

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

Yulia Otmakhova Matteo Guida Lea Frermann

School of Computing and Information Systems,

The University of Melbourne

y.otmakhova@unimelb.edu.au,

guida@student.unimelb.edu.au, lea.frermann@unimelb.edu.au

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 WAR or THEORIES ARE BUILDINGS. Following their seminal work, *conceptual* metaphor theory (CMT) became a productive method to generalize from specific metaphori-

¹Code, data and statistical scripts are available at <https://github.com/julia-nixie/ConceptFrameMet>.

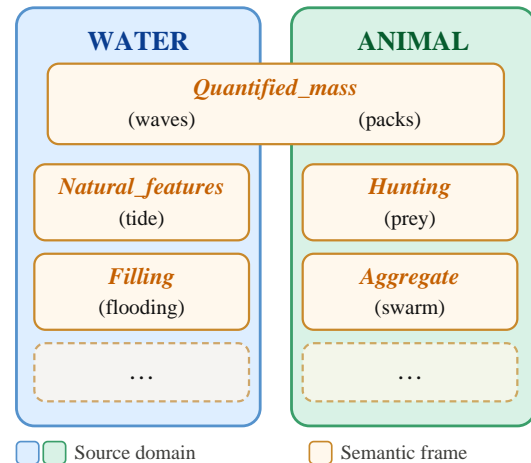


Figure 1: Interaction between semantic frames and domains

cal 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 “overflow-

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

ing” and can be considered a natural *feature of landscape*, thus instilling a more positive, accepting attitude towards immigration.

Such associations have been explained by linguistic, constructionalist theories of metaphors (Sullivan, 2025, 2013), positing that the associations derived from the source domain are related to the *semantic frame* of a lexical item used in the metaphor (Fillmore and Baker, 2001). In our example, *flood* has a semantic frame of *Filling* which emphasizes the movement of water and the negative result of such movement. On the other hand, *wave* and *tide* have semantic frames of *Quantified_mass* and *Natural_features* with more neutral associations. Semantic frames are generic structures that can be applied to many source domains. For example, the *Quantified_mass* semantic frame can also be applied to the ANIMAL source domain (*packs of immigrants*) or PRESSURE/PHYSICAL_BURDEN source domain (*loads of illegals*). Thus, it is the interplay of the source domain and semantic frame that uniquely defines the associations of the metaphor, where the source domain points to a cluster of associations, and the semantic frame allows to pick out specific ones (Figure 1).

This nuanced understanding of associations is particularly important when metaphors are being used as *framing devices* to emphasize and highlight a particular aspect of an issue or a debate, i.e. to frame it (Entman, 1993).³ Such “relatively stable metaphorical projections that function as key framing devices within a particular discourse over a certain period of time” which are termed *discourse metaphors* (Nerlich and Jaspal, 2012, p. 134) are the focus of our study. Specifically, we show that differences in metaphor use between political ideologies — which has not been fully explained in prior work (Mendelsohn, 2024; El Refaie, 2001) — can be effectively analyzed through prevalent choice of semantic frames that funnel particular associations from the source domain.

Our contributions are as follows:

- We are the first to employ FrameNet-style semantic frames to pinpoint the differences in associations carried through source domains,

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*.

and to analyse metaphorical framing at the discourse level rather than “local”, individual metaphors.

- To support the above, we implement a framework that consists of two components: a pre-trained language model that detects metaphors and predicts their semantic frame and source domain, and a statistical module that uses log-likelihood ratio estimation (Rayson and Garside, 2000) to highlight salient source domains and semantic frames of metaphors within a particular discourse.
- We apply our framework across two domains showing it can be used to *discover* prevalent discourse metaphors and fine-grained associations within a particular topic such as climate change, and to *highlight nuanced differences* in superficially similar metaphor usage, such as usage of ANIMAL, WATER and WAR metaphors within the immigration discourse by liberals and conservatives.
- We release a new corpus of climate change metaphors annotated with source domains, and re-open the discussion on the optimal taxonomy of source domains, proposing semantic frames as a principled way to derive it.

2 Related work

Metaphor analysis in NLP Metaphors have attracted significant attention, with an extensive body of work focusing on annotating metaphors and creating corpora (Group, 2007; Steen et al., 2010; Boisson et al., 2025), automatically detecting them in text (Wang et al., 2025; Zhang and Liu, 2022; Uduehi and Bunescu, 2024; Tian et al., 2024; Reimann and Scheffler, 2025), understanding (Tong et al., 2024; Ye et al., 2025), generation (Joseph et al., 2023; Chakrabarty et al., 2021; Veale, 2016), and interpretation, such as the ability to handle them in inference and question answering tasks (Liu et al., 2022; Sanchez-Bayona and Agerri, 2025) (see Ge et al. (2023); Rai and Chakraverty (2020) for an extensive survey).

Concept metaphor theory (CMT; Lakoff and Johnson (2008)) has been widely adopted as a tool to improve detection, explanation, and generation of metaphors (Jones, 1992; Stowe et al., 2021a; Ge et al., 2022; Mao et al., 2023; Tian et al., 2025; Wachowiak and Gromann, 2023). Some works use

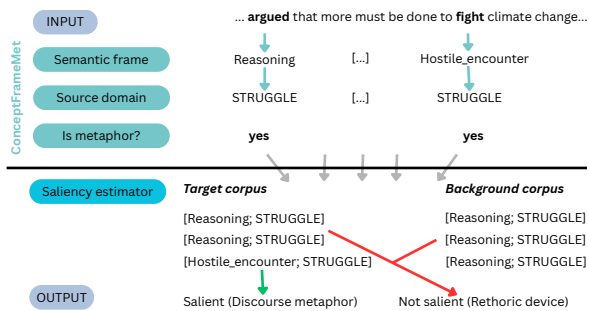


Figure 2: Framework overview with semantic frames and source domains for the input metaphors (**bold**).

semantic frames as a proxy 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 consist 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;

Ahrens and Zeng, 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 associations 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 target word in a sentence (details in Section A.1.1). We compare two input configurations: the target word masked out (*immigrants MASK the cities*), and the target word separated 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 – to choose one of 797 classes (Section A.2.1). We include results reported on the same test set for the same task in two recent papers (An et al., 2023;

	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).

Devasier et al., 2024), which, however, require substantial data-augmentation and are thus unsuitable 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, while LLMs significantly under-perform. The relatively low macro-F1 score is explained by a large number of minor classes which are semantically similar (i.e. *Capability vs Possibility*; Section A.3 lists the most often confused semantic frames). However, their overall proportion is low and thus unlikely to significantly affect the downstream analysis. We use the SEP-style model in all further experiments.

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. In Frames_CONCAT, we simply concatenate this feature vector with the RoBERTa-encoded input vector. In Frames_ATTEN, we maintain two copies of semantic frame vectors. The first is frozen, and the second is updated during training. During inference, we apply atten-

	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.

tion over the trainable matrix (with source domain embeddings as query) to highlight those semantic frames that are important for the source domain prediction, and add the frozen frame vectors as residuals. We also compare the 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 boosts source domain prediction performance, with the attention-based model performing best overall. In particular, they achieve an improvement of 20 points of macro F1 on under-represented classes (<10 instances in training data; see Appendix Figure 5). 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. To improve on that, 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*

	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 (metaphor 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.

the city and *immigrants flooded the city*, i.e. still be somewhat close to a metaphorical context embedding. Thus in ConceptFrameMet we replace the word embedding with that of the predicted source domain (WATER), amplifying the 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 (Section A.2.3). To ensure fair comparison with pre-trained models, we only evaluate the results on target words.⁴

Results. Results in Tables 3 and 7 show that MelBert performs best in-domain (i.e., when tested on the test portion of the same data set) but fails to generalize to other data sets. LLMs perform best on less varied, verb-only 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 (22%). However, the results on VUA-18 are indicative of practical usage scenarios where one needs to distinguish metaphors from mostly literal words. This,

⁴For positive cases (metaphor present), we consider an LLM prediction to be correct if the target word belongs to a extracted span; for negative cases (no metaphor), we confirm that the target word does not occur in any of extracted spans.

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

Since this study focuses on metaphors as a media framing device, we distinguish between *discourse metaphors* — which carry stable, prevalent associations used to frame some issue (Scheufele and Scheufele, 2010) — and metaphors that are merely figures of speech. Unfortunately, they cannot be differentiated based on their raw counts in text since some metaphors are used frequently across different contexts. Rather, a metaphor has to be salient in a given discourse (Zinken et al., 2008), i.e. occur in it more frequently than in other discourses (see lower half of Figure 2). Unlike Mason (2004) who compare raw metaphor frequencies between corpora, 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 back ground corpus (Section 4.1); and *contrast* discourse metaphors in texts associated with polar opposite political leanings within a single topic (Section 4.2).

	Acc	P	R	F1
Metaphor prediction	0.94	0.94	1	0.97
Source domain prediction	0.88	0.62	0.66	0.64

Table 4: Performance of ConceptFrameMet on the climate dataset. Scores are binary for metaphor detection, and macro-averaged for source domain prediction.

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 analyze them through salient semantic frames.

4.1.1 Data and methods

We collect a dataset of 47K paragraphs about climate change from a corpus of New York Times articles published between 1986 and 2020 (Fast and Horvitz, 2017; Mendelsohn et al., 2020) and an equivalent generic corpus randomly sampled from the same source (details in Section A.6). 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 416 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.9 for more details). We collect four

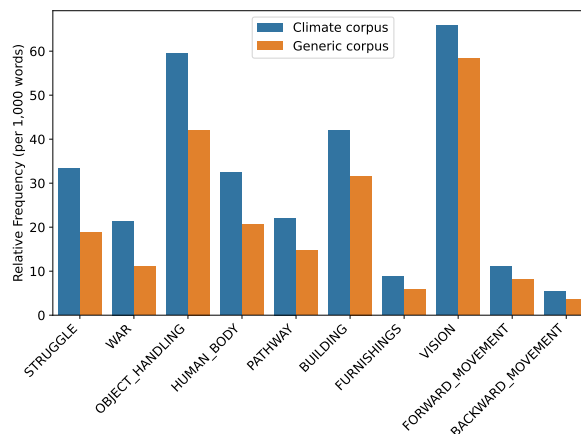


Figure 3: Relative frequency of source domains in climate and generic corpora. The source domains are sorted left to right in order of diminishing saliency

annotators per sample and achieve a reliable averaged pair-wise annotator agreement of 68% , 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, showing that ConceptFrameMet generalizes well to our new domain (Table 4).⁵

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.7.

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

⁵Note that the metaphor prediction recall and F1 are 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

We use a corpus of 400K US tweets on immigration with automatically predicted metaphoricity scores and scores of political affiliation (liberal vs conservative; Mendelsohn and Budak (2025)). We select three well-studied domains (WATER, ANIMAL and WAR) for analysis.

We extract and filter out candidate tweets as explained in Section A.8, and use ConceptFrameMet to annotate them for metaphors, their semantic frames and source domains (the resulting statistics in Table 9). We apply log-likelihood ratio to semantic frames within the ANIMAL, WATER and WAR 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.

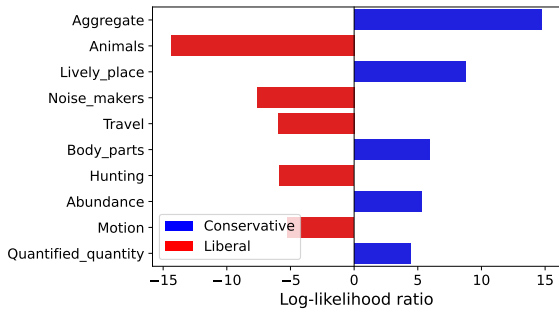


Figure 4: The five semantic frames with highest saliency in the liberal (red) and conservative (blue) data, in the ANIMAL domain.

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 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* (“IL-LEGALS 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*).

Within the WATER source domain, the conservative side features only a few 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 **trickle** of immigrants”), *Natural feature* (“Many different immigrant **streams**”), 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”). Similarly, the WAR metaphors on the conservative side demonstrate semantic paucity (despite their frequent usage), with only use 2 semantic frames of *Invading* and *Arriving* (“invaders”, “invasion” etc). On the other hand, liberals have a broader variety of semantic frames highlighting their persecution by authorities such as *Hostile_encounter* (“waging war”), *Attack* (“raid”, “assault”, “onslaught”), *Revenge* (“crack-down”, “retaliate”) etc).

5 Discussion

We highlight several open-ended questions regarding metaphor analysis that fall out of our work.

Can semantic frames provide a more principled way of choosing source domains? To date, there exists no principled way of constructing the ontology of conceptual mappings between source and target domains. The Master Metaphor List (Lakoff et al., 1991) has been criticized for conflating concepts across different taxonomy levels, resulting in categories that are neither consistent in granularity nor mutually exclusive (Lönneker-Rodman, 2008; Shutova and Teufel, 2010). We also observe this in the LCC corpus (Mohler et al., 2016), and estimate which source domains are likely to be problematic using our annotated corpus (see Section A.9). We calculate normalized point-wise mutual information (NPMI) of pairs of labels that are chosen for the same metaphor by a single annotator, or by all annotators. To account for noise, we weight the resulting score by the frequency of a particular label pair (Section A.10). This highlights such source domains as BATTLE vs WAR, BARRIER vs PATHWAY, and LIGHT vs VISION as being the hardest to distinguish (top 10 confusions in Tables 11 and 12).

The distribution of semantic frames that support these competing source domains, however, reveals two distinct patterns. Some of the source domains such as BATTLE are completely subsumed by another domain (WAR): all four semantic frames of BATTLE are also part of the WAR source domain

(65 semantic frames). Thus, BATTLE is a child of the WAR source domain, and can be considered redundant. On the other hand, while some source domains are related in meaning (BARRIER vs PATHWAY), the bulk of semantic frames constituting each of them is distinct: only 30% of PATHWAY semantic frames overlap with BARRIER, and 45% in the reverse direction. Here, the composition of source domains could help to choose a more appropriate source domain where multiple mappings are possible. For example, the metaphor *roadblocks (to net zero)* is closer to such semantic frames as *Hindering, Prohibiting, Thwarting* that belong to BARRIER domain, rather than to *Means, Manner, Trajectory* that characterize PATHWAYS domain. Thus, semantic frames can help to establish the relationship between the source domains and provide a more systematic way of distinguishing between them.

Can the semantic analysis of source domains be reproduced by LLMs end-to-end? While the results above demonstrate sub-par LLM performance on the individual tasks in our pipeline (metaphor detection, source domain prediction, semantic frame identification), this does not rule out the possibility that LLMs are able to perform the full task end-to-end and identify the most salient metaphors within the same instruction workflow. To test this, we conducted an experiment on two subsets of the immigration tweet corpus, namely those with source domain WATER or ANIMAL. Thus, we simplify the task, only asking the model to identify salient semantic frames.

We used Claude Sonnet 4.0 with instructions that combined all steps of the pipeline (detect metaphors, identify semantic frames, decide which semantic frames are salient for the conservative vs liberal side). We provided the list of semantic frames, and supplied the sets of liberal and conservatives tweets for a particular domain, clearly marking them as such (prompt in Section A.11).

We found the results lacking in multiple respects. Some semantic frames (e.g. “dehumanization” for the ANIMAL domain, Section A.12.1) and the examples to support them were hallucinated. The LLM also attributed the same lexical item (for example, “flood”) to multiple semantic frames, or the same semantic frames (“Catastrophe”) to both liberals and conservatives (Section A.12.2), unable to compare and weigh the semantic frames in a principled way. Most importantly, we suspect that

LLM predictions are strongly affected by their pre-training knowledge, and not primarily driven by the input corpus. To test this, we presented the LLM with tweets that were unlikely to include WATER metaphors but prompted it to identify WATER metaphors (Section A.12.3). The model still generated semantic frames, and they were nearly identical to those generated when given context, though the examples generated by the model either lacked WATER metaphors or were not found in the provided context (Section A.12.2). This suggests that the model has apriori “knowledge” about metaphors and retrofits the examples to it, either picking them from the texts (if they are available) or hallucinating them (when they are not). For large-scale analyses it is unclear how to check for hallucination and/or outputs that are not grounded in the discourse of interest, how to rank the source domains and semantic frames in a reproducible way, and how to objectively judge the results considering that the model also adds its interpretation which, while superficially plausible, is detrimental to a principled analysis. In contrast, the advantages of our approach are interpretability, grounding in existing resources, and the ability to assess the statistical significance of the findings. Our work also provides an important reference for future work to benchmark and improve LLMs to make their analysis more grounded and less interpretative.

6 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.

7 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 image-schematic approaches (Wachowiak et al., 2022) or the MetaNet (Dodge et al., 2015) framework, which encode the relationships between the components of the source domain, to enable more sophisticated analysis of semantic frames that support complex, extended discourse metaphors.

Acknowledgments

This work was supported by the Australian Research Council Discovery Early Career Research Award (Grant No. DE230100761) and University of Melbourne’s Early Career Researcher Grant (Grant No. 2025ECRG137).

References

- Kathleen Ahrens and Winnie Huiheng Zeng. 2022. Referential and evaluative strategies of conceptual metaphor use in government discourse. *Journal of Pragmatics*, 188:83–96.
- 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 nonliterary 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.
- Ethan Fast and Eric Horvitz. 2017. Long-term trends in the public perception of artificial intelligence. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- Charles J Fillmore and Collin F Baker. 2001. Frame semantics for text understanding. In *Proceedings of WordNet and Other Lexical Resources Workshop, NAACL*, volume 6, pages 59–64.
- Bálint Forgács and Csaba Pléh. 2022. The fluffy metaphors of climate science. In *Metaphors and analogies in sciences and humanities: Words and worlds*, pages 447–477. Springer.
- Maucha Gamonal. 2022. [A descriptive study of metaphors and frames in the multilingual shared annotation task](#). In *Proceedings of the Workshop on Dimensions of Meaning: Distributional and Curated Semantics (DistCurate 2022)*, pages 1–7, Seattle, Washington. Association for Computational Linguistics.
- Mengshi Ge, Rui Mao, and Erik Cambria. 2022. Explainable metaphor identification inspired by Conceptual Metaphor Theory. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 10681–10689.
- Mengshi Ge, Rui Mao, and Erik Cambria. 2023. A survey on computational metaphor processing techniques: From identification, interpretation, generation to application. *Artificial Intelligence Review*, 56(Suppl 2):1829–1895.
- 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 *Proceedings of the Third Workshop on Metaphor in NLP*, pages 56–66, Denver, Colorado. Association for Computational Linguistics.
- Pragglejaz Group. 2007. MIP: A method for identifying metaphorically used words in discourse. *Metaphor and symbol*, 22(1):1–39.
- Yolanda Guan and Winnie Huiheng Zeng. 2024. [Changes in the sentiments and metaphors in COVID-19 news discourse \(2019-2024\)](#). In *Proceedings of the 38th Pacific Asia Conference on Language, Information and Computation*, pages 810–819, Tokyo, Japan. Tokyo University of Foreign Studies.
- Carl G Herndl and Stuart C Brown. 1996. *Green culture: Environmental rhetoric in contemporary America*.
- Jingyi Huang and Ming Liu. 2025. Metaphorical framing of climate change in Chinese and American news media: A corpus-assisted discourse study. *Critical Arts*, pages 1–20.
- Hyeju Jang, Keith Maki, Eduard Hovy, and Carolyn Rosé. 2017. [Finding structure in figurative language: Metaphor detection with topic-based frames](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 320–330, Saarbrücken, Germany. Association for Computational Linguistics.
- Mark Jones. 1992. Generating a specific class of metaphors. In *30th Annual Meeting of the Association for Computational Linguistics*, pages 321–323.
- 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 *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10090–10104, Toronto, Canada. Association for Computational Linguistics.
- 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 *The Role of Language in the Climate Change Debate*, pages 90–109. Routledge.
- George Lakoff, Jane Espenson, and Alan Schwartz. 1991. Master metaphor list (technical report). *Cognitive Linguistics Group University of California, Berkeley*.
- George Lakoff and Mark Johnson. 2008. *Metaphors we live by*. University of Chicago press.
- Mark J Landau, Daniel Sullivan, and Jeff Greenberg. 2009. Evidence that self-relevant motives and metaphoric framing interact to influence political and social attitudes. *Psychological science*, 20(11):1421–1427.
- Chee Wee (Ben) Leong, Beata Beigman Klebanov, and Ekaterina Shutova. 2018. [A report on the 2018 VUA metaphor detection shared task](#). In *Proceedings of the Workshop on Figurative Language Processing*, pages 56–66, New Orleans, Louisiana. Association for Computational Linguistics.

- Yanlin Li, Jing Chen, Kathleen Ahrens, and Chu-Ren Huang. 2024. [The evolving use of WAR metaphors in businesswomen-focused media discourse](#). In *Proceedings of the 38th Pacific Asia Conference on Language, Information and Computation*, pages 1377–1386, Tokyo, Japan. Tokyo University of Foreign Studies.
- Yucheng Li, Shun Wang, Chenghua Lin, Frank Guerin, and Loic Barrault. 2023. [FrameBERT: Conceptual metaphor detection with frame embedding learning](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1558–1563, Dubrovnik, Croatia. Association for Computational Linguistics.
- Emmy Liu, Chenxuan Cui, Kenneth Zheng, and Graham Neubig. 2022. [Testing the ability of language models to interpret figurative language](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4437–4452, Seattle, United States. Association for Computational Linguistics.
- 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. *arXiv preprint arXiv:1907.11692*.
- Birte Lönneker-Rodman. 2008. The hamburg metaphor database project: issues in resource creation. *Language resources and evaluation*, 42(3):293–318.
- Rui Mao, Xiao Li, Kai He, Mengshi Ge, and Erik Cambria. 2023. [MetaPro online: A computational metaphor processing online system](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 127–135, Toronto, Canada. Association for Computational Linguistics.
- Zachary J Mason. 2004. CorMet: A computational, corpus-based conventional metaphor extraction system. *Computational linguistics*, 30(1):23–44.
- Julia Mendelsohn. 2024. *Theory-grounded Computational Analysis of Political Framing in Online Media*. PhD dissertation, University of Michigan.
- Julia Mendelsohn and Ceren Budak. 2025. [When people are floods: Analyzing dehumanizing metaphors in immigration discourse with large language models](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8079–8103, Vienna, Austria. Association for Computational Linguistics.
- Julia Mendelsohn, Yulia Tsvetkov, and Dan Jurafsky. 2020. A framework for the computational linguistic analysis of dehumanization. *Frontiers in artificial intelligence*, 3:55.
- 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. *Frontiers in Psychology*, 16:1591408.
- Saif Mohammad, Ekaterina Shutova, and Peter Turney. 2016. [Metaphor as a medium for emotion: An empirical study](#). In *Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics*, pages 23–33, Berlin, Germany. Association for Computational Linguistics.
- Michael Mohler, Mary Brunson, Bryan Rink, and Marc Tomlinson. 2016. Introducing the LCC metaphor datasets. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 4221–4227.
- Brigitte Nerlich and Rusi Jaspal. 2012. Metaphors we die by? Geoengineering, metaphors, and the argument from catastrophe. *Metaphor and symbol*, 27(2):131–147.
- Yulia Otmakhova, Shima Khanehzar, and Lea Frermann. 2024. [Media framing: A typology and survey of computational approaches across disciplines](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15407–15428, Bangkok, Thailand. Association for Computational Linguistics.
- Vinodkumar Prabhakaran, Marek Rei, and Ekaterina Shutova. 2021. How metaphors impact political discourse: A large-scale topic-agnostic study using neural metaphor detection. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 15, pages 503–512.
- Sunny Rai and Shampa Chakraverty. 2020. A survey on computational metaphor processing. *ACM Computing Surveys (CSUR)*, 53(2):1–37.
- Paul Rayson and Roger Garside. 2000. [Comparing corpora using frequency profiling](#). In *The Workshop on Comparing Corpora*, pages 1–6, Hong Kong, China. Association for Computational Linguistics.
- Sebastian Reimann and Tatjana Scheffler. 2025. [Using large language models to perform MIPVU-inspired automatic metaphor detection](#). In *Proceedings of the 2nd Workshop on Analogical Abstraction in Cognition, Perception, and Language (Analogy-Angle II)*, pages 10–21, Vienna, Austria. Association for Computational Linguistics.
- Zachary Rosen. 2018. [Computationally constructed concepts: A machine learning approach to metaphor interpretation using usage-based construction grammatical cues](#). In *Proceedings of the Workshop on Figurative Language Processing*, pages 102–109, New Orleans, Louisiana. Association for Computational Linguistics.
- Josef Ruppenhofer, Michael Ellsworth, Miriam RL Petruck, and Christopher R Johnson. FrameNet II: Extended theory and practice.

- Elisa Sanchez-Bayona and Rodrigo Agerri. 2025. [Metaphor and large language models: When surface features matter more than deep understanding](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 17462–17477, Vienna, Austria. Association for Computational Linguistics.
- 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 *Doing News Framing Analysis*, pages 126–150. Routledge.
- 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. *Applied linguistics*, 39(5):625–645.
- Meghdut Sengupta, Milad Alshomary, and Henning Wachsmuth. 2022. [Back to the roots: Predicting the source domain of metaphors using contrastive learning](#). In *Proceedings of the 3rd Workshop on Figurative Language Processing (FLP)*, pages 137–142, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Meghdut Sengupta, Roxanne El Baff, Milad Alshomary, and Henning Wachsmuth. 2024. [Analyzing the use of metaphors in news editorials for political framing](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3621–3631, Mexico City, Mexico. Association for Computational Linguistics.
- Ekaterina Shutova and Simone Teufel. 2010. Metaphor corpus annotated for source-target domain mappings. In *LREC*, volume 2, pages 2–2.
- Jorunn Skinnemoen. 2009. Metaphors in climate change discourse. Master’s thesis.
- Gerard J Steen, Aletta G Dorst, Tina Krennmayr, Anna A Kaal, and J Berenike Herrmann. 2010. A method for linguistic metaphor identification.
- Kevin Stowe, Nils Beck, and Iryna Gurevych. 2021a. [Exploring metaphoric paraphrase generation](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 323–336, Online. Association for Computational Linguistics.
- Kevin Stowe, Tuhin Chakrabarty, Nanyun Peng, Smaranda Muresan, and Iryna Gurevych. 2021b. [Metaphor generation with conceptual mappings](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6724–6736, Online. Association for Computational Linguistics.
- Karen Sullivan. 2013. *Frames and constructions in metaphoric language*. John Benjamins Publishing Company.
- Karen Sullivan. 2023. Three levels of framing. *Wiley Interdisciplinary Reviews: Cognitive Science*, 14(5):e1651.
- Karen Sullivan. 2025. *Metaphors and Constructions*, page 129–146. Cambridge Handbooks in Language and Linguistics. Cambridge University Press.
- Swabha Swayamdipta, Sam Thomson, Chris Dyer, and Noah A Smith. 2017. Frame-semantic parsing with softmax-margin segmental RNNs and a syntactic scaffold. *arXiv preprint arXiv:1706.09528*.
- Yuan Tian, Minzheng Wang, Nan Xu, and Wenji Mao. 2025. [ImaRA: An imaginative frame augmented method for low-resource multimodal metaphor detection and explanation](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 3953–3967, Albuquerque, New Mexico. Association for Computational Linguistics.
- Yuan Tian, Nan Xu, and Wenji Mao. 2024. [A theory guided scaffolding instruction framework for LLM-enabled metaphor reasoning](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7738–7755, Mexico City, Mexico. Association for Computational Linguistics.
- Xiaoyu Tong, Rochelle Choenni, Martha Lewis, and Ekaterina Shutova. 2024. [Metaphor understanding challenge dataset for LLMs](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3517–3536, Bangkok, Thailand. Association for Computational Linguistics.
- Oseremen Uduehi and Razvan Bunescu. 2024. [An expectation-realization model for metaphor detection](#). In *Proceedings of the 4th Workshop on Figurative Language Processing (FigLang 2024)*, pages 79–84, Mexico City, Mexico (Hybrid). Association for Computational Linguistics.
- Gisela Vallejo, Christine de Kock, Timothy Baldwin, and Lea Frermann. 2025. [Human interest framing across cultures: A case study on climate change](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 11380–11398, Abu Dhabi, UAE. Association for Computational Linguistics.
- Tony Veale. 2016. [Round up the usual suspects: Knowledge-based metaphor generation](#). In *Proceedings of the Fourth Workshop on Metaphor in NLP*, pages 34–41, San Diego, California. Association for Computational Linguistics.
- Lennart Wachowiak and Dagmar Gromann. 2023. [Does GPT-3 grasp metaphors? Identifying metaphor mappings with generative language models](#). In *Proceedings of the 61st Annual Meeting of the Association for*

Computational Linguistics (Volume 1: Long Papers), pages 1018–1032, Toronto, Canada. Association for Computational Linguistics.

Lennart Wachowiak, Dagmar Gromann, and Chao Xu. 2022. [Drum up SUPPORT: Systematic analysis of image-schematic conceptual metaphors](#). In *Proceedings of the 3rd Workshop on Figurative Language Processing (FLP)*, pages 44–53, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Dian Wang, Yang Li, Suge Wang, Xin Chen, Jian Liao, Deyu Li, and Xiaoli Li. 2025. CKEMI: Concept knowledge enhanced metaphor identification framework. *Information Processing & Management*, 62(1):103946.

Yunxiao Wang. 2024. [Metaphorical framing of refugees, asylum seekers and immigrants in UK’s left and right-wing media](#). In *Proceedings of the 8th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature (LaTeCH-CLfL 2024)*, pages 18–27, St. Julians, Malta. Association for Computational Linguistics.

Nicolas Webersinke, Mathias Kraus, Julia Anna Binger, and Markus Leippold. 2021. ClimateBERT: A pretrained language model for climate-related text. *arXiv preprint arXiv:2110.12010*.

Kaiwen Yang and Ya Sun. 2025. A tale of identities: Environmental identities based on a deliberate metaphor analysis of US energy companies’ social media. *Journal of Business and Technical Communication*, 39(2):149–190.

Fengying Ye, Shanshan Wang, Lidia S Chao, and Derek F Wong. 2025. Unveiling LLMs’ metaphorical understanding: Exploring conceptual irrelevance, context leveraging and syntactic influence. *arXiv preprint arXiv:2510.04120*.

Shenglong Zhang and Ying Liu. 2022. [Metaphor detection via linguistics enhanced Siamese network](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 4149–4159, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Jörg Zinken, Iina Hellsten, and Brigitte Nerlich. 2008. Discourse metaphors. *Body, language and mind*, 2:363–385.

A Appendix

A.1 Pre-training models finetuning details

A.1.1 Semantic frame prediction

Fine-tuning. We use input length of 256, batch size of 32, Adam W optimization, and 2 warm-up epochs. All models are trained for the maximum of 10 epochs with patience of 2 as determined by

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

A.1.2 Source domain prediction

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

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

A.1.3 Metaphor prediction

Data. VUA-18: we use the original splits of 6323 (training)/1550 (development)/2694 (test). We use all samples (metaphors across all parts-of-speech), the average metaphor rate being 22%. We use the training split for fine-tuning, and report the results on the test set.

TroFi: we use the whole corpus (3737 sentences) for evaluation (no fine-tuning, in zero-shot way). The corpus focuses on verbs and has a metaphor rate of around 43%.

MOH-X: we use the whole corpus (647 sentences) for evaluation for evaluation (no fine-tuning, in zero-shot way). The corpus focuses on verbs and has a metaphor rate of around 49%.

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

A.2 Classifier prompts

A.2.1 Semantic frame classification

We use the following prompt, where **labels** is substituted with the list of semantic frames and their 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.

A.2.2 Source domain classification

We use the following prompt, which is based on human annotation instructions (see Section A.9), where **labels** is substituted with the list of source domains:

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.

A.2.3 Metaphor prediction

We use the MIP protocol for the prompt, which our preliminary experiments have shown to be more effective than other metaphor identification protocols such as SPV or CMT.

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.

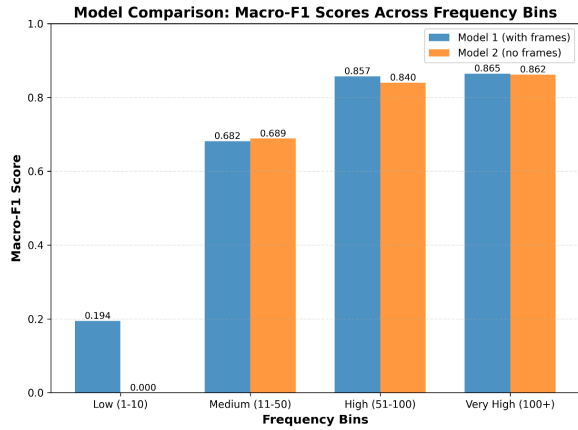


Figure 5: Performance comparison on test set for classes with different frequency in the training set.

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 Most often confused semantic frames

The table below shows the pairs of most often confused semantic frames by RoBERTA SEQ, as well as their raw counts in the test set (out of 6714 samples).

Sem. frame 1	Sem. frame 2	Count
Buildings	Locale_by_use	23
Locating	Becoming_aware	14
Possibility	Capability	13
Calendric_unit	Measure_duration	12
Ride_vehicle	Bringing	11
Stage_of_progress	Temporal_collocation	10
Posture	Change_posture	10
Have_associated	Possession	9
Capability	Possibility	8
Touring	Travel	8

Table 6: Frame pairs and counts

A.4 Full results for metaphor classifier performance

A.5 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.6 Collecting the climate change corpus and generic corpus

We use a corpus of New York Times articles published between 1986 and 2020 and extract all paragraphs which mention “climate change” or “global

warming”. Next, we filter them using a ClimateBERT pre-trained model (Webersinke et al., 2021) which was fine-tuned to detect climate-related texts (Binger et al., 2023), and, additionally, use the prompt from Vallejo et al. (2025) with Gemini 2.5 Flash to detect if a paragraph focuses on climate change (rather than just contains the keywords).

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

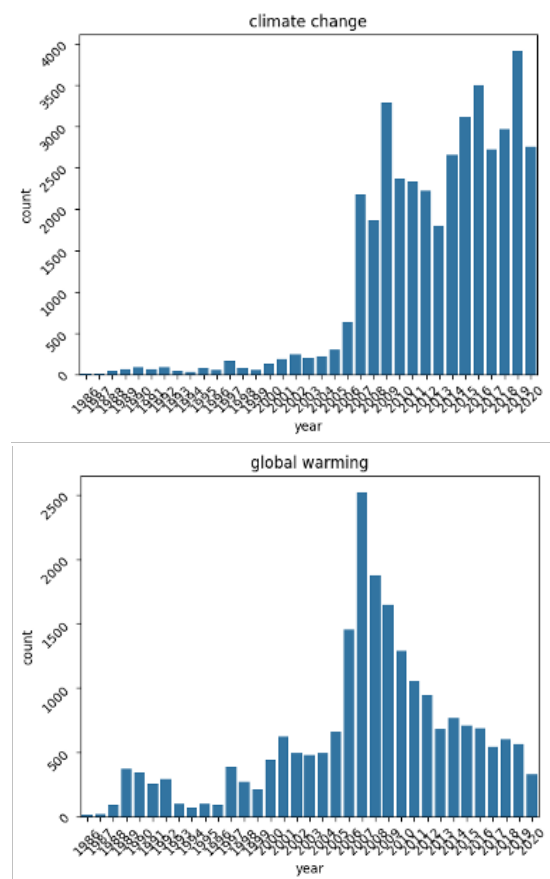


Figure 6: Distribution of mentions of “climate change” or “global warming” across 1986-2020

For saliency analysis, we also collect a generic corpus randomly sampled from the same NYT data set ensuring that the paragraphs do not contain our climate keywords. We sample the same number of paragraphs per year as in the climate corpus (47K in total).

A.7 Top source domains in climate corpus

Table 8 shows the most salient source domains in NYT climate corpus, together with their examples, alternative names, and selected theoretical studies that identified them.

	VUA-18				TroFi				MOH-X			
	Acc	P	R	F1	Acc	P	R	F1	Acc	P	R	F1
Random baseline	0.499	0.142	0.501	0.222	0.500	0.436	0.501	0.466	0.484	0.468	0.505	0.486
MelBert	0.938	0.778	0.786	0.782	0.593	0.521	0.801	0.631	0.800	0.753	0.871	0.806
FrameBert	0.933	0.757	0.776	0.766	0.595	0.524	0.761	0.620	0.769	0.720	0.859	0.780
ConceptFrameMet	0.931	0.732	0.806	0.767	0.605	0.532	0.787	0.634	0.803	0.746	0.899	0.814
Gemini 2.5	0.786	0.304	0.390	0.341	0.601	0.528	0.789	0.633	0.847	0.809	0.893	0.849
Claude Sonnet 4.0	0.827	0.367	0.301	0.330	0.602	0.526	0.852	0.651	0.806	0.739	0.926	0.822

Table 7: Full results for metaphor prediction task for pre-trained models fine-tuned on VUA-18 dataset. The reported metrics are binary (for the metaphor class).

Source domain	Alternative names	Examples	Reference
STRUGGLE	CHALLENGE	fight, confront	Kapranov (2017)
WAR	BATTLE	battle, war, on the front lines	Skinemoen (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	Skinemoen (2009)
BACKWARD MOVEMENT		backwards, rollback	Skinemoen (2009)

Table 8: 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.

A.8 Immigration source domains

The corpus provided by Mendelsohn and Budak (2025) contains tweets automatically annotated with predicted metaphoricity scores across seven dehumanization source domains. The authors also release a smaller dataset with 200 tweets per source domain, where each tweet is judged by 10 human annotators as containing or not containing a metaphor from that particular domain.

As the larger dataset only contains automatically predicted probabilities, we convert them to (potential) source domain labels as follows. We regard a tweet t as potentially containing a metaphor with a particular source domain s if the predicted probability of that source domain $P(s|t)$ is larger than the average metaphoricity score $\bar{m}_s = \frac{1}{|T_s|} \sum_{t' \in T_s} m_{t'}$ of tweets that were judged by annotators as belonging to that source domain (these cut-off scores are in Section A.8).

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, ANIMAL or WAR source can refer to another, irrelevant target rather than immigration (for example, as in *left-winged politicians*), we further filter out tweets where the metaphor does not refer to immigrants.

To filter out metaphors which have the source domain we are interested in but the wrong (irrelevant) target, we first use Claude Sonnet 4.5 with the prompt below, where we ask it to resolve the target which the specific metaphor refers to:

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!

In this way, we collect a list of target domains used in tweets, such as "undocumented immigrant", "illegals", "Liberals", "economy", or "Trump", and filter it out manually to collect a list of expressions that are likely to refer to the target domain we are interested in (immigration). In particular, we arrive

at the following list of substrings which then match against the predicted target domains: “immi”, “illegal”, “alien”, “ICE”, “migr”, “foreign”, “refug”, “detention”, “border”, “asylum”, “famil”.

Table 9 below shows the source domains used in metaphorical framing of tweets about immigration, their definitions, the average metaphoricity scores from the annotated corpus by (Mendelsohn and Budak, 2025), the number of tweets from the unlabeled part of that corpus with metaphoricity scores about that cut off, as well as the final number of tweets used for analysis after filtering.

A.9 Annotation task details

A.9.1 Task overview

We collected human annotations for metaphor source domains using a custom web-based annotation platform.⁶ Annotators were presented with sentences containing highlighted metaphors and asked to identify up to 3 source domains from a provided list of options for each metaphor. The annotation interface included a mandatory 10 second reading period before options appeared, and copy-paste was not allowed, designed to encourage careful consideration of the metaphorical usage in context and to avoid AI-powered tools. The order of samples within the batch was randomized to avoid batch effects.

A.9.2 Annotator recruitment

Annotators were recruited through Prolific (<https://www.prolific.com/>), a crowdsourcing platform commonly used for academic research. The study collected participants’ Prolific IDs, study IDs, and session IDs to enable proper tracking and compensation. All participants were required to provide informed consent through their Prolific ID before proceeding to the annotation task. However, all participants data (Prolific IDs) was anonymised after collection and quality checks. Participants were required to be native English speakers, with English listed as their first language. Eligible annotators reported their country of birth as either the United Kingdom or the United States. To ensure a high level of reading comprehension and annotation quality, participants were additionally required to have completed at least an undergraduate degree (BA/BSc or equivalent), with many holding graduate or doctoral degrees. Only Prolific users

⁶Available at <https://metaphor-annotation-source.onrender.com/>

with an approval rate between 98% and 100% were permitted to participate.

This study was approved by the Human Ethics Committee (Reference No. 2025-32051-65749-4) and has been conducted according to the corresponding ethical standards.

A.9.3 Annotation instructions

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

A.9.4 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:** participants 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.

Source domain	Examples	Avg metaphoricity score	# of candidates	# after filtering
ANIMAL	shelter, cage, swarm	0.3045	9845	544
WAR	army, hordes	0.3974	16455	3695
WATER	flood, tide, pour	0.4375	3661	1402

Table 9: Statistics for ANIMAL, WAR and WATER source domains

Batch	Agreement rate	Majority vote rate
1	0.68	0.74
2	0.66	0.73
3	0.71	0.70
4	0.75	0.72
5	0.73	0.75
6	0.55	0.68
7	0.77	0.82
8	0.61	0.73
9	0.69	0.63
Mean	0.68	0.72

Table 10: Per batch IAA rate

A.9.5 Compensation

Each annotation batch was completed by four independent annotators. Participants were allotted up to 45 minutes to complete a batch of up to 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.9.6 Annotation agreement and adjudication

Table 10 shows averaged pair-wise annotation agreement between 4 annotators (6 comparisons) within each batch, as well as overall (mean) agreement across batches. Since annotators were allowed to choose multiple labels per sample, agreement is defined as at least one of those labels being the same, i.e. any overlap between the selected labels. To evaluate how consistent the labels were among all four annotators, we also report the rate of strong majority vote agreement, i.e. percentage of samples where the same (and only one) label was chosen by at least three out of four annotators.

Whenever the majority vote could not be established (or the majority vote was OTHER DOMAIN, i.e. not one of the suggested classes), the labels were adjudicated by the first author of the paper, a postdoctorate researcher with background in linguistics and pragmatics. The adjudicator considered the selected labels, as well as the labels for

samples containing related words (for example, if the label for “vision” could not be established, the labels for “visions” were checked). In cases where none of the five suggested labels was appropriate, the adjudicator was allowed to choose the label from the full list of source domains. The adjudicator also ensured that the source domains were assigned consistently between multiple samples containing the same metaphor.

A.10 Most often confused source domains

To get the list of the most often confused labels, we calculate normalized point-wise mutual information and weight the score by the frequency of label pairs:

$$\text{NPMI}_w(a, b) = \frac{\log \frac{p(a,b)}{p(a)p(b)}}{-\log p(a, b)} \times \log c(a, b)$$

where $c(a,b)$ is the co-occurrence count of labels a and b .

A.11 Prompts for end-to-end detection of salient semantic frames within WATER source domain

A.11.1 System prompt

You are a linguist specializing in discourse analysis. You will be given two collection of tweets - by liberals and by conservatives. You need to analyse how WATER metaphors are used by these sides. Output the most salient (prominent) semantic frames used for WATER metaphors by liberal and conservative sides.

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 two lists, each containing salient semantic frames and their examples. COPY specific examples from tweets. ””

Label A	Label B	Count	NPMI	NPMI _w
BATTLE	WAR	33	0.751	2.627
COMPETITION	STRUGGLE	30	0.554	1.884
OBJECT HANDLING	PHYSICAL OBJECT	37	0.423	1.528
BATTLE	STRUGGLE	17	0.393	1.114
STRUGGLE	WAR	16	0.367	1.019
MOVEMENT	VEHICLE	19	0.286	0.843
BARRIER	PATHWAY	9	0.374	0.822
FORWARD MOVEMENT	UPWARD MOVEMENT	5	0.390	0.628
PHYSICAL HARM	TOOL	4	0.433	0.601
LIGHT	VISION	9	0.264	0.579

Table 11: Top co-occurring label pairs by weighted NPMI (within the labels chosen by the same annotator)

Label A	Label B	Count	NPMI	NPMI _w
COMPETITION	STRUGGLE	17	0.658	1.863
BATTLE	WAR	12	0.716	1.778
OBJECT HANDLING	PHYSICAL OBJECT	21	0.442	1.344
BARRIER	PATHWAY	9	0.494	1.086
LIGHT	VISION	11	0.443	1.062
BUILDING	SHAPE	7	0.429	0.835
STRUGGLE	WAR	8	0.388	0.807
BATTLE	STRUGGLE	7	0.367	0.715
FORCEFUL EXTRACTION	OBJECT HANDLING	7	0.357	0.695
AVERSION	STRUGGLE	4	0.450	0.624

Table 12: Top co-occurring label pairs by weighted NPMI (within the labels chosen by all annotators)

A.11.2 User prompt

What are the prevalent semantic frames used in WATER metaphors by liberals and conservatives?

Liberal tweets: {liberals}

Conservative tweets: {conservatives}

A.12 Examples of end-to-end LLM metaphor analysis

A.12.1 ANIMAL source domain

Liberal Semantic Frames:

- Abusing** - Liberals frequently use animal metaphors to describe how immigrants are being mistreated, with phrases like "treating immigrants like animals," "caging children like animals," and "treating them worse than dogs"
- Containing** - Extensive use of metaphors describing detention facilities as "cages," "kennels," "dog pounds," and "pens" where immigrants are held
- Hunting** - Metaphors describing ICE

and border patrol activities as "hunting down immigrants," "stalking," and "preying on" vulnerable populations

- Dehumanization** - References to how Trump and officials call immigrants "animals," "snakes," "cockroaches," and other dehumanizing terms

- Protecting** - Calls to protect immigrants from being treated like animals, with references to animal welfare standards being better than immigrant treatment

Conservative Semantic Frames:

- Categorization** - Frequent use of "animals" to categorize MS-13 gang members and violent criminals, distinguishing them from law-abiding citizens

- Being_in_category** - Classifying certain immigrants as "illegal aliens" versus legal immigrants, often using animal terms for the former group

- Swarming** - Describing immigration as animal-like mass movement with terms like "swarms," "herds," "flocks," and "stampedes" of immigrants

- Hunting** - References to "hunting down" and "rounding up" illegal immigrants for deportation

5. ****Proliferating_in_number**** - Using breeding metaphors to describe immigrant population growth, with terms like "breeding like rabbits" and concerns about demographic changes

A.12.2 WATER source domain - based on tweets that contain WATER metaphors

****Liberal Side - Most Salient Semantic Frames:****

1. ****Fluidic_motion**** - This is the most prominent frame, with liberals using water metaphors to describe the movement and flow of immigrants in neutral or sympathetic terms. Examples include "flow of migrants," "stream of immigrants," "waves of immigration," and "tide of migrants." These metaphors often emphasize the natural, ongoing nature of migration.
2. ****Motion**** - Closely related to fluidic motion, this frame captures the directional movement of people, often described as "pouring in," "flooding," or "streaming" across borders, but typically used to critique anti-immigration policies rather than the immigrants themselves.
3. ****Catastrophe**** - Used to describe the humanitarian crisis and dangerous conditions migrants face, such as "drowning" in the Rio Grande, being "swept away," or facing "floods" of persecution in their home countries.

****Conservative Side - Most Salient Semantic Frames:****

1. ****Fluidic_motion**** - Also the most prominent frame for conservatives, but used with negative connotations to describe an overwhelming, uncontrolled influx. Examples include "flood of illegals," "pouring across the border," "stream of illegal aliens," and "tide of illegal immigration." The emphasis is on the volume and lack of control.
2. ****Catastrophe**** - Used to describe the perceived negative impact of immigration on American society, such as the country being "flooded," "swamped," or "drowning" in illegal immigrants, suggesting an overwhelming disaster.
3. ****Emptying**** - This frame appears in contexts of "draining" resources, the economy, or social services due to immigration, suggesting that immigrants are depleting American resources like water draining from a container.

A.12.3 WATER source domain - based on tweets that do NOT contain WATER metaphors

****Liberal Tweets - Most Salient Water Metaphor Semantic Frames:****

1. ****Fluidic_motion**** - Examples: "immigrants...pour in and infest our country", "flooding our country with

immigrants", "immigrants are swarming to the border", "migrants fleeing violence in El Salvador still plan for the U.S."

2. ****Catastrophe**** - Examples: "more hopeless. I feel like we are about to get buried by a tsunami...Iran/Roe/Immigration", "We're not being invaded from any direction, we're not under attack by immigrants"

3. ****Abounding_with**** - Examples: "Trump refers to immigrants as an infestation 46 times", "immigrants...pour in and infest our country", "calling migrants invaders and an infestation"

****Conservative Tweets - Most Salient Water Metaphor Semantic Frames:****

1. ****Fluidic_motion**** - Examples: "illegals are flooding states", "Democrats divert 137B tax\$ ANNUALLY, to support & defend illegals", "illegals flooding the country", "flood our country with immigrants", "illegals pouring in"

2. ****Catastrophe**** - Examples: "NATIONAL EMERGENCY: 2,000+ Migrants Quarantined After Bringing Disease", "This is an invasion", "illegals flooding your border overwhelming your system"

3. ****Scarcity**** - Examples: "illegals cost America over 300 Billion", "illegals are costing us billions", "Each illegal immigrant costs taxpayers \$70k per year", "illegals cost US well over 300 Billion a yr"

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.