Broadening the coverage of computational representations of metaphor through Dynamic Metaphor Theory

Xiaojuan Tan[♣] and Jelke Bloem[♢♠]

♣ Department of Language, Literature and Communication, Vrije Universiteit Amsterdam
♢ Institute for Logic, Language and Computation, University of Amsterdam
♠ Data Science Centre, University of Amsterdam
tanxiaojuan01@gmail.com, j.bloem@uva.nl

Abstract

Current approaches to computational metaphor processing typically incorporate static representations of metaphor. We aim to show that this limits the coverage of such systems. We take insights from dynamic metaphor theory and discuss how existing computational models of metaphor might benefit from representing the dynamics of metaphor when applied to the analysis of conflicting discourse. We propose that a frame-based approach to metaphor representation based on the model of YinYang Dynamics of Metaphoricity (YYDM) would pave the way to more comprehensive modeling of metaphor. In particular, the metaphoricity cues of the YYDM model could be used to address the task of dynamic metaphor identification. Frame-based modeling of dynamic metaphor would facilitate the computational analysis of perspectives in conflicting discourse, with potential applications in analyzing political discourse.

Keywords: Dynamic metaphor theory, conflicting discourse, metaphoricity, extended metaphor

1. Introduction

The ubiquity and power of metaphor in discourse have served as an impetus for computational linguists to develop ways to automatically identify and process metaphors. Although various computational approaches to metaphor have been developed, the problem is far from solved partially due to incomplete coverage of the nature and mechanism of metaphor. Computational work typically focuses on representing the surface realizations of metaphor (called 'linguistic metaphor' in the typology of Shutova, 2015) or the metaphorical mappings underlying them as an inventory of mappings (called 'conceptual metaphor'), but these representations are static.

Conversely, cognitive linguists have made breakthroughs in developing dynamic metaphor models, but these models are only applied in manual semantic and/or pragmatic analysis. This article aims to bring together these two lines of work, by connecting recent developments in metaphor theory to recent computational approaches. In particular, we take insights from dynamic metaphor theory and discuss how existing computational models of metaphor might benefit from representing the dynamics of metaphor, particularly in the case of conflicting discourse. With this theoretical contribution, we aim to outline a theoretically informed path towards computational representations of metaphors that go beyond static metaphors and to introduce cognitive linguists to the possibility of the computational modeling of dynamic metaphors.

2. Metaphors are Dynamic

The most well-known metaphor theory is Lakoff and Johnson's (1980) Conceptual Metaphor Theory (CMT). According to CMT, the essence of metaphor is understanding and experiencing one kind of thing in terms of another. Accordingly, metaphor structures a cross-domain mapping of thought, from a relatively concrete target domain to an abstract source domain. As an example, in everyday life we often come across expressions that reason the target domain of 'love' in terms of the source domain of JOURNEY:

(1) We're *at* a *crossroads*.
We can't *move forward*.
I don't think our relationship is *going anywhere*.

The italicized metaphors here are called linguistic metaphors, while the mapping between the source domain and target domain (LOVE IS JOURNEY) is termed a conceptual metaphor.

Dynamic metaphor theory is a key recent development in the field of metaphor studies. A commonly held assumption in many earlier and contemporary metaphor theories is that metaphor is static. From the static viewpoint, linguistic metaphors are either dead (conventional) or alive (novel). For instance, Black (1979) regards conventional metaphorical expressions as dead and only novel metaphorical expressions as alive. In *Metaphors We Live By*, Lakoff and Johnson (1980) imply that the category of 'live' metaphor is much larger than generally assumed and should encompass the conventional metaphors of ordinary lan-

First Workshop on Reference, Framing, and Perspective, pages 40–50

guage. However, their perspective on metaphors is still static in that metaphors are restricted to a fixed cognitive structure of thought. Lakoff (1993, p. 210) characterizes the mapping of two frames as universally "fixed patterns of ontological correspondences" between two conceptual domains. Such a static view of metaphor has been attacked by discourse analysts because Lakoffian works take conceptual metaphors as highly stabilized conceptual mappings across speech communities. Rejecting the static view, Müller (2008) was the first to argue that the property of metaphor has the potential for activation and thus metaphor is dynamic. She claims that in certain discourse contexts, the source domains and conventionalized linguistic metaphors may be active for a given speaker at a given moment in time. This argument can be illustrated via the following two examples.

(2) We have to do these things to make America great again. Because we can't *lose* almost $800 billion on the start of the trade dispute, like has been done for many years. (Donald Trump's speech, 24/01/2019; italicization added)

(3) We were *losing* all our cases in the World Trade Organization. Almost every case, we were *lost*, *lost*, *lost*. (Donald Trump's speech, 13/08/2019; italicization added)

In example 2, the conventional linguistic metaphor 'lose' indicates a possible mapping between the source domain of COMPETITION and the target domain of 'trade dispute'. The metaphoricity is static since the metaphor occurs only once without any other semantic elaboration. However, in example 3, 'lose' becomes a more salient linguistic metaphor and COMPETITION becomes a more salient source domain in Trump's trade speech, through a strategy that foregrounds metaphor use – the repetition of the lexeme 'lose' in different verb tenses. This is what we call a metaphoricity cue. Through this cue, the metaphoricity of 'lose' is activated by former president Trump and/or the speech writers at that moment. Therefore, 'lose' is not static but dynamic, as it is no longer strictly restricted to the rigid category of *conventional* linguistic metaphor. Its metaphoricity achieves a higher degree of activation in example 3 than in example 2.

Challenging the static perspective of metaphor which has been taken for granted for decades, pivotal dynamic metaphor scholars have put forth different usage-based models (e.g., Cameron, 2010; Jensen and Cuffari, 2014; Müller, 2008). However, these dynamic perspectives are limited due to their focus on either the change of linguistic metaphor or the change of source domain. For instance, Cameron (2010) focuses on patterns of development of metaphorical expressions, while Kyratzis (1997) focuses on the chains of the source domains. Until recently, there has been little attention paid to the mechanisms by which changes in both source domains and/or target domains can activate metaphoricity.

The recently proposed YinYang Dynamics of Metaphoricity (YYDM, Tan, 2023; Tan and Cienki, 2024) addresses this theoretical gap. This usage-based model, created for metaphor analysis in texts, emphasizes how change within and between source and/or target domains can activate metaphoricity. As this model defines dynamicity in terms of explicit metaphoricity cues that are textually expressed, we consider YYDM a promising theoretical framework for the computational modeling of dynamic metaphors. It puts forward that metaphor develops with the emotions and attitudes of discourse participants, which outlines a way to empirically reconstruct the inner mechanism of dynamic metaphors, the motivation behind their use, and therefore their effect on society.

2.1. YinYang Dynamics of Metaphoricity

The model of YinYang Dynamics of Metaphoricity (YYDM) assumes that metaphorical expressions range from Yin-inactive metaphors to Yang-active metaphors; there is no strict boundary between them. Yin-inactive metaphors have a low degree of metaphoricity because they are not surrounded by any metaphoricity cues (cf. example 2). On the contrary, Yang-active metaphors have a high degree of metaphoricity because they are surrounded by metaphoricity cues (cf. repetition in example 3). The same metaphorical expression can be inactive in one context and active in another. The degree of activation of metaphoricity can be documented through Tan and Cienki's (2024) metaphoricity cues. In this section, we show a limited sample of different types of metaphoricity cues through examples from conflicting discourse.[1]

2.1.1. Cues highlighting the same source domain

Clustering of metaphorical expressions in the source domain This is exemplified by the clustered metaphorical expressions ('capitulation', 'submission', and 'retreat') that all highlight the WAR source domain in the following example: "On the question of foreign trade, previous leaders were guided by a shameful policy of capitulation, submission, and retreat."

[1] Examples taken from the Trump subcorpus and Xi subcorpus of Tan and Cienki's (2023) corpus on US-China trade conflict.

Explicit mapping Presenting the source domain explicitly. Consider the explicit mapping of CATASTROPHE (WTO IS CATASTROPHE): "World Trade Organization is a catastrophe."

Marking devices are used (Goatly, 1997, pp. 262-263; Cameron and Deignan, 2003) to mark the source domain, e.g., 'sort of', 'like', 'kind of', 'really', 'imagine', 'so to speak', 'actually', 'literally', 'if you like', 'in a way', 'as it were'. For instance, using "as" to mark the source domain LEVERAGE: "...the previous administration was unwilling to use our huge trade deficit with China as leverage..."

Repetition of the same linguistic metaphor with the same source domain. In this example, the cue is repeating 'stole' within the source domain of CRIME: "...other countries stole our factories, stole our plants, stole our wealth, and stole our jobs."

2.1.2. Cues indicating the change of source domain, but non-change of target domain

Diversification Using different source domains to refer to the same target domain is a metaphoricity cue. Consider the diversified source domains (POISON; GOOD PRESCRIPTION) in the following example, i.e., "Trade protectionism is a poison rather than a good prescription."

Novelization Using novel linguistic metaphors to refer to the new source domain is a metaphoricity cue. Considering the novel linguistic metaphor 'top student' to refer to the new source domain (TOP STUDENTS): "China is a top student among the members of the World Trade Organization."

2.1.3. Cue indicating the change of target domain, but non-change of source domain

Multivalency Using the same source domain to refer to different target domains exemplifies this metaphoricity cue. Consider the repeated source domain PILLAR and different target domains 'trade policy' and 'trade regulation' in the following example, i.e., "In addition to trade policy, trade regulation is also the pillar of our economic development"

2.1.4. Cue indicating the change of both source domain and target domain

Mixing different source domains mapped to different target domains illustrates this cue. Consider the mixed mappings (CHINA IS A TOP STUDENT; WTO IS CONTAINER) in the following example, i.e., "China has been a top student since its entry into the World Trade Organization".

2.2. Exemplification of YYDM through data on conflicting discourse

Generally, the more semantic representation of a source domain is present in a discourse, the more the source domain is foregrounded and a higher degree of activated metaphoricity is achieved. The more metaphoricity cues that point to a linguistic metaphor, the more the linguistic metaphor is highlighted and the higher the degree of activation (cf. Tan, 2023). The following examples of 'war' metaphors from discourse on the recent U.S.-China trade war will illustrate more clearly how the dynamic model functions. These examples are selected because they use commonly studied frames that also exist in frame repositories such as FrameNet (Baker et al., 1998) and these frames can show how the divergent opinions of the Chinese and American governments evolve in the process of trade negotiation.

(4) Q Talking about a trade *war*? PRESIDENT TRUMP: I don't think you'll have a trade *war*. Q No trade *war*? PRESIDENT TRUMP: I don't think so. I don't think you're going to have trade *war*, no. (Remarks, 05/03/2018; italicization added)

(5) Q On the tariffs, the President tweeted that trade *wars* are good, easy to *win*. Can you explain what he meant by that? MS. SANDERS: Look, the President, I think, is very confident that if that's where we ended up, we certainly would *win*. But that's not the *goal*. The *goal* is to get free, fair, and reciprocal trade, and hope that other countries will join in. (Press Briefing of Press Secretary, 05/03/2018; italicization added)

From example 4 to example 5, the Yin-inactive metaphor 'war', framing the target domain of 'trade negotiation', became a Yang-active metaphor on March 5th, 2018 through different metaphoricity cues. In example 4, the *repetition* of 'war' activated both the metaphorical expression 'war' and the source domain of 'WAR'. It shows former president Trump's position in the morning that he could threaten China to make concessions in trade negotiations without launching a trade war. However, in example 5, 'war' and the source domain of WAR are further highlighted in a press briefing.

The journalist first activates 'war' and WAR through *a cluster of* WAR metaphors ('war', 'win'). Then the Press Secretary foregrounds them further through the *repetition* of 'win', *the change of the source domain* (the change from WAR to JOURNEY via the change from 'war' to 'goal'), and *the change of both the source domain and target domain* (from TRADE IS JOURNEY to GET FREE, FAIR, AND RECIPROCAL TRADE IS GOAL). With the activation of the metaphoricity, Trump's administration changed their attitude and triggered a nationalist sentiment of winning the war, i.e., from the non-necessity of a trade war to the determination to get free trade through a trade war. This kind of trade discourse from the Trump administration was attacked by the Chinese government, which can be shown through the activation of Chinese dynamic metaphors below.

(6) 历史已经证明，贸易战没有赢家，中国不愿意打贸易战。'The history has proved that there is no *winner* in the trade *war*. China is not willing to *fight* a trade *war*' (Reports of Leaders' Activities, 06/03/2018; italicization added)

(7) Q 中方目前的态度比较克制，但并不代表没有好牌... 贸易战会发展到什么程度，要看美国走到哪一步。中国要反击这场贸易战的 "牌" 有不少，从大豆到汽车、飞机，可以打出组合拳来回击，这些商品的可替代性都比较强。对于美方挑起的贸易战，我们完全有底气采取强有力措施精准还击。'China's current attitude is relatively restrained, but it does not mean that there is no *good card*...the extent of the trade *war* depends on the procedure taken by the U.S. China has many *"cards"* to *fight back* against this trade *war*. From soybean to car and airplane, it can *hit back* with a *combination combo*. These commodities are highly replaceable. For the trade *war* provoked/*shouldered* by the U.S., We have the confidence to take strong measures to *fight back accurately*.' (China Daily, 26/03/2018; italicization added)

From examples 6 to 7, the source domain WAR becomes more and more salient, and 战 'war' changes from a Yin-inactive metaphor to a Yang-active metaphor within a month. In example 6, the repetition of 战 'war' and the clustered metaphors of WAR (打 'fight', 战 'war', 赢家 'winner') activate 战 'war' within the same source domain of WAR. Built on the activation of metaphor, the Chinese government conveyed its stance that China was unwilling to go to war on March 6th, 2018, which replies to America's decision to launch a trade war on March 5th, 2018. On March 26th, the Chinese attitude evolved to counterattack, which aroused a nationalist sentiment of protecting China through a trade war. This was shown by a higher activation of 战 'war' and WAR through many metaphoricity cues across sentences in example 7.

At the beginning of example 7, China is framed as a card player having a set of good cards in the CARDS game and then is *reframed* as a defender fighting back the U.S. aggression with a series of WAR weapons. With the *change of source domain* (from CARDS to WAR), the cards of soybeans, cars, and airplanes are *reframed* as weapons. As the news report continues, the *collocated* metaphors (打出 'hit' and 组合拳 'combination combo') introduce a *new* source domain for trade (BOXING COMBO). That is, the actions of playing cards are reconstructed as blows in a BOXING COMBO, which reconstructs China as a boxer hitting back the U.S. through the combo, and foregrounds WAR and 战 'war' even further through another *reframing*. In the next sentences, WAR and 战 'war' are even more foregrounded through *aggregated* metaphoricity cues. Namely, a *new* reframing reconstructs trade war as a PHYSICAL OBJECT through 挑起 'shoulder'. With the *change of source domain and target domain*, China attributes the guilt of starting the trade war to the U.S. The following *clustered metaphors* of WAR (e.g., 反击/反制 'counterattack', 精准还击 'fight back accurately') as well as the *repetition* of 战 'war' and 精准 'accurate(ly)', portray China's strong skills in counterattacking and its confidence in winning the trade war.

Applying the dynamic model (YYDM) to authentic data gathered from discourse on trade conflicts, these examples reveal that metaphors can be activated and become dynamic through additional semantic representations of the source domain, and additional changes of the source domain and/or target domains. Dynamic metaphors can connect various thoughts and participants over stretches of texts and even an entire large-scale corpus. With the development of metaphors, the sentiments and attitudes on the China-U.S. trade war also changed. This shows discourse is not a matter of detached meaning construction but instead a dynamic system intertwined with intersected levels (e.g., linguistic, conceptual, sociopolitical) which needs to be understood as processes, flows, or movements (Larsen-Freeman and Cameron, 2008). Since the dynamics in the micro level of language use function all the way up to the discourse dynamics at the social group level (Tan et al., 2024), automatic identification of metaphoricity cues at the micro level can lead to the prediction of different changes driving the production and reproduction of conflicting discourses.

3. Computational representations

Next, we examine the extent to which the dynamics of metaphor might be represented in the field of NLP. In recent years, several excellent surveys on the state of computational metaphor processing have been written (Rai and Chakraverty, 2020; Tong et al., 2021; Ge et al., 2023), which we will not reiterate here. We instead aim to survey computational work that constructs detailed or extended computational representations of metaphor, which may cover aspects of the dynamics of metaphor. Broadly, computational work on metaphor has centered around two tasks: automated metaphor identification and automated interpretation. Typically, identification is operationalized as a sequence labeling task, and interpretation is operationalized as a paraphrasing task.

For many years, computational metaphor processing relied mainly on hand-crafted resources such as MetaNet (Dodge et al., 2015). MetaNet is a multilingual repository of conceptual metaphors that is linked to FrameNet (Baker et al., 1998), enabling computational representation of conceptual metaphors in terms of source and target domains, and theoretically grounding those domains in terms of frame semantics (Fillmore, 1976). For metaphor identification, metaphoric expressions can be linked to conceptual metaphors as listed in MetaNet. However, such approaches have limited coverage, with few possibilities to generalize beyond the hand-crafted metaphor inventory.

More recent approaches rely on the use of dense vector representations as features for predicting metaphoricity labels. Typically, the data and labeling from the VU Amsterdam (VUA) Metaphor Corpus (Steen et al., 2010) are used. This corpus contains token-level binary annotation indicating metaphoric or non-metaphoric use. These tools identify a wider range of metaphoric expressions than those relying on metaphor repositories, but lack explanatory power and theoretical grounding in metaphor theory. Such tools do not tell us why an expression is metaphoric, e.g. by performing conceptual mapping, identifying it as an instance of a particular conceptual metaphor with a particular source and target domain. An example of this approach is Gong et al.'s (2020) RoBERTA-based system.

A particularly accessible example of this approach is Mao et al.'s (2023) MetaPro 2.0, an end-to-end metaphor processing system incorporating the tasks of identification and interpretation with state of the art performance on standard benchmarks. The identification module is trained on the VUA corpus, and the interpretation paraphrasing is done by having RoBERTa mask a metaphorically used word and predict a synonym or hypernym of that word in its place. The approach is limited to substituting metaphorically used words with a synonym or hypernym that fits the context literally, excluding more creative metaphoric uses. At the time of writing, this system is available as a functioning online demo[2]. It is thus a good way for cognitive linguists to assess the state of the art in computational metaphor processing.

[2]https://metapro.ruimao.tech/

3.1. Conceptual mapping

A few computational studies do address the conceptual mapping task in addition to metaphor identification. Firstly, the aforementioned MetaNet (Dodge et al., 2015) was used to perform conceptual mapping by identifying candidate metaphoric expressions through grammatical patterns and then matching the words in the source domain slot and target domain slot to frames using MetaNet, FrameNet, Wordnet or Wiktionary. If those frames have metaphoric mappings in the hand-coded MetaNet repository, it is identified as metaphoric. This formalizes connections between different instances of a conceptual metaphor but does not generalize to novel mappings. Although this is not discussed in the paper, due to the use of a standardized resource, different instances of the same conceptual metaphor occurring throughout various discourses can easily be connected using this approach.

Shutova et al. (2017, p. 79) emphasize the importance of conceptual metaphoric mappings, stating that "one needs to address conceptual properties of metaphor, along with the surface ones". They use semi-supervised clustering to create source and target domains based on seed expressions (e.g. "grasp theory", "ideology embraces"). These expressions are used to learn how to map these domains, allowing the detection of mappings between other expressions within these domains that are not in the seed set. This is extended to an unsupervised method where the same clusters are automatically organized hierarchically by inferring connections at a hypernym level which can represent the kind of broad conceptual mappings that metaphors consist of. However, the mappings are static and limited to ones based on two-word verb-subject and verb-object relations without context. Domains also remain unlabeled and not linked to a word sense, frame or metaphor repository.

Ge et al. (2022) take conceptual mapping one step further by explicitly defining concepts as WordNet hypernyms, with the goal of increasing the explainability of metaphor identification methods. They use an algorithm to determine the level of hypernymy in the WordNet hierarchy that sufficiently covers most senses of a noun without being too abstract. This approach is also incorporated in the aforementioned MetaPro 2.0 (Mao et al., 2023) system to map to the source domain. For a given metaphor, paraphrases from MetaPro's interpretation module are regarded as labels of the metaphor's target domain. The concept resulting from the application of Ge et al.'s (2022) algorithm is regarded as the metaphor's source domain, yielding a mapping. However, the analysis is limited to decontextualized pairs of dependent words (verb-noun or adjective-noun), and the mapping to WordNet only exists for the source domain, not the target domain.

Wachowiak and Gromann (2023) predict source domains from GPT-3 given a sentence and a target domain in a one-shot text completion task. This form of metaphor mapping is fairly flexible by not being connected to pre-defined domains and by drawing upon a huge amount of training data. However, it presupposes that a specific linguistic metaphor statically maps to a source domain regardless of discourse context and the approach does not consider other aspects of dynamic metaphor theory such as the amount of activation of the metaphor. It may also yield non-standard source domain descriptions that do not map to metaphor inventories.

3.2. Extended metaphors

While no computational work directly addresses the dynamicity of metaphor, the related concept of extended metaphor is discussed by a few authors. However, these works diverge on what extended metaphors are. The aforementioned surveys of computational metaphor illustrate that "broader tasks of identifying conceptual metaphors, extended metaphors, and metaphoric framing, have been largely ignored" (Tong et al., 2021, p. 4679) and "identifying other types of metaphors, such as extended metaphors or MWEs, has yet to be well solved" (Ge et al., 2023, p. 1857).

Klebanov and Beigman (2010) probably comes closest to describing the dynamics of metaphor, discussing the case of bargaining in political communication. They describe an extended metaphor of the European Union as a train. This metaphor received various politically charged positive and negative extensions over time and was the subject of counter-metaphors in European political discourse. The authors note the extensive attempts at bargaining over the same metaphor rather than re-framing the discussion with a novel metaphor, and aim to explain this bargaining with a game-theoretical model. Klebanov and Beigman's (2010) model represents extended metaphor as an abstract set of frames that can be negotiated about, and this is the representation of extended metaphor used throughout the work. However, the discussion proceeds in terms of game-theoretic negotiation about these frames rather than in terms of cognitive metaphor theory – there is no representation of source and target domains. The work does illustrate the importance of acknowledging diverging perspectives on a metaphor in the domain of political communication and the need for a model that can represent this. The authors note the difficulty of automatic detection of extended metaphor due to a lack of sufficient training data for particular metaphors.

Subsequently, Shutova (2015, p. 585) states that "a computational method for identification and interpretation of extended metaphor in real-world discourse is yet to be proposed. A discourse-level metaphor processing system would need to identify a chain of metaphorical expressions in a text, which indicates a systematic association of the text topic with a particular domain. These chains would then demonstrate how continuous scenarios can be transferred across domains". This line of work regards extended metaphors as a group of linguistic metaphors elaborating on the same conceptual metaphor, which is close to the type of dynamic metaphor in section 2.1.1. However, they do not encompass other types of dynamic metaphors in which there is a change of frame (sections 2.1.2-2.1.4). Dynamic metaphor theory can provide a framework for operationalizing this.

Dankers et al. (2020) address another aspect of dynamic metaphor by emphasizing the importance of discourse context in the identification task. However, the task itself remains a binary identification of linguistic metaphor, just with a larger context window for the sequence prediction task. The analysis does not show that the model recognizes elements such as clusters of linguistic metaphors that refer to a particular conceptual metaphor, and it would not explicitly identify conceptual metaphors anyway as there is no conceptual mapping. The main improvements using this method come from having more information on the topic of the text and implicit coreference resolution.

We are aware of only one study that computationally models extended metaphor occurring across sentences. Jang (2017) argues that metaphor should be defined in terms of frames in order to be able to process metaphor at the discourse level. Metaphor is then defined as a phenomenon that occurs when "a speaker brings one frame into a context/situation governed by another, and explicitly relates parts of each". Extended metaphor is described as metaphors that can be around related metaphors. Specifically, in a dissertation chapter not published as a paper, Jang (2017, Ch. 8) applies template induction to find frame elements in a connected discourse, using this frame information for metaphor detection. In this context, the concept of extended metaphor is defined as a switch of source frames in discourse (corresponding to section 2.1.2 in our theoretical description): "metaphor performs social functions through the switching of frames" Jang (2017, p. 79). Other types of dynamic metaphor (sections 2.1.3-2.1.4) are not addressed, and the concept of dynamic metaphor is not mentioned.

In this work, the frame elements are extracted from lexico-grammatical patterns in a seeded but unsupervised way. The template induction then identifies more frame elements for the target frame, which may occur in the vicinity of the candidate linguistic metaphor. The final task is once again metaphor identification, thus the identified frame elements are only used as features to solve this task, and explicit representation of clusters of metaphoric expressions with the same source frame (as in Section 2.1.1) is not demonstrated or evaluated directly. Although the notions of frames and frame elements ('frame facets') and their linguistic instances ('slot instances') are used, they are not formalized. The frames are not connected to repositories such as FrameNet and it is not clear where the frame seed words come from.

4. Dynamic metaphors and frames

After comparing state-of-the-art theory on dynamic metaphor with the state of the art in computational metaphor processing, we identified several gaps that would need to be addressed in order to computationally represent the dynamics of metaphor.

4.1. Mapping clustered and repeated references to a metaphor

In the computational literature, instances of metaphor are largely viewed in isolation, and when they are not, surrounding instances within a discourse are mainly viewed as features to aid in the identification of a targeted instance. In the new dynamic model (YYDM), linguistic metaphors are connected and may partially instantiate different elements of the metaphor's source or target domain. These connected linguistic metaphors indicate the dynamic activation of the metaphor. A computational representation of dynamic metaphor should be able to formalize these connections, for example, by linking them to a common metaphoric mapping as in MetaNet. At the same time, it should represent the fact that some linguistic metaphors serve as cues marking the activation of a certain linguistic/conceptual metaphor, which is similar to the frame elements of Jang (2017). This should be possible even when the conceptual metaphor changes due to changes in source and/or target domains (cf. section 2.1.2- 2.1.4), unlike in static MetaNet representations.

4.2. Representing changing domains

An important aspect of dynamic metaphors is that a metaphor's source domain, target domain or both can change between instances. Firstly, this requires a computational model of metaphor to have explicit representations of a source and target domain, which requires conceptual mapping. This was only performed by a few studies until now, as we saw in section 3.1. These studies either map to unsupervised clusters, WordNet hypernym levels, or FrameNet frames.

Secondly, mappings should either be able to have different source and target domains, or should be a set of related mappings that map different source and target domains as part of a dynamic metaphor. This could be operationalized through something like FrameNet's inheritance relation or MetaNet's related metaphor property. Mappings can change even across longer spans of discourse, as a metaphor can be used dynamically throughout a political discourse across time. Therefore, it seems useful to have standardized frame and metaphor identifiers that instances of a metaphor can be linked to (e.g. Klebanov and Beigman's (2010) EU train metaphor).

Next, these frame changes are often used in conflicting discourse to represent or emphasize divergent perspectives as in our example 7 on the US-China trade war. This could be represented using something like FrameNet's perspectivization relation, where frames describing the same situation from different perspectives are linked. Metaphoric mappings might also be considered perspectivized in this way.

Lastly, in the YYDM model, metaphors have a degree of activation based on the amount of metaphoricity cues in their context. A computational model of dynamic metaphor would also have to account for this. This possibility is already addressed to some extent by Jang (2017) – the frame elements they detect in the vicinity of the frame involved in the metaphoric mapping can be considered as metaphoricity cues. A computational representation of a metaphor that can be active or inactive following the YYDM model could include the number of metaphoricity cues found to indicate whether it is a Yin-inactive metaphor or a Yang-active metaphor.

4.3. Choosing a representation

We propose that frame-based approaches to metaphor representation are the best choice for modelling dynamic metaphors computationally. These approaches can incorporate the necessary elements from the YYDM model in order to represent detailed aspects of metaphor that may change dynamically throughout a discourse (e.g., different frame-evoking words denoting different frame elements). Frame-based approaches can enable their alignment between source and target domains to facilitate analysis. They can also incorporate perspectivization, which is important if computational metaphor processing is to be applied to the study of conflicting discourse.

Existing approaches that get closer to representing dynamic aspects of metaphor, such as the conceptual mapping of MetaNet or the frame elements of Jang (2017), already draw ideas from Frame Semantics (Fillmore, 1985). Jang (2017, p. 92) also argues that "modeling metaphor through the lens of frame theory could be the first step in detecting extended metaphor". Frame-based approaches are well grounded in theory, particularly in cognitive linguistics, which is a framework that aligns well with the current data-driven and distributional paradigm in the field of natural language processing (Levshina and Heylen, 2014; van Trijp, 2017; Rambelli et al., 2019). Frame-based representations could also handle multi-word units, a weakness of most current computational approaches.

In previous work, frame-based approaches such as MetaNet have shown a lack of scalability due to their dependence on hand-crafted linguistic resources. However, with the recent increase of interest in grounding elements of large language models in linguistic theory, we are starting to see efforts to induce even complex linguistic representations such as frames from data. Yamada et al. (2021) perform semantic frame induction from contextual word embeddings, paving the way for automatic frame construction. They show that clusters of contextualized word representations can be used to distinguish the difference between multiple frames invoked by the same verb. Furthermore, recent work has shown that evidence for Construction Grammar constructions, another branch of linguistic theory related to cognitive linguistics, can be found in transformer-based sentence embedding models. Li et al. (2022) observed that argument structure constructions get clustered by their construction type (e.g. ditransitive, caused-motion) rather than by their verb, and Veenboer and Bloem (2023) note that constructions in embedding space are surrounded by nearest neighbors with similar constructional semantics, also generalizing to instances containing verbs not seen in example constructions.

Therefore, it may be possible to create frame-based representations of dynamic metaphors in an unsupervised way in the future, especially given some seed set of frames, metaphoric mappings, frame elements or annotated metaphoricity cues.

5. Towards identification

Besides the representation of dynamic metaphors, the YYDM model can also be operationalized to aid in the popular task of metaphor identification. We propose that the metaphoricity cues discussed in section 2.1 can be used not only to characterize, but also to identify dynamic metaphor use.

Specifically, cues such as **repetition** and **clustering** of metaphoric expressions (section 2.1.1) could be detected by searching for multiple frame elements in the context of an expression, as done by Jang (2017, Ch. 8). **Marking devices** indicating metaphoricity such as 'so to speak' could be used directly as identification features. This was backed with empirical corpus data by Cameron and Deignan (2003) ('tuning devices'), but surprisingly this work has never been cited in the field of NLP. **Explicit mapping** can also be detected if a model is able to perform conceptual mapping or draws upon a repository of metaphoric mappings - it would only require a small step to match the conceptual metaphor of CATASTROPHE (WTO IS CATASTROPHE) to the lexemes in "World Trade Organization is a catastrophe.", especially when a frame-semantic parser can be used.

The cues indicating frame changes are more abstract, as they require frame representations that may be beyond the capabilities of current frame-semantic parsers. We tried our example of the **diversification** cue, "Trade protectionism is a poison rather than a good prescription" with the Frame Semantic Transformer parser (Chanin, 2023). It does detect both source frames: the Toxic_substance frame, with *poison* filling the TOXIC_SUBSTANCE FE, and Usefulness frame, with *prescription* filling the ENTITY FE. However, a frame related to 'trade protectionism' is not detected, as this is a concept rather specific to the domain of trade war. As for existing metaphor representations, TRADE PROTECTIONISM IS POISON can be categorized as a subcase of the NEGATIVELY EVALUATED CONDITIONS ARE HARMFUL AGENTS mapping in MetaNet. The latter does exist, but the former is not in the repository, and neither is something corresponding to the domain-specific TRADE PROTECTIONISM IS PRESCRIPTION. However, missing frames could be induced (Yamada et al., 2021), substituted by domains defined as WordNet hypernyms or clusters of related words, as was done in work discussed in section 3.1.

With such frame-based representations of metaphor, metaphoricity cues indicating frame changes can be detected in theory. **Diversification** could be detected by identifying the various source frames in the discourse, and checking whether they map to the same target domain in an available inventory of metaphoric mappings. Detecting **multivalency** would be the inverse of this, identifying target domain frames instead. **Novelization** is more difficult to detect as novel metaphors would not be in an inventory of conventional metaphors. The novelization cue might be found by detecting linguistic patterns that look like metaphors (e.g. from induced templates as in Jang, 2017), where the target domain has been used before in the discourse and the source domain is unknown in the local context of that target domain. The **mixing** of source and target domains appears to be a difficult cue to detect, but if we take metaphor to be dynamic over a larger discourse, we could find them by restricting the search space to only metaphors that have already been used in the discourse. Instances of mixing of the source and target domains of these previously used metaphors can then be found by detecting novel combinations of the source and target domain frames within a limited context window.

Such work on cue-based identification of dynamic metaphor or induction of missing MetaNet mappings may be aided by an annotated metaphor corpus that includes annotation of metaphoricity cues. Tan and Cienki (2023) annotated a 6M word corpus of texts relating to US-China trade conflicts with detailed features of the YYDM model, including metaphoricity cues. This labeled data could be used to train a classifier that can use the metaphoricity cues as features for the task of dynamic metaphor identification. It could also be used to extend the computational task of static conceptual mapping to the task of dynamic mapping, where multiple different but related metaphoric mappings may exist within the same discourse.

6. Discussion

We have sketched a proposal for more comprehensive computational representations of metaphor. Using the model of YinYang Dynamics of Metaphoricity as a theoretical framework, we demonstrated that metaphors are dynamic rather than static. Next, we surveyed the state of the art in computational metaphor representation. We found that, although dynamic metaphor theory was never explicitly addressed computationally, some of its ingredients, such as conceptual mappings, are represented. We then proposed ways to incorporate the main elements from the YYDM model into computational representations of metaphor using frame representations. Lastly, we discussed how the metaphoricity cues of the YYDM model could be used to address the task of dynamic metaphor identification. Overall, we hope to have shown that approaches based on Frame Semantics, such as FrameNet, with the addition of an inventory of metaphoric mappings, such as MetaNet, provide the necessary ingredients for computational representation of dynamic metaphor. The main weakness of this approach is limited coverage, but combining recent work on frame induction with an annotated corpus of dynamic metaphor may help to address this.

The importance of representing metaphor dynamically in the computational domain lies in the increasing importance of representing different perspectives on events and issues. This is true in NLP where the real-world application of large language models has shown that aggregating all data points into a single distribution or ground truth label erases minority perspectives (Cabitza et al., 2023). In political discourse analysis, metaphor researchers holding the static view fail to demonstrate how metaphors can develop together with political perspectives. Political discourse is a dynamic system where metaphors developing at the micro-level of language use are dynamically intertwined with the hidden political interests and power at the macro-level of discourse context which influences the evolution of political perspectives.

Dynamic metaphors evolve in discourse over time and can be sustained over many years. Having computational representations thereof would open up the possibility of performing diachronic metaphor analysis by comparing diachronic representations. Research on lexical semantic change using diachronic word embeddings has been quite successful (Tahmasebi et al., 2021), but similar approaches have not been developed for metaphor.

Representations of dynamic metaphors may also have benefits for downstream NLP tasks. Metaphoric expressions in conflicting discourse are often used to express polarized sentiment, and detecting this could contribute to better sentiment analysis. When metaphors are explicitly resisted (van Poppel and Pilgram, 2023), they may carry negative sentiments and conflicting perspectives.

Event detection is another possible application area – dynamic use of metaphor can involve many mentions of a particular event from various perspectives, each adding more information about the event. Our analysis of example 4 and 5 shows that the former American president and the Press Secretary make multiple metaphoric references to a trade war in two statements on the same day. This points toward the possibility of detecting the evolution of big political conflicts, which complements the existing computational techniques that focus on the detection of detached events.

In a nutshell, this study aims to bring cognitive linguists and computational linguists together, by showing recent developments in metaphor theory as well as a path towards computational application. Contrary to the static view of cognitive linguists and computational linguists, this paper argues that the cognitive dimension (frames), affective dimension (sentiments), and social-political dimension (perspectives) are constantly interacting. This inherent variability of the discourse system has implications for experts from both fields. Future computational operationalizations of this new dynamic model applied to different datasets could have impactful applications in analyzing political discourse in general and in analyzing conflicting discourse in particular.

7. Bibliographical References

Collin F Baker, Charles J Fillmore, and John B Lowe. 1998. The Berkeley FrameNet project. In *COLING 1998 Volume 1: The 17th International Conference on Computational Linguistics*.

Max Black. 1979. More about metaphor. In A. Ortony, editor, *Metaphor and thought*, pages 19–43. Cambridge University Press.

Federico Cabitza, Andrea Campagner, and Valerio Basile. 2023. Toward a perspectivist turn in ground truthing for predictive computing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 6860–6868.

Lynne Cameron. 2010. The discourse dynamics framework for metaphor. In R. Maslen L. Cameron, editor, *Metaphor analysis: Research practice in applied linguistics, social sciences and the humanities*, pages 77–94. Equinox London.

Lynne Cameron and Alice Deignan. 2003. Combining large and small corpora to investigate tuning devices around metaphor in spoken discourse. *Metaphor and symbol*, 18(3):149–160.

David Chanin. 2023. Open-source frame semantic parsing. *arXiv preprint arXiv:2303.12788*.

Verna Dankers, Karan Malhotra, Gaurav Kudva, Volodymyr Medentsiy, and Ekaterina Shutova. 2020. Being neighbourly: Neural metaphor identification in discourse. In *Proceedings of the second workshop on figurative language processing*, pages 227–234.

Ellen K 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.

Charles Fillmore. 1976. Frame semantics and the nature of language. *Annals of the New York Academy of Sciences*, 280(1):20–32.

Charles Fillmore. 1985. Frames and the semantics of understanding. *Quaderni di semantica*, 6:222–254.

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.

Andrew Goatly. 1997. *The language of metaphors*. Routledge.

Hongyu Gong, Kshitij Gupta, Akriti Jain, and Suma Bhat. 2020. Illinimet: Illinois system for metaphor detection with contextual and linguistic information. In *Proceedings of the Second Workshop on Figurative Language Processing*, pages 146–153.

Hyeju Jang. 2017. *Computational Modeling of Metaphor in Discourse*. Ph.D. thesis, Carnegie Mellon University.

Thomas Wiben Jensen and Elena Cuffari. 2014. Doubleness in experience: toward a distributed enactive approach to metaphoricity. *Metaphor and Symbol*, 29(4):278–297.

Beata Beigman Klebanov and Eyal Beigman. 2010. A game-theoretic model of metaphorical bargaining. In *Proceedings of the 48th annual meeting of the Association for Computational Linguistics*, pages 698–709.

Sakis Kyratzis. 1997. *Metaphorically speaking, sex, politics, and the Greeks*. Ph.D. thesis, Lancaster University.

George Lakoff. 1993. The contemporary theory of metaphor. In A. Ortony, editor, *Metaphor and thought*, page 202–251. Cambridge University Press.

George Lakoff and Mark Johnson. 1980. *Metaphors we live by*. University of Chicago Press.

Diane Larsen-Freeman and Lynne Cameron. 2008. *Complex systems and applied linguistics*. Oxford University Press.

Natalia Levshina and Kris Heylen. 2014. A radically data-driven construction grammar: Experiments with Dutch causative constructions. *Extending the scope of Construction Grammar*, 54:17.

Bai Li, Zining Zhu, Guillaume Thomas, Frank Rudzicz, and Yang Xu. 2022. Neural reality of argument structure constructions. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7410–7423.

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.

Cornelia Müller. 2008. *Metaphors dead and alive, sleeping and waking: A dynamic view*. University of Chicago Press.

Sunny Rai and Shampa Chakraverty. 2020. A survey on computational metaphor processing. *ACM Computing Surveys (CSUR)*, 53(2):1–37.

Giulia Rambelli, Emmanuele Chersoni, Philippe Blache, Chu-Ren Huang, and Alessandro Lenci. 2019. Distributional semantics meets Construction Grammar. Towards a unified usage-based model of grammar and meaning. In *First International Workshop on Designing Meaning Representations (DMR 2019)*.

Ekaterina Shutova. 2015. Design and evaluation of metaphor processing systems. *Computational Linguistics*, 41(4):579–623.

Ekaterina Shutova, Lin Sun, Elkin Darío Gutiérrez, Patricia Lichtenstein, and Srini Narayanan. 2017. Multilingual metaphor processing: Experiments with semi-supervised and unsupervised learning. *Computational Linguistics*, 43(1):71–123.

Gerard Steen, Aletta G Dorst, J Berenike Herrmann, Anna Kaal, Tina Krennmayr, Trijntje Pasma, et al. 2010. A method for linguistic metaphor identification. *Amsterdam: Benjamins*.

Nina Tahmasebi, Lars Borin, and Adam Jatowt. 2021. Survey of computational approaches to lexical semantic change detection. *Computational approaches to semantic change*, 6(1).

Xiaojuan Tan. 2023. *Static and dynamic metaphoricity in US-China trade discourse: A transdisciplinary perspective*. Ph.D. thesis, Vrije Universiteit Amsterdam.

Xiaojuan Tan and Alan Cienki. 2023. Dynamic metaphor theory.

Xiaojuan Tan and Alan Cienki. 2024. A dynamic metaphor perspective on Trump and Xi's trade negotiation in governmental discourse. *Language & Communication*, 96:42–53.

Xiaojuan Tan, Alan Cienki, and Bertie Kaal. 2024. The diachronic and cross-linguistic use of trade metaphors in us-china governmental discourse: A socio-cognitive approach. *Metaphor and the Social World*, 14:130–153.

Xiaoyu Tong, Ekaterina Shutova, and Martha Lewis. 2021. Recent advances in neural metaphor processing: A linguistic, cognitive and social perspective. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4673–4686.

Lotte van Poppel and Roosmaryn Pilgram. 2023. Types of resistance to metaphor. *Metaphor and Symbol*, 38(4):311–328.

Remi van Trijp. 2017. A computational construction grammar for English. In *2017 AAAI Spring Symposium Series*.

Tim Veenboer and Jelke Bloem. 2023. Using collostructional analysis to evaluate bert's representation of linguistic constructions. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 12937–12951.

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.

Kosuke Yamada, Ryohei Sasano, and Koichi Takeda. 2021. Verb sense clustering using contextualized word representations for semantic frame induction. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4353–4362.