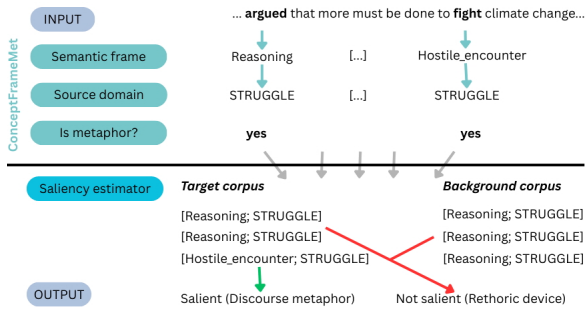


Figure 2: Framework overview showing semantic frames and source domains for the input metaphors (**bold**).

or substitute, for conceptual mappings, in order to detect (Li et al., 2023), or generate (Stowe et al., 2021b) metaphors. We are the first to leverage the interaction between the two frameworks.

According to the constructionist view of metaphors, their conceptual understanding is bound by linguistic constructs (Sullivan, 2013, 2025). Rosen (2018) use construction grammar cues such as argument structure to improve prediction of source domains of metaphors, Jang et al. (2017) collect syntactical patterns associated with a particular source domain for improved metaphor detection, while several studies encode semantic roles or domains (*agent, physical_affliction*) (Gordon et al., 2015; Dodge et al., 2015), establishing schema that consists of multiple lexical instantiations. However, with the exception of a small descriptive study which looks at the differences in semantic frames used for the same CMT metaphor across two languages (Gamonal, 2022), we are the first to use FrameNet-style semantic frames to contrast specific metaphor instantiations and pinpoint the differences in associations carried through them.

Metaphors as framing devices While metaphors are well-known framing devices (Landau et al., 2009; Boeynaems et al., 2017; Semino et al., 2018; Brugman et al., 2019), the link between metaphors and framing is still under-explored in NLP. Several studies approximated framing effects of metaphors through other variables: Prabhakaran et al. (2021) analyzed effects of metaphor usage on reader engagement, while Baleato Rodríguez et al. (2023) modeled metaphors to improve propaganda detection.

A number of works focus on metaphors as framing devices by analysing their source domains in

a particular topic or discourse (Chen et al., 2022; Sengupta et al., 2022; Guan and Zeng, 2024; Li et al., 2024; Meng et al., 2025) or by statistically comparing source domains across political leanings (Sengupta et al., 2024; Wang, 2024; Mendelsohn and Budak, 2025). None of these studies, however, explain why metaphors from the same source domain are used across polar opposite ideologies (Mendelsohn and Budak, 2025), or considered finer distinctions within the source domain. We bridge this gap by adopting semantic frames to extract particular associations from source domains, and saliency as a tool to discover prevalent and distinctive combinations of source domains and semantic frames.

3 Methods

We introduce our two-part framework: a model that combines metaphor detection with predicting their source domain and underlying semantic frames (Section 3.1), and the log-likelihood ratio method used to estimate the saliency of source domains and semantic frames in the corpus (Section 3.2). Figure 2 shows an overview of the framework.

3.1 Predicting metaphors, source domains and semantic frames

In this section we describe ConceptFrameMet – a RoBERTa (Liu et al., 2019) model for detecting metaphors and predicting their source domain and semantic frames (Figure 2, top). The model relies on three components – Semantic frame classifier (Section 3.1.1), Source domain classifier (Section 3.1.2), and metaphor classifier (Section 3.1.3).

3.1.1 Semantic frame classifier

Data. We use FrameNet 1.7 (Ruppenhofer et al.) which contains sentences with semantic frames annotations over 797 classes, with the train, dev, and test splits of 19391/2272/6714 samples (Swayamdipta et al., 2017).

Models. We fine-tune a RoBERTa-base model (Liu et al., 2019) to output a probability distribution over 797 semantic frames for a particular word in a sentence (details in Section A.1.1). We compare two input configurations, with the target word masked out (*immigrants MASK the cities*) and with the target word repeated separately from the context sentence (*SEP immigrants flood the cities SEP flood SEP*). We also prompt two LLMs – Gemini 2.5 and Claude Sonnet 4.0, which we prompt to choose one of 797 classes (details in Section A.2.1).

	Acc	micro-F1	macro-F1
RoBERTa MASK	0.806	0.806	0.053
RoBERTa SEP	0.861	0.866	0.648
Gemini 2.5	0.508	0.508	0.430
Claude Sonnet 4.0	0.736	0.736	0.600
An et al. (2023)	0.899	–	–
Devasier et al. (2024)	–	0.887	–

Table 1: Semantic frame prediction performance on FN1.7 test across three runs (sd < 0.05).

We also include results reported on the same test set for the same task in two recent papers (An et al., 2023; Devasier et al., 2024), which, however, are knowledge-heavy requiring substantial data-augmentation and are thus unsuitable to deploy for our large-scale topic-specific analyses.

Results. As shown in Table 1, our best performing model (RoBERTa SEP) performs comparatively to the knowledge-heavy models in prior work. The LLMs significantly under-perform our fine-tuned RoBERTa models. Accordingly, we use the SEP-style model in all further experiments, which is light-weight and scalable to large data sets.

3.1.2 Source domain classifier

Data. We use the LCC Metaphor Dataset (large) (Mohler et al., 2016), the largest English dataset annotated for source domains, and pre-process it as described in A.1.2. The resulting dataset has 99 source domains, and is randomly sampled with stratification into train, dev, and test of 11704/2509/2509 samples.

Models. We fine-tune a RoBERTa-base model to predict one of 99 source domains (details in A.1.2). Similar to Section 3.1.1, we compare the masked input style with the inputs where the metaphor and the context are separated by SEP. Moreover, since some of the source domain classes are semantically close and easily confused (Shutova and Teufel, 2010; Mohler et al., 2016), we hypothesize that including the semantic frame of the metaphor will help the model to differentiate between domains. To test this, we pass the predicted probability distribution of semantic frames as a frozen feature vector to the model in two ways: by either concatenating it with the RoBERTa-encoded input vector (Frames_CONCAT), or by applying attention over semantic frames (together with the frozen vector as residuals) to highlight those semantic frames that are important for the source domain prediction (Frames_ATTEN). We also compare the

	Acc	P	R	F1
RoBERTa MASK	0.307	0.203	0.184	0.182
RoBERTa SEP	0.833	0.745	0.742	0.740
Frames_CONCAT	0.837	0.759	0.758	0.754
Frames_ATTEN	0.838	0.764	0.757	0.756
Gemini 2.5	0.528	0.419	0.356	0.345
Claude Sonnet 4.0	0.528	0.517	0.452	0.445

Table 2: Source domain prediction performance on the LCC dataset across three runs (sd < 0.05). All metrics are macro-averaged.

fine-tuned models with zero-shot Gemini 2.5 and Claude Sonnet 4.0, where the prompt includes the list of 99 classes (Section A.2.2).

Results. Table 2 confirms that adding semantic frame information to the input boosts source domain prediction performance, with the attention-based model performing best overall. In particular, they achieve improvement of 20 points of macro F1 on underrepresented classes (<10 instances in training data; see Figure 5 in the appendix). Most fine-tuned RoBERTa models again outperform zero-shot LLMs. We use the Frames_ATTEN model in all further analyses.

3.1.3 Metaphor prediction

Data. We use VUA-18, the largest available metaphor prediction dataset (Leong et al., 2018) to train and test our models, and evaluate generalizability through zero-shot testing on two smaller benchmarks – MOH-X (Mohammad et al., 2016) and TroFi (Birke and Sarkar, 2006) (see Section A.1.3 for data statistics).

Models. We integrate our semantic frame and source domain classifiers with a metaphor prediction model. For the latter, we choose MelBert (Choi et al., 2021), which is the strongest pre-trained model among "knowledge-lean" models that do not require augmenting data with additional features. MelBert is a RoBERTa model which captures the inconsistency between the contextual meaning of the metaphor and its literal meaning through contrasting embeddings of the word in context and in isolation. We posit that the encoding of a word in isolation still mixes representations of literal and metaphorical usage; for example, the embedding of the isolated word *flooded* will likely reflect both the senses of *the river flooded the city* and *immigrants flooded the city*, i.e. still be somewhat close to a metaphorical context embedding. We thus replace the word embedding with that of the predicted source domain (WATER), amplifying the

	VUA-18	TroFi	MOH-X
Random baseline	0.222	0.466	0.486
MelBert	0.782	0.631	0.806
FrameBert	0.766	0.620	0.780
Gemini 2.5	0.341	0.633	0.849
Claude Sonnet 4.0	0.330	0.651	0.822
ConceptFrameMet (ours)	<u>0.767</u>	<u>0.634</u>	0.814

Table 3: Binary F1 (positive class) performance on metaphor prediction for models fine-tuned on VUA-18 (train). We show results on VUA-18 (test) and two other test sets. LLM results are zero-shot for all three datasets. Bold=best; underline=second best.

difference between contexts in the metaphorical case (*immigrants* and *water*), and minimizing the gap in literal cases (*river* and *water*).

To account for source domain prediction noise, instead of replacing a word with its predicted source domain, we blend their embeddings, using the confidence score of the source domain prediction (α) as a weight:

$$\mathbf{e}_{\text{blended}} = \alpha \cdot \mathbf{e}_{\text{source_domain}} + (1 - \alpha) \cdot \mathbf{e}_{\text{word}}.$$

We compare the resulting model with FrameBert (Li et al., 2023), a modified version of MelBert and Claude Sonnet 4.0 and Gemini 2.5 which are prompted to identify all metaphors in a given text (Section A.2.3).

Results. Results in Table 3 show that MelBert performs best in-domain (i.e., when tested on the test portion of its training data set) but fails to generalize to other data sets. LLMs perform best on the smaller test sets. ConceptFrameMet (our method) shows the best overall performance, only slightly underperforming the best method across all three test sets, while being substantially more resource efficient than the LLMs. The LLMs tend to over-predict metaphors, so they struggle on VUA-18 where the rate of metaphors is low (15%). However, the results on VUA-18 are indicative of practical usage scenarios where one needs to distinguish metaphors from mostly literal words. This, together with low performance of LLMs on semantic frame and source domain prediction, motivates us to use our light-weight, integrated PLM model for metaphor detection and analysis.

3.2 Detecting discourse metaphors

To distinguish between *discourse metaphors* — which carry stable, prevalent associations that allow to resolve an ambivalence of interpretation of

some issue (Scheufele and Scheufele, 2010) — and metaphors that are merely figures of speech, we define two principles. First, the target domain of discourse metaphors must be a polarizing issue. Second, discourse metaphors are salient in a given discourse (Zinken et al., 2008). We operationalize *saliency* as the log-likelihood ratio (Rayson and Garside, 2000), which is a common technique to compare distributions of items (such as words or, in our case, semantic frames and source domains) between two corpora. In particular, given two corpora C_1 and C_2 , the log-likelihood ratio identifies candidate metaphors that are significantly over-represented in one of the corpora, i.e., reject the null hypothesis H_0 that they are represented equally in both corpora (C_1 and C_2).

$$-2 \ln \lambda = -2[\ell(\theta_0) - \ell(\hat{\theta})] = 2 \sum O_i \ln \left(\frac{O_i}{E_i} \right)$$

where λ = likelihood ratio, $\ell(\theta_0)$ = log-likelihood under H_0 , $\ell(\hat{\theta})$ = maximum log-likelihood (from (C_1, C_2)), O_i = observed frequency, E_i = expected frequency under H_0 .

Source domains and semantic frames with log-likelihood values that reject H_0 with $p > 0.05$ are identified and their relative frequency is compared to decide which of the corpora (C_1 or C_2) they are more strongly associated with.

4 Metaphorical framing through source domains and semantic frames

We now apply our framework explained in Section 3 to (1) *discover* discourse metaphors in texts about a particular topic by comparing metaphors in a topic-specific vs. a general background corpus (Section 4.1); and *contrast* discourse metaphors in texts associated with polar opposite political leanings within a single topic (Section 4.2).

4.1 Metaphorical framing in climate change news

Existing methods for metaphorical framing analysis relies on manual discovery of prevalent source domains in a particular discourse, while NLP approaches require such knowledge of source domains a priori. Here, we apply our framework to automatically identify salient source domains from a large collection of news on climate change, and further analyze them through salient semantic frames.

	Accuracy	F1
Metaphor prediction	0.955	0.976
Source domain prediction	0.935	0.728

Table 4: Performance of ConceptFrameMet on the climate change metaphors dataset. We report binary F1 for the metaphor prediction task, and macro-F1 for the source domain prediction task.

4.1.1 Data and methods

We use a corpus of New York Times articles published between 1986 and 2020 (Fast and Horvitz, 2017; Mendelsohn et al., 2020). We extract all paragraphs which mention "climate change" or "global warming", and then additionally filter them using a ClimateBERT pre-trained model (Webersinke et al., 2021) which was fine-tuned to detect climate-related texts (Bingler et al., 2023). This results in 47K paragraphs (distribution over years in Section A.4). For saliency analysis, we also collect a generic corpus which contains the same number of paragraphs randomly sampled from the same underlying NYT data set ensuring that the paragraphs do not contain our climate keywords.

Next, we use ConceptFrameMet (Section 3.1) to detect metaphors, their semantic frames and source domains, in both the climate and generic corpus. To ensure that the identified metaphors refer to the topic of interest, we retain only sentences that contain terms "climate" or "warming". Finally, we calculate log-likelihood ratios for source domains, and then for semantic frames within each source domain, as described in Section 3.2.

4.1.2 Evaluation

The process above resulted in a corpus of 6,859 sentences that contain metaphors with source domains and semantic frames salient in the climate change corpus ($p < 0.05$). To evaluate the performance of ConceptFrameMet on that corpus, we randomly select a sample of 400 sentences that covers all combinations of salient source domains and semantic frames within them. We presented crowd workers on Prolific with sentences containing highlighted metaphors and asked them to identify up to three source domains from a list of 5 options, comprising the top two domains predicted by ConceptFrameMet and three random distractors plus options OTHER DOMAIN and NO METAPHOR. See Section A.8 for more details including quality control, compensation and instructions We collected four annotators per sample and achieve a

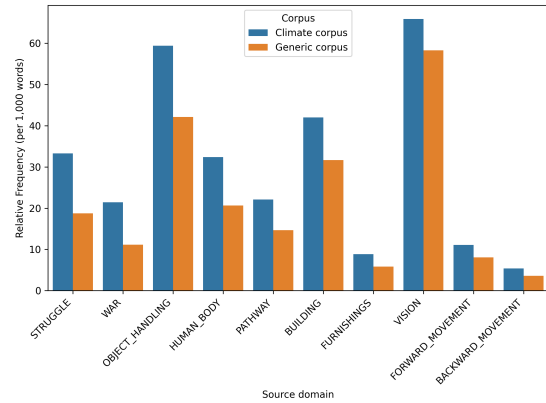


Figure 3: Relative frequency of source domains in climate and generic corpora. The source domains are sorted left to right in order of diminishing saliency

reliable averaged pair-wise annotator agreement of 69% , and use majority voting (with adjudication by one author this paper with a degree in linguistics) to determine the final source domain label. We finally compute model performance against human labels, judging a prediction as correct if the human majority vote corresponds to the top-1 model prediction (Table 4), showing that ConceptFrameMet generalizes well to our new domain.⁴

4.1.3 Results

Source domains Figure 3 presents the 10 most salient source domains in the climate corpus, ordered by log-likelihood ratio, from left to right, while the y-axis indicates frequency, noting that saliency highlights *differences in frequency* of source domains between corpora rather than the *absolute frequency* of a source domain within a corpus. Thus, a frequent source domain (such as VISION) can be ranked lower than a less frequent domain (STRUGGLE), and the most frequent source domain in the climate corpus (MOVEMENT) is not among the most salient ones. Importantly, the discovered source domains are well-aligned with the ones reported in theoretical critical discourse analysis studies. We present examples of metaphors from each of the top 10 source domains, together with the theoretical studies, in Section A.5.

Semantic frames Even within a salient source domain, some metaphors can be used as a figure of speech rather than a discourse (framing) metaphor.

⁴Note that the metaphor prediction F1 is inflated since all samples were predicted as containing metaphors by the model; in this we follow the process and limitations of the metaphor corpus creation suggested by (Mohler et al., 2016)

	Climate corpus		Generic corpus	
	Semantic frame	Metaphors	Semantic frame	Metaphors
STRUGGLE	<i>Hostile_encounter</i> <i>Topic</i> <i>Relation</i>	fight, confront address at odds	<i>Difficulty</i> <i>Resolve_problem</i> <i>Cause_to_end</i>	challenges, hard settle ended
WAR	<i>Hostile_encounter</i> <i>Boundary</i> <i>Judgment_communication</i>	battle, waged war on the front lines assailed, crusaded	<i>Change_of_leadership</i> <i>Invading</i> <i>Aiming</i>	rebel, revolution invasion, intervene target
OBJECT HANDLING	<i>Attempt_suasion</i> <i>Taking_sides</i> <i>Intentionally_act</i>	push (for), press take (a stand), handle take (action, steps)	<i>Getting</i> <i>Intercepting</i> <i>Entity</i>	gets snapping up, caught things, stuff
HUMAN BODY	<i>Confronting_problem</i> <i>Taking_sides</i> <i>Part_orientational</i>	face, confront stance, embrace at the heart of	<i>Social_connection</i> <i>Body_parts</i> <i>First_rank</i>	contacts, close eye for, at the wrist core

Table 5: Examples of salient semantic frames (with SOURCE DOMAINS and Metaphors) in the climate change corpus (left) and the generic corpus (right).

For example, the source domain STRUGGLE (which compares mental hardships to physical ones such as fighting) includes metaphorical use of *argue* (*The ministry argued that...*) which are used in the climate corpus with similar frequency as in general texts, as well as climate-specific STRUGGLE metaphors such as *fight* (*more must be done to fight climate change*).

When comparing a topic-focused corpus against a randomized, generic corpus, log-likelihood ratio estimation allows to achieve two things: first, metaphors that are common figures of speech, used equally frequently irrespective of the text’s topic, which will have a statistically insignificant differences. Second, it points towards prominent *framings* of the topic (as repeatedly used metaphors), as well as absent (but possible) framings (which metaphors are more rare than in general usage).

Table 5 shows examples of semantic frames for top source domains associated with climate change which are salient in climate discourse (left) or are avoided (right). In some source domains, this difference is particularly revealing – for example, climate metaphors in the OBJECT HANDLING source domain tend to use semantic frames denoting intentional actions such as *Attempt_suasion* (push for) or rather than more "passive" semantic frames such as *Intercepting* (caught). Interestingly, while the metaphors on both sides of the WAR source domain express fierce activity, climate change actions are not framed as a revolution (*Change_of_leadership*) or intervention (*Invading*).

In sum, semantic frames as mental templates allow to pick out specific parts of a source domain that correspond to the ways we relate to and act upon the climate crisis. However, they cannot represent the metaphorical frame on their own – as

Table 5 shows, the same semantic frame (for example *Taking_sides*) can be used across different source domains, imbued with different physical imagery. Thus, discovering prevalent metaphorical frames requires both components.

4.2 Conservative vs liberal framing in immigration discourse on Twitter/X

Previous research has shown that, perhaps surprisingly, the same dehumanizing source domains are used to frame immigration by both conservatives and liberals, albeit with different frequency and intensity (Mendelsohn and Budak, 2025). Here, we study the interplay of source domains and semantic frames to shed light on the more nuanced differences of metaphor use across political camps.

4.2.1 Data and methods

To examine difference in metaphorical framing between users with different political affiliation, we use a corpus of 400K US tweets on immigration with automatically predicted metaphoricity scores across seven dehumanizing source domains and predicted scores of political affiliation (liberal vs conservative; Mendelsohn and Budak (2025)). We select two well-studied domains (WATER and ANIMAL) for analysis.

We extract candidate tweets containing metaphors from our chosen source domains as explained in Section A.6, and use ConceptFrameMet to annotate the resulting 14K tweets for metaphors, their semantic frames and source domains. For precision, we only use tweets which have the same source domain annotation at both tweet (original predictions) and metaphor level (our model). Finally, since a metaphor with a WATER or ANIMAL source can refer to another, irrelevant

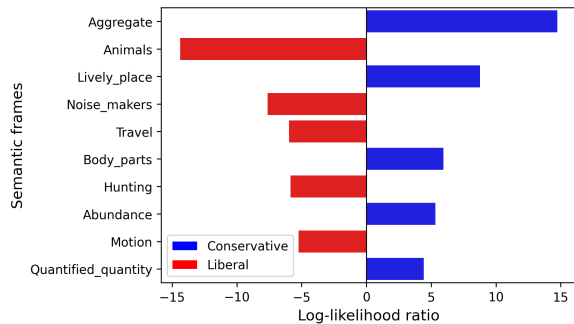


Figure 4: The five semantic frames with highest saliency in the liberal (red) and conservative (blue) data, in the ANIMAL domain.

target rather than immigration (for example, as in *left-winged politicians*), we further filter out tweets where the metaphor does not refer to immigrants (see Section A.7). This results in 544 tweets with source domain ANIMAL and 1,400 for WATER.

We apply log-likelihood ratio to semantic frames within the ANIMAL and WATER source domains to highlight those which are more representative of liberal vs conservative discourse. To do so, we split each subset of tweets with metaphors of a particular source domain into two parts according to their predicted ideology, and use them as C_1 and C_2 , respectively.

4.2.2 Results

While many semantic frames are used by both political leanings (e.g. the top most used semantic frame in the ANIMAL source domain on is *Animals* for both sides, which is a reference to an animal or the word "animal"), their *saliency* differs. Figure 4 shows the top five most salient ANIMAL semantic frames for both leanings. The semantic frame *Animals* is salient in liberal tweets, where it is used to criticize dehumanization ("Families shouldn't be ripped apart and treated like **animals**."). Other prevalent ANIMAL metaphors in liberal tweets are based on semantic frames of *Hunting* ("ICE agents are **predators**, undocumented immigrants are **prey**"), or *Moral_evaluation* ("How is that not **predatory**?"). Conversely, the ANIMAL metaphors in conservative tweets are predominately using semantic frames like *Aggregate* ("Anytime you disrupt a rats **nest** [...] they scatter; "the **swarm** of Muslim migrants"), *Body parts* ("ILLEGALS taking [...] every benefit they can lay their **paws** on"), or *Abundance* ("the illegal immigrants will run **rampant**").

Overall, we observe that the semantic frames

used by conservatives pick out more specific, "colorful", but also more negative parts of the ANIMAL source domain, while liberals often use semantic frames that show lack of agency and power, with immigrants portrayed as victims or passive (i.e. immigrants are *caged, hunted, treated like animals*).

Similarly, within the WATER source domain, the conservative features salient semantic frames like *Filling* ("**Flooding** America with illegals"; "those cities won't take **overflow** Illegal Immigrants") and *Abounding with* ("It's now a **cesspool** of crime", "CA is **saturated** with illegals"), which depict immigration as "too much" that needs to be controlled. On the other hand, the semantic frames on the liberal side have more neutral associations: *Quantified mass* ("The **waves** of desperate migrants", "the **trickle** of immigrants"), *Natural feature* ("Many different immigrant **streams**", "Migration is like the **tide**. You can put a wall up, but it will go around it."), or, again, frame immigrants as powerless victims through such semantic frames as *Killing* and *Death* ("the federales tried to **drown** families") or *Removing* ("this [policy] is in part to **flush out** immigrant parents").

5 Conclusions

Inspired by a constructionist view of metaphor theory, which posits that metaphors are bounded by their linguistic expression, we proposed a framework, ConceptFrameMet, that models metaphors through both their source domains and semantic frames and show how salient semantic frames within source domains allow to capture nuanced differences in framing evoked by the metaphors. We release a pre-trained model with predicts metaphors, their source domains and semantic frames, as well as a statistical model for estimating the saliency of source domains and semantic frames within a particular topic or discourse. We demonstrate the potential of our framework by applying it to two different tasks – *discovering* source domains and semantic frames used in framing climate change news, and *highlighting nuanced differences* in semantic frames used by liberals and conservatives within the same metaphorical source domain. Future work can include integrating it with framework such as MetaNet (Dodge et al., 2015) to enable more sophisticated analysis of semantic frames that support complex, extended discourse metaphors.

5.1 Limitations

Our study has several limitations. First, we use a model trained on a set of source domains used in the LCC corpus (Mohler et al., 2016), which might not be comprehensive enough to cover new topics that are very different from the ones occurring in it. Thus, a larger – or even open – set of source domains is needed, and a method which is sensitive enough to differentiate between a large number of source domains. Relatedly, the question of distinguishing between related or semantically similar source domains (Shutova and Teufel, 2010) is still unanswered, though in Section 3.1.2 we show an example of how semantic frames can serve as a tool to make such distinctions in a more principled way. Finally, in this study we only look at individual semantic frames within a particular domain, and do not consider how they interplay and link with each other in the same document to create more powerful extended metaphors. Thus, we hope to connect our FrameNet-frames based analysis – which operates at lexical level – with the MetaNet (Dodge et al., 2015) framework, which encodes the relationships between the components of the source domain.

References

Kaikai An, Ce Zheng, Bofei Gao, Haozhe Zhao, and Baobao Chang. 2023. **Coarse-to-fine dual encoders are better frame identification learners**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13455–13466, Singapore. Association for Computational Linguistics.

Daniel Baleato Rodríguez, Verna Dankers, Preslav Nakov, and Ekaterina Shutova. 2023. **Paper bullets: Modeling propaganda with the help of metaphor**. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 472–489, Dubrovnik, Croatia. Association for Computational Linguistics.

Julia Bingler, Mathias Kraus, Markus Leippold, and Nicolas Webersinke. 2023. How cheap talk in climate disclosures relates to climate initiatives, corporate emissions, and reputation risk. Working paper, Available at SSRN 3998435.

Julia Birke and Anoop Sarkar. 2006. **A clustering approach for nearly unsupervised recognition of nonliteral language**. In *11th Conference of the European Chapter of the Association for Computational Linguistics*, pages 329–336, Trento, Italy. Association for Computational Linguistics.

Amber Boeynaems, Christian Burgers, Elly A Konijn, and Gerard J Steen. 2017. The effects of metaphorical framing on political persuasion: A systematic

literature review. *Metaphor and Symbol*, 32(2):118–134.

Joanne Boisson, Arif Mehmood, and Jose Camacho-Collados. 2025. **METAPHORSHARE: A dynamic collaborative repository of open metaphor datasets**. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (System Demonstrations)*, pages 509–521, Albuquerque, New Mexico. Association for Computational Linguistics.

Britta C Brugman, Christian Burgers, and Barbara Vis. 2019. Metaphorical framing in political discourse through words vs. concepts: A meta-analysis. *Language and Cognition*, 11(1):41–65.

Tuhin Chakrabarty, Xurui Zhang, Smaranda Muresan, and Nanyun Peng. 2021. **MERMAID: Metaphor generation with symbolism and discriminative decoding**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4250–4261, Online. Association for Computational Linguistics.

Jieyu Chen, Kathleen Ahrens, and Chu-Ren Huang. 2022. **Framing legitimacy in CSR: A corpus of Chinese and American petroleum company CSR reports and preliminary analysis**. In *Proceedings of the First Computing Social Responsibility Workshop within the 13th Language Resources and Evaluation Conference*, pages 24–34, Marseille, France. European Language Resources Association.

Minjin Choi, Sunkyung Lee, Eunseong Choi, Heesoo Park, Junhyuk Lee, Dongwon Lee, and Jongwuk Lee. 2021. Melbert: Metaphor detection via contextualized late interaction using metaphorical identification theories. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1763–1773.

Jacob Devasier, Yogesh Gurjar, and Chengkai Li. 2024. **Robust frame-semantic models with lexical unit trees and negative samples**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6930–6941, Bangkok, Thailand. Association for Computational Linguistics.

Ellen Dodge, Jisup Hong, and Elise Stickles. 2015. **MetaNet: Deep semantic automatic metaphor analysis**. In *Proceedings of the Third Workshop on Metaphor in NLP*, pages 40–49, Denver, Colorado. Association for Computational Linguistics.

Elisabeth El Refaie. 2001. Metaphors we discriminate by: Naturalized themes in austrian newspaper articles about asylum seekers. *Journal of Sociolinguistics*, 5(3):352–371.

Robert M Entman. 1993. Framing: Towards clarification of a fractured paradigm. *McQuail’s reader in mass communication theory*, 390:397.

743	Ethan Fast and Eric Horvitz. 2017. Long-term trends in the public perception of artificial intelligence. In <i>Proceedings of the AAAI conference on artificial intelligence</i> , volume 31.	799
744		800
745		801
746		
747	Charles J Fillmore and Collin F Baker. 2001. Frame semantics for text understanding. In <i>Proceedings of WordNet and Other Lexical Resources Workshop, NAACL</i> , volume 6, pages 59–64.	802
748		803
749		804
750		805
751	Bálint Forgács and Csaba Pléh. 2022. The fluffy metaphors of climate science. In <i>Metaphors and analogies in sciences and humanities: Words and worlds</i> , pages 447–477. Springer.	806
752		807
753		808
754		809
755	Maucha Gamonal. 2022. A descriptive study of metaphors and frames in the multilingual shared annotation task. In <i>Proceedings of the Workshop on Dimensions of Meaning: Distributional and Curated Semantics (DistCurate 2022)</i> , pages 1–7, Seattle, Washington. Association for Computational Linguistics.	810
756		811
757		812
758		813
759		814
760		
761	Mengshi Ge, Rui Mao, and Erik Cambria. 2022. Explainable metaphor identification inspired by conceptual metaphor theory. In <i>Proceedings of the AAAI conference on artificial intelligence</i> , volume 36, pages 10681–10689.	815
762		816
763		817
764		818
765		819
766	Mengshi Ge, Rui Mao, and Erik Cambria. 2023. A survey on computational metaphor processing techniques: From identification, interpretation, generation to application. <i>Artificial Intelligence Review</i> , 56(Suppl 2):1829–1895.	820
767		821
768		822
769		823
770		824
771		825
772	Jonathan Gordon, Jerry Hobbs, Jonathan May, Michael Mohler, Fabrizio Morbini, Bryan Rink, Marc Tomlinson, and Suzanne Wertheim. 2015. A corpus of rich metaphor annotation. In <i>Proceedings of the Third Workshop on Metaphor in NLP</i> , pages 56–66, Denver, Colorado. Association for Computational Linguistics.	826
773		827
774		828
775		829
776		830
777		831
778		832
779	Pragglejaz Group. 2007. Mip: A method for identifying metaphorically used words in discourse. <i>Metaphor and symbol</i> , 22(1):1–39.	833
780		834
781		835
782	Yolanda Guan and Winnie Huiheng Zeng. 2024. Changes in the sentiments and metaphors in COVID-19 news discourse (2019–2024). In <i>Proceedings of the 38th Pacific Asia Conference on Language, Information and Computation</i> , pages 810–819, Tokyo, Japan. Tokyo University of Foreign Studies.	836
783		837
784		838
785		839
786		
787	Carl G Herndl and Stuart C Brown. 1996. <i>Green culture: Environmental rhetoric in contemporary America</i> .	840
788		841
789	Jingyi Huang and Ming Liu. 2025. Metaphorical framing of climate change in chinese and american news media: A corpus-assisted discourse study. <i>Critical Arts</i> , pages 1–20.	842
790		843
791		844
792		845
793	Hyeju Jang, Keith Maki, Eduard Hovy, and Carolyn Rosé. 2017. Finding structure in figurative language: Metaphor detection with topic-based frames. In <i>Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue</i> , pages 320–330, Saarbrücken, Germany. Association for Computational Linguistics.	846
794		847
795		
796		
797		
798		
	Mark Jones. 1992. Generating a specific class of metaphors. In <i>30th Annual Meeting of the Association for Computational Linguistics</i> , pages 321–323.	799
		800
		801
	Rohan Joseph, Timothy Liu, Aik Beng Ng, Simon See, and Sunny Rai. 2023. NewsMet : A ‘do it all’ dataset of contemporary metaphors in news headlines. In <i>Findings of the Association for Computational Linguistics: ACL 2023</i> , pages 10090–10104, Toronto, Canada. Association for Computational Linguistics.	802
		803
		804
		805
		806
		807
	Oleksandr Kapranov. 2017. Conceptual metaphors associated with climate change in corporate reports in the fossil fuels market: Two perspectives from the united states and australia. In <i>The Role of Language in the Climate Change Debate</i> , pages 90–109. Routledge.	808
		809
		810
		811
		812
	George Lakoff and Mark Johnson. 2008. <i>Metaphors we live by</i> . University of Chicago press.	813
		814
	Mark J Landau, Daniel Sullivan, and Jeff Greenberg. 2009. Evidence that self-relevant motives and metaphoric framing interact to influence political and social attitudes. <i>Psychological science</i> , 20(11):1421–1427.	815
		816
		817
		818
		819
	Chee Wee (Ben) Leong, Beata Beigman Klebanov, and Ekaterina Shutova. 2018. A report on the 2018 VUA metaphor detection shared task. In <i>Proceedings of the Workshop on Figurative Language Processing</i> , pages 56–66, New Orleans, Louisiana. Association for Computational Linguistics.	820
		821
		822
		823
		824
		825
	Yanlin Li, Jing Chen, Kathleen Ahrens, and Chu-Ren Huang. 2024. The evolving use of WAR metaphors in businesswomen-focused media discourse. In <i>Proceedings of the 38th Pacific Asia Conference on Language, Information and Computation</i> , pages 1377–1386, Tokyo, Japan. Tokyo University of Foreign Studies.	826
		827
		828
		829
		830
		831
		832
	Yucheng Li, Shun Wang, Chenghua Lin, Frank Guerin, and Loic Barrault. 2023. FrameBERT: Conceptual metaphor detection with frame embedding learning. In <i>Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics</i> , pages 1558–1563, Dubrovnik, Croatia. Association for Computational Linguistics.	833
		834
		835
		836
		837
		838
		839
	Emmy Liu, Chenxuan Cui, Kenneth Zheng, and Graham Neubig. 2022. Testing the ability of language models to interpret figurative language. In <i>Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 4437–4452, Seattle, United States. Association for Computational Linguistics.	840
		841
		842
		843
		844
		845
		846
		847
	Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. <i>arXiv preprint arXiv:1907.11692</i> .	848
		849
		850
		851
		852
	Rui Mao, Xiao Li, Kai He, Mengshi Ge, and Erik Cambria. 2023. MetaPro online: A computational	853
		854

855	metaphor processing online system . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)</i> , pages 127–135, Toronto, Canada. Association for Computational Linguistics.	910
856		911
857		912
858		913
859		
860	Julia Mendelsohn. 2024. <i>Theory-grounded Computational Analysis of Political Framing in Online Media</i> . Phd dissertation, University of Michigan.	
861		
862		
863	Julia Mendelsohn and Ceren Budak. 2025. When people are floods: Analyzing dehumanizing metaphors in immigration discourse with large language models . In <i>Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 8079–8103, Vienna, Austria. Association for Computational Linguistics.	914
864		915
865		916
866		917
867		918
868		919
869		920
870	Julia Mendelsohn, Yulia Tsvetkov, and Dan Jurafsky. 2020. A framework for the computational linguistic analysis of dehumanization. <i>Frontiers in artificial intelligence</i> , 3:55.	
871		
872		
873		
874	Haohan Meng, Xiaoyu Li, and Jinhua Sun. 2025. Large language models prompt engineering as a method for embodied cognitive linguistic representation: a case study of political metaphors in trump’s discourse. <i>Frontiers in Psychology</i> , 16:1591408.	
875		
876		
877		
878		
879	Saif Mohammad, Ekaterina Shutova, and Peter Turney. 2016. Metaphor as a medium for emotion: An empirical study . In <i>Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics</i> , pages 23–33, Berlin, Germany. Association for Computational Linguistics.	
880		
881		
882		
883		
884		
885	Michael Mohler, Mary Brunson, Bryan Rink, and Marc Tomlinson. 2016. Introducing the lcc metaphor datasets. In <i>Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)</i> , pages 4221–4227.	
886		
887		
888		
889		
890	Brigitte Nerlich and Rusi Jaspal. 2012. Metaphors we die by? geoen지니어ing, metaphors, and the argument from catastrophe. <i>Metaphor and symbol</i> , 27(2):131–147.	
891		
892		
893		
894	Yulia Otmakhova, Shima Khanehzar, and Lea Frermann. 2024. Media framing: A typology and survey of computational approaches across disciplines . In <i>Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 15407–15428, Bangkok, Thailand. Association for Computational Linguistics.	
895		
896		
897		
898		
899		
900		
901	Vinodkumar Prabhakaran, Marek Rei, and Ekaterina Shutova. 2021. How metaphors impact political discourse: A large-scale topic-agnostic study using neural metaphor detection. In <i>Proceedings of the International AAAI Conference on Web and Social Media</i> , volume 15, pages 503–512.	
902		
903		
904		
905		
906		
907	Sunny Rai and Shampa Chakraverty. 2020. A survey on computational metaphor processing. <i>ACM Computing Surveys (CSUR)</i> , 53(2):1–37.	
908		
909		
	Paul Rayson and Roger Garside. 2000. Comparing corpora using frequency profiling . In <i>The Workshop on Comparing Corpora</i> , pages 1–6, Hong Kong, China. Association for Computational Linguistics.	910
		911
		912
		913
	Sebastian Reimann and Tatjana Scheffler. 2025. Using large language models to perform MIPVU-inspired automatic metaphor detection . In <i>Proceedings of the 2nd Workshop on Analogical Abstraction in Cognition, Perception, and Language (Analogy-Angle II)</i> , pages 10–21, Vienna, Austria. Association for Computational Linguistics.	914
		915
		916
		917
		918
		919
		920
	Zachary Rosen. 2018. Computationally constructed concepts: A machine learning approach to metaphor interpretation using usage-based construction grammatical cues . In <i>Proceedings of the Workshop on Figurative Language Processing</i> , pages 102–109, New Orleans, Louisiana. Association for Computational Linguistics.	921
		922
		923
		924
		925
		926
		927
	Josef Ruppenhofer, Michael Ellsworth, Miriam RL Petruck, and Christopher R Johnson. <i>Framenet ii: Extended theory and practice</i> .	928
		929
		930
	Elisa Sanchez-Bayona and Rodrigo Agerri. 2025. Metaphor and large language models: When surface features matter more than deep understanding . In <i>Findings of the Association for Computational Linguistics: ACL 2025</i> , pages 17462–17477, Vienna, Austria. Association for Computational Linguistics.	931
		932
		933
		934
		935
		936
	Bertram T Scheufele and Dietram A Scheufele. 2010. Of spreading activation, applicability, and schemas: Conceptual distinctions and their operational implications for measuring frames and framing effects. In <i>Doing News Framing Analysis</i> , pages 126–150. Routledge.	937
		938
		939
		940
		941
		942
	Elena Semino, Zsófia Demjén, and Jane Demmen. 2018. An integrated approach to metaphor and framing in cognition, discourse, and practice, with an application to metaphors for cancer. <i>Applied linguistics</i> , 39(5):625–645.	943
		944
		945
		946
		947
	Meghdut Sengupta, Milad Alshomary, and Henning Wachsmuth. 2022. Back to the roots: Predicting the source domain of metaphors using contrastive learning . In <i>Proceedings of the 3rd Workshop on Figurative Language Processing (FLP)</i> , pages 137–142, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.	948
		949
		950
		951
		952
		953
		954
	Meghdut Sengupta, Roxanne El Baff, Milad Alshomary, and Henning Wachsmuth. 2024. Analyzing the use of metaphors in news editorials for political framing . In <i>Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)</i> , pages 3621–3631, Mexico City, Mexico. Association for Computational Linguistics.	955
		956
		957
		958
		959
		960
		961
		962
		963
	Ekaterina Shutova and Simone Teufel. 2010. Metaphor corpus annotated for source-target domain mappings. In <i>LREC</i> , volume 2, pages 2–2.	964
		965
		966

1075 absence of improvement in terms of binary-F1. The
1076 starting learning rate is $2e^{-5}$ with linear learning
1077 rate scheduler and weight decay of 0.1.

1078 A.1.2 Source domain prediction

1079 **Data.** We use the large version (ENGLISH large)
1080 of the LCC Metaphor Dataset (large version)
1081 (Mohler et al., 2016). To ensure high quality of
1082 training data, we only use human-validated subsets
1083 of the dataset (ANN, REC, SYS), and only the sam-
1084 ples where at least half of the annotators regarding
1085 the source domain, or there is only one annotation.
1086 We also remove those source domains which have
1087 less than 3 samples, to ensure valid evaluation and
1088 testing. This results in 16722 samples, out of which
1089 11154 have single annotations, and 5568 are anno-
1090 tated by at least two people (with a high agreement
1091 in terms of Krippendorff’s α 0.76 and Fleiss κ of
1092 0.76).

1093 **Fine-tuning details.** All models have the same
1094 hyperparameters in terms of input length (256),
1095 batch size of 32, Adam W optimization, and 2
1096 warm-up epochs. All models are trained for the
1097 maximum of 30 epochs with patience of 5 as de-
1098 termined by absence of improvement in terms of
1099 macro-F1. The starting learning rate is $2e^{-5}$ with
1100 linear learning rate scheduler and weight decay of
1101 0.1.

1102 A.1.3 Metaphor prediction

1103 For fair comparison with Melbert, we reproduce
1104 fine-tuning parameters used in the original imple-
1105 mentation (Choi et al., 2021). In particular, we use
1106 the learning rate of $3e-05$ with 2 warm-up epochs
1107 and the drop ratio of 0.2, and train for 3 epochs.

1108 A.2 Classifier prompts

1109 A.2.1 Semantic frame classification

1110 We use the following prompt, where **labels** is sub-
1111 stituted with the list of semantic frames and their
1112 definitions:

You are a linguist specializing in
semantic frames (FrameNet).

Please choose the semantic frame out of
the following list: **labels**. Do not use
any other labels, and do not change the
wording of the label. Do not remove “_”
in the label if it exists.

Please return a json object which consists
of the following field:

“frame”: one of the values from the list.

Do not output anything else. Do not output
any reasoning steps or explanations.

1114 A.2.2 Source domain classification

1115 We use the following prompt, which is based on
1116 human annotation instructions (see Section A.8),
1117 where **labels** is substituted with the list of source
1118 domains:
1119

You are a linguist specializing in
metaphors. You will be given a metaphor
and a sentence it occurs in. You will be
asked to identify the source domain of
the metaphor.

A metaphor is when you describe something
by saying it is something else, even
though it is not literally true. For
example, we can say “They are forced to
make a decision” while there is no actual
physical force applied to them.

When using a metaphor, we are carrying
over associations from a more tangible and
specific source domain (such as physical
force or pressure) to a more abstract
domain (such as obligation).

In this study, you will need to identify
the source domain of the metaphor. It is
helpful to remember that the source domain
is usually a more specific, physical
thing (force, pressure) while the target
domain is more abstract (obligation).

Please choose the source domain out of
the following list: **labels**. Do not use
any other labels, and do not change the
wording of the label, Do not remove “_”
in the label if it exists.

Please return a json object which consists
of the following field:

“source”: one of the values from the list.

Do not output anything else. Do not output
any reasoning steps or explanations.

1120 A.2.3 Metaphor prediction

1121 We use the MIP protocol for the prompt, which our
1122 preliminary experiments have shown to be more ef-
1123 fective than other metaphor identification protocols
1124 such as SPV or CMT.
1125

You are a linguist specializing in
metaphors. You will be given a sentence
and asked to find all metaphors in it.
Please extract all text spans that have a

metaphorical meaning.

Go through the following steps to determine if a word is used in a metaphorical meaning:

1. Read the entire sentence to establish a general understanding of the meaning.
2. (a) Establish the word's meaning in context, that is, how it applies to an entity, relation, or attribute in the situation evoked by the text (contextual meaning). Take into account what comes before and after the lexical unit.
 - (b) Determine if the target word has a more basic contemporary meaning in other contexts than the one in the given context. For our purposes, basic meanings tend to be
 - More concrete; what they evoke is easier to imagine, see, hear, feel, smell, and taste.
 - Related to bodily action.
 - More precise (as opposed to vague)
 - Historically older.
 Basic meanings are not necessarily the most frequent meanings of the target word.
 - (c) If the target word has a more basic current-contemporary meaning in other contexts than the given context, decide whether the contextual meaning contrasts with the basic meaning but can be understood in comparison with it.
3. If yes, the word is metaphorical. Otherwise it is literal.

Please return a json object which consists of the following field:

metaphors: a list of extracted metaphor spans.

Do not output any explanations or reasoning steps.

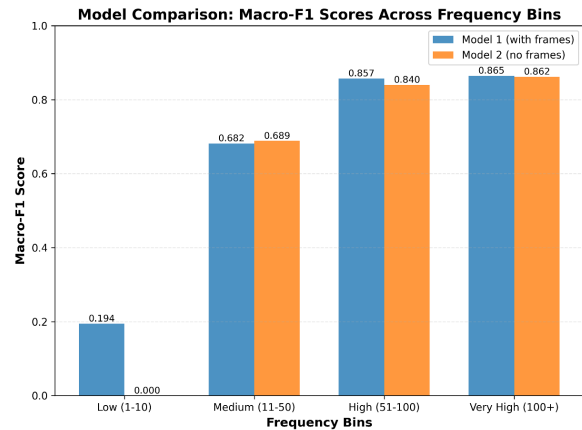


Figure 5: Performance comparison on test set for classes with different frequency in the training set.




Figure 6: Distribution of mentions of "climate change" and "global warming" across 1986-2020

To evaluate the extracted spans, we check if they overlap with golden spans in the dataset. We consider partial overlaps to be correct.

A.3 Source prediction performance depending on the class frequency

Figure 5 compared performance of SEP-style RoBERTa models with and without semantic frames on test set for classes with different frequency in the training set.

A.4 Distribution of mentions of "climate change" and "global warming" across 1986-2020

Figure 6 shows distribution of mentions of "climate change" and "global warming" in New York Times over the period from 1986 to 2020.

1143 A.5 Top source domains in climate corpus

1144 Table 6 shows the most salient source domains in
1145 NYT climate corpus, together with their examples,
1146 alternative names, and selected theoretical studies
1147 that identified them.

1148 A.6 Immigration source domains

1149 The corpus provided by (Mendelsohn and Budak,
1150 2025) contains tweets automatically annotated with
1151 predicted metaphoricity scores across seven dehu-
1152 manization source domains. The authors also re-
1153 lease a smaller dataset with 200 tweets per source
1154 domain, where each tweet is judged by 10 hu-
1155 man annotators as containing or not containing a
1156 metaphor from that particular domain.

1157 As the larger dataset only contains automatically
1158 predicted probabilities, we convert them to (poten-
1159 tial) source domain labels as follows. We regard a
1160 tweet t as potentially containing a metaphor with a
1161 particular source domain s if the predicted probabili-
1162 ty of that source domain $P(s|t)$ is larger than the
1163 average metaphoricity score $\bar{m}_s = \frac{1}{|T_s|} \sum_{t' \in T_s} m_{t'}$
1164 of tweets that were judged by annotators as belong-
1165 ing to that source domain (these cut-off scores are
1166 in Section A.6).

1167 Table 7 below shows the source domains used
1168 in metaphorical framing of tweets about immigra-
1169 tion, their definitions, the average metaphoricity
1170 scores from the annotated corpus by (Mendelsohn
1171 and Budak, 2025), the number of tweets from the
1172 unlabeled part of that corpus with metaphoricity
1173 scores about that cut off, as well as the final number
1174 of tweets used for analysis after filtering

1175 A.7 Filtering out metaphors which refer to 1176 non-relevant target domains

1177 To filter out metaphors which have the source do-
1178 main we are interested in but the wrong (irrele-
1179 vant) target, we first use Claude Sonnet 4.5 with
1180 the prompt below, where we ask it to resolve the
1181 target which the specific metaphor refers to:

```
1182 You are a linguist specializing in metaphors. You will be given a sentence and a metaphor in it. Identify and output the target of the metaphor, i.e. the entity, person, or object it refers to. For example, given the following input:
```

```
"text": "Immigrants are flooding into our country and ruining our economy"  
"metaphor": "flooding"
```

```
output the following:
```

```
"target": "immigrants"
```

```
Make sure to use json format with this field. Do not output anything else, no explanations! You are NOT allowed to say "no metaphor is present" since the metaphor is given to you!
```

1183 In this way, we collect a list of target do-
1184 mains used in tweets, such as "undocumented im-
1185 migrants", "illegals", "Liberals", "economy", or
1186 "Trump", and filter it out manually to collect a list
1187 of expressions that are likely to refer to the target
1188 domain we are interested in (immigration). In par-
1189 ticular, we arrive at the following list of substrings
1190 which then match against the predicted target do-
1191 mains: 'immi', 'illegal', 'alien', 'ICE', 'migr',
1192 'foreign', 'refug', 'detention', 'border', 'asylum',
1193 'famil'.

1194 A.8 Annotation Task Details

1195 A.8.1 Task Overview

1196 We collected human annotations for metaphor
1197 source domains using a custom web-based anno-
1198 tation platform.⁵ Annotators were presented with
1199 sentences containing highlighted metaphors and
1200 asked to identify up to 3 source domains from a
1201 provided list of options for each metaphor. The
1202 annotation interface included a mandatory 10 sec-
1203 ond reading period before options appeared, and
1204 copy-paste was not allowed, designed to encourage
1205 careful consideration of the metaphorical usage in
1206 context and to avoid AI-powered tools.

1207 A.8.2 Annotator Recruitment

1208 Annotators were recruited through Prolific (<https://www.prolific.com/>), a crowdsourcing plat-
1209 form commonly used for academic research. The
1210 study collected participants' Prolific IDs, study IDs,
1211 and session IDs to enable proper tracking and com-
1212 pensation. All participants were required to pro-
1213 vide informed consent through their Prolific ID
1214 before proceeding to the annotation task. Partici-
1215 pants were required to be native English speakers,
1216 with English listed as their first language. Eligible
1217 annotators reported their country of birth as either
1218 the United Kingdom or the United States. To en-
1219 sure a high level of reading comprehension and
1220 annotation quality, participants were additionally
1221 required to have completed at least an undergrad-
1222 uate degree (BA/BSc or equivalent), with many
1223

⁵Available at <https://metaphor-annotation-source.onrender.com/>

Source domain	Alternative names	Examples	Reference
STRUGGLE	CHALLENGE	fight, confront	Kapranov (2017)
WAR	BATTLE	battle, war, on the front lines	Skinmoen (2009)
OBJECT HANDLING	OBJECT	push (for), press, handle, take (action)	Herndl and Brown (1996)
HUMAN BODY	BODY	face, confront, embrace, at the heart of	Nerlich and Jaspal (2012)
PATHWAY	PATH	passing, path	Yang and Sun (2025)
BUILDING	CONSTRUCTION	build, lay a foundation	Huang and Liu (2025)
FURNISHINGS		climategate, take a back seat	Forgács and Pléh (2022)
VISION		view, presage	
FORWARD MOVEMENT		ahead, progress, gain	Skinmoen (2009)
BACKWARD MOVEMENT		backwards, rollback	Skinmoen (2009)

Table 6: Top 10 salience source domains in NYT climate corpus, together with their alternative names in the literature, examples and references. For brevity we provide only one reference per source domain.

Source domain	Examples	Avg metaphoricity score	# of candidates	# after filtering
Animal	shelter, cage, swarm	0.3045	9845	544
Water	flood, tide, pour	0.4375	3661	1402

Table 7: Statistics for ANIMAL and WATER source domains

holding graduate or doctoral degrees. Only Prolific users with an approval rate between 98% and 100% were permitted to participate.

A.8.3 Annotation Instructions

Figure 7 features the instructions provided to annotators during the interactive tutorial phase.

A.8.4 Quality Assurance

We implemented a multi-stage quality assurance protocol to ensure high annotation reliability. Prior to beginning the main annotation task, participants were shown five example annotations and then completed five comprehension (attention check) questions with known correct answers. To pass, participants were required to answer at least three out of five questions correctly (i.e., a maximum of two errors). Participants who failed the comprehension check were permitted one retry after reviewing the task instructions. A second failure resulted in permanent rejection and participants were instructed to return their submission on Prolific.

During the annotation task, participants were required to complete each sample sequentially before proceeding to the next. For each sample, annotators could select a maximum of three source domains. They could also select *No Metaphor*, which was mutually exclusive with all other options, or *Other Domain*, which prompted them to manually specify a domain not listed among the predefined options.

The annotation platform explicitly tracked three rejection scenarios:

1. **First comprehension check failure:** partic-

ipants were allowed to retry after reviewing the instructions.

2. **Second comprehension check failure:** participants were permanently rejected and instructed to return their submission on Prolific.

3. **Incomplete submissions:** participants who did not complete all 50 annotations in their assigned batch were automatically instructed to return their submission.

Participants were informed during the consent process that failing comprehension checks would result in submission return or rejection, and that repeated failures would lead to ineligibility for future tasks within the project unless manually re-approved.

A.8.5 Compensation

Each annotation batch was completed by four independent annotators. Participants were allotted up to 45 minutes to complete a batch of 50 annotations. Compensation was set at an hourly rate of £12, in line with Prolific’s recommended fair pay guidelines and well above the local standard minimum wage.

A.8.6 Annotation agreement per batch

Table 8 shows averaged pair-wise annotation agreement within each batch, as well as overall (mean) agreement across batches

Batch	Agreement rate
1	0.65
2	0.63
3	0.71
4	0.73
5	0.76
6	0.53
7	0.84
8	0.66
Mean	0.69

Table 8: Per batch IAA rate

Metaphor Annotation Study - Instructions

Please read through all instructions carefully before starting the study.

What is this study about?

A metaphor is when you describe something by saying it is something else, even though it's not literally true. For example, we can say "**They are forced to make a decision**" while there is no actual physical force applied to them.

When using a metaphor, we are carrying over associations from a more tangible and specific source domain (such as physical force or pressure) to a more abstract target domain (such as obligation).

In this study, you will need to identify the **source** domain of the metaphor.

It is helpful to remember that the source domain is usually a more specific, physical thing (force, pressure) that provides the conceptual structure for understanding the metaphor.

NEXT

Metaphor Annotation Study - Instructions

Please read through all instructions carefully before starting the study.

How does the task work?

You will see sentences with **highlighted words** that are used metaphorically, along with a list of 5 potential source domains. Your task is to select the source domain for the highlighted metaphor which represents associations and meaning carried over to the target domain.

You can select up to 3 source domains that apply to the metaphor.

For each sentence, you will first see the sentence to be annotated with source. The process is:

If you think that the highlighted word is not used as a metaphor in this sentence, select "NO METAPHOR"

- **Source Annotation:** Read the sentence carefully, focus on the highlighted metaphorical word, and select up to 3 appropriate source domains from the given options
- You can select "OTHER DOMAIN" along with other domains if needed
- If you select "OTHER DOMAIN", please fill in your suggestion in the provided text box

PREVIOUS

NEXT

Figure 7: Screenshots of the annotation instruction interface. Annotators were guided through an interactive tutorial explaining (top) the concept of metaphor source domains and (bottom) the annotation task workflow before beginning comprehension checks.