

Intersecting Word Vectors to Take Figurative Language to New Heights

Andrea Gagliano & Emily Paul & Kyle Booten & Marti A. Hearst

University of California at Berkeley

{andrea.gagliano, emily.paul, kbooten, hearst}@berkeley.edu

Abstract

This paper proposes a technique to create figurative relationships using Mikolov et al.'s word vectors. Drawing on existing work on figurative language, we start with a pair of words and use the intersection of word vector similarity sets to blend the distinct semantic spaces of the two words. We conduct preliminary quantitative and qualitative observations to compare the use of this novel intersection method with the standard word vector addition method for the purpose of supporting the generation of figurative language.

1 Introduction

*“I sit in my chair all day and work and work
Measuring words against each other.”*
- Conrad Aiken *Improvisations: Light And Snow*

While metaphors are part of everyday language, in poetry they are vital. Metaphorical language, in contrast with literal or non-metaphorical language, “mak[es] use of structure imported from a completely different conceptual domain” (Lakoff and Turner, 1989).

Lakoff and Turner analyze famous poems and show how they can be understood as the blending of concepts from multiple metaphorical frames. For example, they state that a common metaphor is DEATH AS DEPARTURE and provide an example of an Emily Dickinson poem in which she merely needs to mention the words “death” and “carriage” in the same set of stanzas for the reader to know that the carriage is not taking a spin around the block, but rather a one-way trip with no return.

*“Because I could not stop for Death –
He kindly stopped for me –*

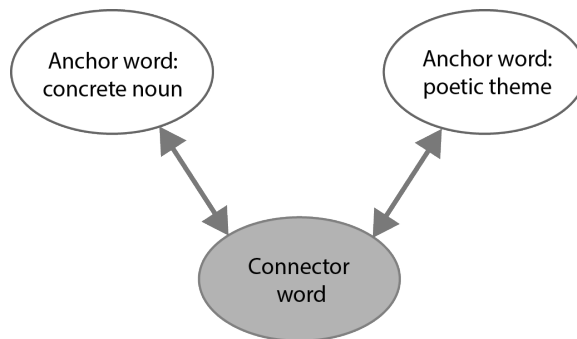


Figure 1: Connector word drawing together the two semantic spaces of the anchor words.

*The Carriage held but just Ourselves –
And Immortality.”*

Lakoff and Turner convincingly argue that there are basic conceptual metaphors that hold for how we conceive of, and therefore talk about, death (e.g., DEATH IS WINTER, DEATH IS REST or DEATH IS FREEDOM FROM BONDAGE) and these are combined in poetry with other metaphors, such as LIFE IS A JOURNEY, A LIFETIME IS A YEAR, NIGHT IS A COVER, PEOPLE ARE PLANTS, and so on. Our goal in this work is to develop new methods of automatically suggesting words that link together concepts across semantic spaces or frames, and so to aid both programs and people in the generation of poetic and figurative language.

We add to the body of work on poetry analysis and generation by exploring a method to generate a set of words (connector words) that can be used to create a figurative relationship with a pair of anchor words, as shown in Figure 1. We do this by making use of recent advances in statistical word similarity generation methods, in particular the word2vec embedding technology.

Mikolov et al.'s work on word representations in

| Anchor word pairs | Connector words |
|----------------------|-----------------|
| surrendering & storm | barrage |
| caring & flame | cook |
| life & road | journey |

Table 1: Examples of anchor word pairs and connector words. *Life & road* connected by *journey* is an example of how the framework in Figure 1 maps to Lakoff’s LIFE IS A JOURNEY metaphor.

vector space, actualized in the word2vec technology, identifies semantic relationships between words using word vector algebra. These word vectors perform remarkably well at identifying semantic analogy relationships, e.g., capital city to country, currency to country, city to state, and man to woman (Mikolov et al., 2013a; Mikolov et al., 2013b; Mikolov et al., 2013c). The classic example showcasing the power of the word vector algebra is:

$$\begin{aligned} & \text{vector}(\textit{‘King’}) - \text{vector}(\textit{‘Man’}) + \text{vector}(\textit{‘Woman’}) \\ & = \text{vector that is closest to } \text{vector}(\textit{‘Queen’}) \end{aligned}$$

In this paper, we extend the use of these vectors beyond their primary application of identifying such analogous type relationships: we explore their use to draw together two semantic spaces for the creation of *figurative* relationships.

Specifically, starting with a pair of anchor words, made up of a concrete noun and a poetic theme, we leverage Mikolov et al.’s word vectors to return the sets of words most similar to either anchor word in the pair. By finding the intersection of these two sets, we can identify connector words that draw together the anchor words to create figurative relationships. Some examples of the anchor and connector words can be seen in Table 1.

For each pair of anchor words we also generate a list of suggested connector words by using Mikolov et al.’s word vector algebra. We then quantitatively observe the difference between the lists of connector words produced by intersection and by addition. For the connector words from the intersection list, we observe a more balanced similarity score between the connector words and each anchor word than we do for the connector words from the addition list. We then construct an initial dataset to explore qualitatively the figurative relationships generated using connector words from the addition list and those

generated using connector words from the intersection list.

The remainder of this paper provides an overview of related work on figurative language and semantic relationships; outlines the computational methods used to retrieve connector words; describes quantitative and qualitative observations of the retrieved connector words; and discusses future work.

2 Related work

2.1 Work on figurative language

“Figurative language,” in particular metaphor, plays a crucial role in poetry. Lakoff and Turner (1989) explicate in great detail how metaphors are combined in everyday language, and how poets extend and elaborate on conventional metaphors in new ways. A metaphor can be thought of as a linguistic structure that creates a “mapping” of two conceptual spaces or frames (Lakoff and Turner, 1989) or the “blending” of two input spaces (Fauconnier and Turner, 2008). Veale et al. (2000) point to this as an important theoretical model for computational work on metaphor, and emphasize the highly structured relationships that metaphors create between two terms. In one of their examples, to see *Scientists* as *Priests* metaphorically could mean to see their *lab-benches* as *altars*. In this case, the general terms are metaphorically connected based on some more specific attribute that is common to each. This commonality, however, is not immediately obvious, so a good metaphor reveals something surprising about a conceptual space by combining it with another conceptual space.

Veale et al. (2000) offer examples of systems that create such relationships between two input terms to form metaphors. More recently, researchers have generated poetic metaphorical relationships between two terms by leveraging large corpora. Veale and Hao (2007) found metaphorical relationships between two terms by mining Google search results for adjectives used to describe both terms. Veale and Hao (2008) applied a similar approach to mine more complex metaphors from WordNet. Such techniques have been deployed in “computer poetry” applications that automatically generate verse (Veale, 2013; Harmon, 2015).

Literary theorist William Empson (2004) argued

for the importance of reading the “ambiguities” present in verse. The most basic type of ambiguity occurs when a metaphor simultaneously draws on different qualities of the items brought into metaphorical relation, and so it is “effective in several ways at once.” For instance, *eyes* are like *sun* for multiple reasons (e.g., both are literally round, and both may be “bright”). More complicated ambiguities may be generated through puns, in which a word simultaneously carries two distinct and ironically opposite meanings, each relevant to the context. To use Empson’s example: in Pope’s *Dunciad*, a character sleeping “in port” may be both safe at harbor and drunk (on port wine). In this case, two distinct conceptual spaces are activated by one word.

Furthermore, as Empson argues, poets themselves are not always fully in control of the meanings of their words and, “discovering his idea in the act of writing” may create “a simile which applies to nothing exactly, but lies half-way between two things,” a subtly mixed metaphor. He gives the example of a passage from a verse-play by John Ford in which the term “gall” at first seems to mean “boldness”—but, when the author later mentions a “well-grown oak,” comes to retroactively signal “oak-galls,” a horticultural disease. “Figurative language” is not merely the sort of ordered and symmetrical matching of cognitive structures evinced by clear metaphors; it is also what happens when poets get caught up in loose and chaotic association between words, the way jazz musicians zig and zag between notes.

2.2 Semantic relationships from word embeddings

Distributional approaches to representing word meaning have a long history in computational linguistics, and are motivated by the notion that “You shall know a word by the company it keeps!” (Firth, 1957) and the Wittgenstein-inspired notion that a concept is not an isolated thing but really a constellation of concepts linked by family resemblances (Rosch and Mervis, 1975).

Schütze (1993) made early attempts to represent the meaning of concepts by creating n-dimensional word spaces. More recent attempts which use larger collections and innovations in algorithms have yielded more accurate representations of semantic relatedness. These include Mikolov et al.’s work on

word embeddings learned using a Continuous Skip-gram Model. It extends beyond bag-of-words models by accounting for the context a word appears in. This model is implemented in Google’s word2vec tool¹ which has shown significant success in finding both syntactic and semantic relationships between words (Mikolov et al., 2013b; Mikolov et al., 2013a; Mikolov et al., 2013c).

Mikolov et al. find semantic relationships from word pairs which can be categorized into 79 specific word relations, such as *Cause:Effect* or *Action:Goal* as identified in the SemEval-2012 Task 2, *Measuring Relation Similarity* (Jurgens et al., 2012; Turney, 2012). For example, the word pair *clothing:shirt* falls into the *Class Inclusion:Singular Collective* relation. Mikolov et al. (2013c) then use a vector offset technique to understand the validity of resulting analogous relationships, such as “*clothing* is to *shirt* as *dish* is to *bowl*” as tested against the word relation data set presented by Jurgens et al. (2012). Similarly, our work aims to identify a semantic relationship between two words. It differs in that we aim to find figurative relationships between two words for poetic purposes, as opposed to analogous relationships between two pairs of words.

3 Computational methods in word2vec to retrieve connector words

In this section we outline two methods – an addition model and an intersection model – to retrieve connector words that support figurative relationships using word2vec.

The standard functionality of word2vec is to retrieve the top-ranked most similar words. Word2vec addition is optimized for such tasks, but here we are more interested in retrieving words to support figurative relationships. We aim to find the words in the overlap of the family resemblance spaces of each of the anchor words.

We do so by retrieving words from word2vec similarity lists that are common to each anchor word. In contrast to the addition model, the intersection model retrieves words from further down on the sim-

¹Throughout this paper, we are using the Gensim implementation of word2vec (Řehůřek and Sojka, 2010), trained on ‘pruned.word2vec.txt’.

| Concrete nouns | | |
|----------------|--------|----------|
| bed | ear | finger |
| horse | sand | hair |
| bell | grass | rock |
| book | rose | breast |
| ship | blood | window |
| wing | girl | snow |
| wood | ring | body |
| room | wine | ground |
| mouth | garden | stone |
| storm | brain | flame |
| town | wave | shadow |
| silver | mist | line |
| stream | dawn | path |
| dust | breath | king |
| color | spring | darkness |
| side | nation | race |
| state | | |

Table 2: Pool of concrete nouns used in the selection of anchor pairs.

| Poetic themes | | |
|---------------|------------|--------------|
| loss | melancholy | anger |
| animals | calmness | compassion |
| confusion | death | envy |
| faith | fear | forgiveness |
| freedom | friendship | god |
| grace | gratitude | grief |
| hate | hope | immortality |
| jealousy | joy | life |
| mothers | nature | peace |
| people | religion | remembrance |
| love | sadness | silence |
| smiling | songs | spirituality |
| spring | suffering | truth |
| unity | vanity | war |
| water | wind | bitterness |
| consciousness | happiness | earth |
| soul | surrender | violence |

Table 3: Pool of poetic themes used in the selection of anchor pairs.

ilarity lists for each anchor word (moving towards the outer edges of their respective family resemblance spaces). The resulting words in the shared space maintain a balance between the two anchor words, thus drawing them together.

To narrow the scope, the work in this paper focuses on anchor word pairs comprising one concrete noun and one poetic theme. We chose this focus from our observations that poetry often relies on a connection between a concrete concept and a more abstract theme, which is consistent with Kao and Jurafsky’s (2015) findings that professional poetry contains more concreteness.

3.1 Selecting anchor word pairs

For the investigation here, we randomly generate anchor word pairs from a list of concrete nouns (see Table 2) and a list of poetic themes² (see Table 3).³

²Created from a list of poetic themes from <http://www.poetseers.org/themes/> then expanded to include the top 5 most similar words, using word2vec. The expanded list was normalized to lower-case words. Overly-specific words, including “rainbows”, “cats,” “pets,” “rabbits,” “dogs,” “Iraq,” and “sewage”, were removed.

³These particular lists were chosen for expediency; the rigorous definition of concrete nouns and poetic themes is not central to our exploration.

The set of concrete nouns comes from existing poetry.⁴ The words are selected based on their word frequency across the corpus, number of noun senses in WordNet (WordNet, 2010), and degree of concreteness using the word concreteness dataset developed by Brysbaert et al. (2013). The frequency measure is normalized across the corpus. The mean concreteness ratings from Brysbaert et al. (2013) range from 0.0 to 5.0 and include standard deviations. The concrete noun list is composed of nouns with a normalized frequency ranging from 0.0 to 0.1; the number of noun senses ranging from 0 to 4; and the degree of concreteness ranging from 3.5 to 5.0, with the degree of concreteness standard deviation ranging from 0.0 to 3.0.

In generating anchor word pairs, we randomly select a concrete noun and a poetic theme pair where the two words occupy distinct semantic spaces. For the purposes of this paper, we use a cosine similarity score of less than 0.4 as a threshold. The similarity scores of candidate concrete nouns to candidate poetic themes range from -0.15 (dissimilar) to 0.45

⁴Existing poetry from a corpus of 2,860 poems downloaded from the “19th Century American Poetry” section of <http://famouspoetsandpoems.com>.

| Top 10 words from word2vec addition for storm + surrendering |
|---|
| surrendered |
| hurricane |
| storms |
| snowstorm |
| rainstorm |
| tornado |
| blizzard |
| typhoon |
| twister |
| squall |

Table 4: Top 10 words retrieved when adding anchor words *storm* and *surrendering* using word2vec addition.

(similar). This similarity check is used to create an anchor word pair comprising two words with different semantic spaces. If the two anchor words are too similar, they will rely on a synonymous connection and thus will not provide two distinct semantic spaces to blend.

3.2 Addition model in word2vec to retrieve connector words

The addition model retrieves a set of connector words which are the most similar to a pair of anchor words using word2vec’s existing vector addition approach.

Implementation of this addition model involves starting with word2vec’s word vector representations of the concrete noun, \vec{c} , and the poetic theme, \vec{t} , of the anchor word pair.

The word vector, \vec{a} , is then defined such that $\vec{a} = \vec{c} + \vec{t}$. Word2vec then searches for the word vectors with the greatest cosine similarity to \vec{a} , which approximates their similarity (Mikolov et al., 2013c). We use this vector addition to find a set A containing n words closest to \vec{a} . For example, if we take the concrete noun “storm” and the poetic theme “surrendering” as an anchor word pair, we can retrieve the list of words in Table 4.⁵

The resulting list contains primarily words that are synonyms to one of the two anchor words. However, our goal is to retrieve connector words

⁵Proper nouns were removed from the list and morphological duplicates removed.

that blend the semantic spaces of the two anchor words, so we investigate an alternative computation in word2vec, the intersection model.

3.3 Intersection model in word2vec to retrieve connector words

For the intersection model, we start with word2vec’s vector representations of the concrete noun, \vec{c} , and the poetic theme, \vec{t} , of the anchor word pair.

Using word2vec, we then find a set, C , which contains the top $n = 1000$ word vectors that have the greatest cosine similarity to \vec{c} . Similarly, we find a set, T , which contains the top $n = 1000$ word vectors that have the greatest cosine similarity to \vec{t} .

Looking at the intersection, I , of the two sets $I = C \cap T$, we find words that relate both to the initial concrete noun, (\vec{c}), and the poetic theme, (\vec{t}).

The resulting set, I , varies significantly in size depending on the concrete noun and poetic theme pair chosen. The depth n also contributes to the size of the set, I . In our analyses, we elected to use $n = 1000$ because it elicited plentiful yet meaningful results.

It is the case that for any sets I and A of similar size, there are likely to be words unique to each set, as well as words that are shared between these sets. A proof of this appears in Appendix A.

Since A and I have overlapping words, but also contain unique words, we remove the overlapping words to focus on the resulting set $U_I = I \setminus A$, containing the words unique to the intersection set, and $U_A = A \setminus I$, containing the words unique to the addition set. We focus on these unique words sets to facilitate our observations of the differences between the two models. Quantitatively, we observe differences in the range of similarity scores between the anchor word-connector word pairs. Qualitatively, we use the unique words from each set to consider the potential to support figurative language by combining the semantic spaces of the two anchor words.

With the example anchor words, *storm* and *surrendering*, we see the resulting unique word lists in Table 5. In the next two sections we quantitatively and qualitatively observe these two lists.

| Unique to Intersection $U_I = I \setminus A$ | Unique to Addition $U_A = A \setminus I$ |
|---|---|
| onslaught | squall |
| stranding | tornado |
| blowing | typhoon |
| dissipating | snowstorm |
| battering | flooding |
| game | rainstorm |
| breastworks | deluge |
| regrouped | downpour |
| batter | blizzard |
| dissipated | ike |
| outburst | twister |
| pounding | hurricane |
| submerging | rain |
| pounded | |
| barrage | |
| regrouping | |
| stalemate | |

Table 5: Connector words for *storm* and *surrendering* retrieved from the words unique to *I* and the words unique to *A*.

4 Quantitative observations of retrieved connector words

In observing the cosine similarities between the words in U_I and each anchor word, and the words in U_A and each anchor word, we begin to see a pattern where the words from U_I fall within a smaller band of similarity than of those in U_A (see Tables 6 and 7 for example with words from Table 5). Note that the highest possible cosine similarity score is 1, indicating maximum similarity, and the lowest is -1, indicating dissimilarity.

Across 10 different randomly selected anchor word pairs, we see that this same pattern holds. The words in U_I fall within a band of similarity ranging from approximately 0.25 to approximately 0.30 where the average spread between the two similarities is 0.06. By comparison, the similarities between connector words in U_A and each anchor word falls within a larger band of similarity ranging from approximately 0.1 to approximately 0.6 where the average spread between the two similarities is 0.44. Table 8 shows these ranges for each anchor word pair. The connector words in U_I are more balanced between both of the anchor words, whereas the con-

| Unique to Addition $U_I = I \setminus A$ | Similarity to noun <i>storm</i> | Similarity to theme <i>surrendering</i> |
|--|------------------------------------|--|
| onslaught | 0.30 | 0.20 |
| stranding | 0.27 | 0.28 |
| blowing | 0.24 | 0.29 |
| dissipating | 0.23 | 0.22 |
| battering | 0.29 | 0.24 |
| game | 0.19 | 0.25 |
| breastworks | 0.19 | 0.20 |
| regrouped | 0.19 | 0.31 |
| batter | 0.22 | 0.25 |
| dissipated | 0.24 | 0.21 |
| outburst | 0.21 | 0.20 |
| pounding | 0.20 | 0.26 |
| submerging | 0.26 | 0.23 |
| pounded | 0.24 | 0.32 |
| barrage | 0.25 | 0.20 |
| regrouping | 0.19 | 0.31 |
| stalemate | 0.19 | 0.21 |
| Average spread between similarity scores: 0.05 | | |

Table 6: Similarity scores between connector words found in U_I to anchor words *storm* and *surrendering*. The average spread between the scores of 0.05 indicates the small band of similarity the words exist in, showing the balanced similarity the connector word has with each of the anchor words.

connector words in U_A are more closely related to a single anchor word.

5 Qualitative observations of retrieved connector words

Next we qualitatively explore the potential of each model to retrieve words in the shared space between two anchor words using a crowd-sourced dataset of figurative relationships. We annotate these relationships based on the types of connections made between the connector words and anchor word pairs.

5.1 Dataset construction

We construct a dataset made up of sentences stating the figurative relationships tying the connector words from the addition and intersection lists to pairs of anchor words. This dataset allows us to explore the potential of the connector words pro-

| Unique to Intersection $U_A = I \setminus A$ | Similarity to noun <i>storm</i> | Similarity to theme <i>surrendering</i> |
|---|------------------------------------|--|
| squall | 0.63 | -0.03 |
| tornado | 0.64 | -0.02 |
| typhoon | 0.62 | -0.01 |
| snowstorm | 0.64 | 0.01 |
| flooding | 0.57 | 0.01 |
| rainstorm | 0.57 | 0.07 |
| deluge | 0.50 | 0.08 |
| downpour | 0.52 | 0.08 |
| blizzard | 0.61 | 0.00 |
| ike | 0.58 | 0.02 |
| twister | 0.62 | -0.01 |
| hurricane | 0.73 | 0.04 |
| rain | 0.46 | 0.10 |
| Average spread between similarity scores: 0.56 | | |

Table 7: Similarity scores between connector words found in U_A to anchor words *storm* and *surrendering*. The average spread between the scores of 0.56 shows the wide range of similarity scores.

vided by each approach to blend the distinct semantic spaces of the two anchor words. To generate this dataset, we presented crowd-sourced workers from Mechanical Turk with a list of words, either those unique to the I set (U_I) or those unique to the A set (U_A). The words in the provided sets were normalized to exclude proper nouns, lower-case all characters, and eliminate morphological duplicates. If the unique word list exceeded 10 words, a random sample of 10 words was shown.

The Mechanical Turk workers then selected a single connector word from the list and wrote a sentence to describe the relationship between the anchor words and the connector word. Mechanical Turk workers were provided the diagram in Figure 1 with the concrete noun and poetic theme words populated. We informed workers that they should select the connector word that “best connects the anchor words in a poetic sense (e.g., using a double meaning, creating a new image, creating an interesting relationship, etc.)”.

The workers were prompted to fill in text to complete a template sentence of the form: “[connector word] connects [concrete noun] and [poetic theme]

| Anchor word pairs | Range of avg. sim. from words in U_I to anchor words | Range of avg. sim. from words in U_A to anchor words |
|-------------------|--|--|
| flame & caring | 0.22 – 0.30 | 0.13 – 0.58 |
| color & earthly | 0.28 – 0.32 | 0.17 – 0.55 |
| hair & anguish | 0.27 – 0.33 | 0.14 – 0.66 |
| flame & killing | 0.23 – 0.26 | 0.09 – 0.55 |
| mouth & comp. | 0.25 – 0.29 | 0.16 – 0.54 |
| storm & surr. | 0.21 – 0.26 | 0.03 – 0.69 |
| ring & mankind | 0.21 – 0.34 | 0.11 – 0.57 |
| hair & envied | 0.27 – 0.31 | 0.17 – 0.58 |
| book & liberties | 0.23 – 0.29 | 0.15 – 0.54 |
| town & grieving | 0.24 – 0.28 | 0.14 – 0.54 |

Table 8: The low end of the ranges is the average of the minimum similarity scores across all the connector words to each of the words in the anchor word pair. The upper end of the ranges is the average of the maximums. A smaller range means that the anchor words have more balanced similarity to the connector word. comp. = compassionate; surr. = surrendering.

because...”. For example, “Barrage connects storm and surrendering because...”.

With 25 workers writing 4 sentences each across 10 anchor word pairs, the constructed dataset contained 100 generated sentences.

The following sections provide examples of the figurative relationships among the selected connector word and anchor word pairs created by Mechanical Turk workers and discussion of whether the relationships created achieve heightened effects by drawing together two distinct semantic spaces.

5.2 Sample data

Using the dataset of generated sentences, we explore the potential for word2vec to provide connector words that blend the two distinct semantic spaces of the two anchor words using the addition and intersection operations.

Below, we present detailed results for 3 of the assessed anchor word pairs. We show which words were chosen by Mechanical Turk workers as the best connector word (bolded), and the sentences describing the relationship among the connector word

| Unique to Intersection | Unique to Addition |
|------------------------|--------------------|
| radiant | celestial |
| reflections | hues |
| clearness | uncolored |
| unclouded | metaphysical |
| blackness | astral |
| loveliness | cosmic |
| refracted | translucence |
| eidetic | divine |
| allusiveness | heavenly |
| creeds | |
| geometrical | |
| exquisiteness | |
| diaphanous | |

Table 9: Figurative ties between *color* and *earthly*. Bolded words were selected by Mechanical Turk workers as the best word to create the figurative tie.

and the anchor words (underlined) as created by the workers.

5.2.1 Sample 1: color and earthly

All connector words chosen are shown in Table 9. Sample connection descriptions from Mechanical Turk workers as follows:

“**Radiant** connects color and earthly because radiant means a bright color that looks like it’s shining and at night, the earthly sky is radiant because it shines brightly with the stars.”

“**Hues** connects color and earthly because hues imply various colors, shades, or characteristics and hues can be earthly in tone, such as blues, greens and browns.”

5.2.2 Sample 2: storm and surrendering

All connector words chosen are shown in Table 10. Sample connection descriptions from Mechanical Turk workers are as follows:

“**Barrage** connects storm and surrendering because a storm is a barrage of bad weather like winds and rain people surrender when they feel a barrage of overwhelming things coming at them.”

“**Hurricane** connects storm and surrendering because it is a type of storm and those who surrender to it are spared, like grass and those who stand against it are devastated, like big trees.”

| Unique to Intersection | Unique to Addition |
|------------------------|--------------------|
| onslaught | squall |
| stranding | tornado |
| blowing | typhoon |
| dissipating | snowstorm |
| battering | flooding |
| game | rainstorm |
| breastworks | deluge |
| regrouped | downpour |
| batter | blizzard |
| dissipated | ike |
| outburst | twister |
| pounding | hurricane |
| submerging | rain |
| pounded | |
| barrage | |
| regrouping | |
| stalemate | |

Table 10: Figurative ties between *storm* and *surrendering*. Bolded words were selected by Mechanical Turk workers as the best word to create the figurative tie.

5.2.3 Sample 3: flame and caring

All connector words chosen are shown in Table 11. Sample connection descriptions from Mechanical Turk workers are as follows:

“**Cook** connects caring and flame because it is related to flame as flames are used in cooking and cooking can be a symbol of caring for someone with good food.”

“**Torch** connects caring and flame because when someone cares about someone else it’s often said they are carrying a torch for them, while the visual of a torch itself tends to have a flame atop it.”

5.3 Discussion of qualitative observations

As stated above, our goal in suggesting connector words is to blend the distinct semantic spaces of the two anchor words to create figurative relationships.

While the sentences with synonyms do contain figurative language, they do not achieve our goal of using the connector word to blend the two anchor words. Instead, the connector word shares a

| Unique to Intersection | Unique to Addition |
|------------------------|--------------------|
| affection | compassion |
| friendship | torch |
| spirit | selfless |
| passion | considerate |
| soul | kindness |
| brotherhood | compassionate |
| aloneness | loving |
| love | devotion |
| cook | |
| undying | |

Table 11: Figurative ties between *flame* and *caring*. Bolded words were selected by Mechanical Turk workers as the best word to create the figurative tie.

semantic space with one of the anchor words and this shared semantic space is then blended with the semantic space of the other anchor word. The sentences reflect a figurative relationship that is present between the two anchor words (which does not depend on the connector word) rather than a new space created by the introduction of the connector word.

We observe that the connector words that create a heightened effect by blending the two anchor words have a balanced cosine similarity to both anchor words (in the range of approximately 0.25 to approximately 0.30 as discussed in section 4 and shown in Table 8). This means that the connector word is not closer to one or the other anchor word, but rather occupies the shared space between the two anchor words. In contrast, the connector words that are synonymous with one of the anchor words, and thus do not blend the semantic spaces of the two anchor words, have imbalanced cosine similarities to the two anchor words. The connector word’s shared semantic space with one of the anchor words is visible in a higher cosine similarity to that anchor word (approximately 0.6) and a much lower cosine similarity to the other anchor word (approximately 0.1). In this latter case, the connector word is not blending the semantic spaces of the two anchor words but is rather sharing the semantic space of one.

5.3.1 Relationships based on synonyms

In the sentences that rely on synonyms for one part of the relationship, the connector word has a metaphorical relationship with one of the anchor

words and a non-metaphorical with the other anchor word.

By looking at word2vec similarity scores of concrete noun to connector word and poetic theme to connector word, we can see that in the relationships that rely on synonym there is a relatively wide spread between the similarity scores.

The examples below show figurative relationships that rely on synonym-based relationships between the connector word and one of the anchor words along with the similarity scores between the connector word and the concrete noun and between the connector word and the poetic theme.

Examples:

“**Torch** connects *caring* and *flame* because caring for someone can feel like a flame or a torch burns inside you for them.”

Torch and flame are connected through a synonymous relationship; these two words are then connected to caring through a metaphor (caring is a torch burning)

Similarity score **torch**-*caring*: 0.06

Similarity score **torch**-*flame*: 0.67

“**Hues** connects *color* and *earthly* because hues imply various colors, shades, or characteristics and hues can be earthly in tone, such as blues, greens and browns.”

Hues and color are connected through a synonymous relationship; these two words are then connected to earthly through a metaphor (colors are earthly).

Similarity score **hues**-*color*: 0.61

Similarity score **hues**-*earthly*: 0.09

5.3.2 Relationships blending distinct semantic spaces

The figurative relationships that result in a heightened effect are created through a connector word retrieved from the overlapping semantic space between the two anchor words.

In these relationships, the word2vec similarity scores of concrete noun to connector word and poetic theme to connector word are close, indicating a balanced relationship.

The examples below show figurative relationships that use a connector word that blends the two distinct semantic spaces of the anchor words.

Examples:

“**Barrage** connects storm and surrendering because a storm is a barrage of bad weather like winds and rain people surrender when they feel a barrage of overwhelming things coming at them.”

A storm is a barrage of bad weather and life can be a barrage to which you surrender.

Similarity score **barrage-storm**: 0.25

Similarity score **barrage-surrendering**: 0.20

“**Cook** connects caring and flame because it is related to flame as flames are used in cooking and cooking can be a symbol of caring for someone with good food.”

Providing nourishment by cooking requires flames and is caring.

Similarity score **cook-caring**: 0.26

Similarity score **cook-flame**: 0.22

6 Summary of observations

In the dataset of sentences generated by Mechanical Turk workers drawing the connections between anchor words and connector words, we observe the following:

- Instances where the cosine similarity scores between each of the anchor words and the connector word are unbalanced tend to lead to a synonymous relationship between one anchor word-connector word pair (a nonmetaphorical relation) and a shared figurative relationship with the second anchor word (a metaphorical relation). In these cases the connector word is not drawing together the family resemblance semantic spaces of the two anchor words, because it already exists in the semantic space of one of them.
- Instances where the cosine similarity scores between each of the anchor words and the connector word are balanced tend to lead to a heightened effect relationship blending the two distinct semantic spaces of the anchor word.

As seen in Table 8, the band of similarity scores resulting from words in U_I is smaller than the band of similarity resulting from connector words in U_A , suggesting a more balanced relationship.

7 Future work

We have observed figurative relationships resulting from the introduction of a connector word to an anchor word pair. We notice that balanced cosine similarity scores between the connector word and each anchor word tend to lead to heightened effects by blending the two distinct semantic spaces of the anchor word.

The words unique to the intersection list proposed here have balanced cosine similarity scores ranging from approximately 0.25 to 0.30, suggesting that finding the words unique to the intersection list prioritizes the retrieval of words that blend the distinct semantic spaces of two anchor words.

The next step in this work is to test this hypothesis with an evaluation. Such an evaluation may include looking at the band of similarity from 0.25 to 0.30 directly, by way of the unique words to intersection and unique words to addition sets, and/or by way of the complete intersection and complete addition sets. We could also conduct threshold testing for varying word2vec settings and top n settings. In evaluating this work, it would be interesting to see if everyday people and practicing poets judge the relationships differently.

Additionally, through further evaluation, the nature of other bands of similarity outside of the 0.25 to 0.30 range could be tested, as well as the presence of such a band of similarity when expanding beyond the concrete noun-poetic theme scope.

Expanding beyond the concrete noun-poetic theme scope could also involve grounding the anchor pair selection more explicitly in the metaphors proposed by Lakoff and Turner (1989).

Further related work may include consideration of more computations within word2vec to see what types of word relations such computations support.

Once a conceptual understanding is more established, research could then be conducted regarding the various applications that such findings could be used for. Such applications may include poetry generation or tools to assist in creative writing.

Overall, we hope that this work will continue to promote the development of computational approaches to figurative language, because:

“By metaphor you paint
A thing.”

- Wallace Stevens *Poem Written At Morning*

Acknowledgements

We thank David Bamman for his guidance and valuable insights and the reviewers for their very thoughtful and thorough feedback.

A Appendix: Proof

Formally, the set A contains the top n most similar word vectors, \vec{w} , such that $\cos(\vec{w}, \vec{a}) \geq \alpha$, where α is a minimum similarity threshold resulting from selecting the top n words. As such:

$$\begin{aligned} \vec{w} \in A, s.t. : \\ \cos(\vec{w}, \vec{a}) = \cos(\vec{w}, (\vec{c} + \vec{t})) = \frac{\vec{w} \cdot (\vec{c} + \vec{t})}{|\vec{w}| |\vec{c} + \vec{t}|} \geq \alpha \end{aligned} \quad (1)$$

The set I contains all word vectors, \vec{w} , such that $\cos(\vec{w}, \vec{c}) \geq \beta$ and $\cos(\vec{w}, \vec{t}) \geq \gamma$, where β and γ are minimum similarity thresholds resulting from selecting the top n words from each list.

$$\begin{aligned} \vec{w} \in I, s.t. : \\ \cos(\vec{w}, \vec{c}) = \frac{\vec{w} \cdot \vec{c}}{|\vec{w}| |\vec{c}|} \geq \beta \\ \cos(\vec{w}, \vec{t}) = \frac{\vec{w} \cdot \vec{t}}{|\vec{w}| |\vec{t}|} \geq \gamma \end{aligned} \quad (2)$$

If we were finding the single word vector that maximized (1) and (2), the two equations would be equivalent, as shown by Levy and Goldberg (2014). Rather, in the addition model, we are finding the set of words that satisfy (1), and, in the intersection model, we are finding the set of words that satisfy (2). We can see that (1) and (2) are not necessarily

equivalent. If they were, we would have a connector word, \vec{w} , such that (1) and (2) were always both satisfied. As such, we would need to satisfy (3):

$$\begin{aligned} \cos(\vec{w}, (\vec{c} + \vec{t})) &\geq \alpha \\ \cos(\vec{w}, \vec{c}) &\geq \beta \\ \cos(\vec{w}, \vec{t}) &\geq \gamma \end{aligned} \quad (3)$$

Note that (3) assumes the word vectors are length-normalized. We then expand (3) as follows:

$$\begin{aligned} w_1 * (c_1 + t_1) + w_2 * (c_2 + t_2) \\ + \dots + w_n * (c_n + t_n) &\geq \alpha \\ w_1 * c_1 + w_2 * c_2 + \\ + \dots + w_n * c_n &\geq \beta \\ w_1 * t_1 + w_2 * t_2 \\ + \dots + w_n * t_n &\geq \gamma \end{aligned} \quad (4)$$

We can solve (4) as follows:

$$\beta + \gamma \geq \alpha \quad (5)$$

(5) is not necessarily always true. Thus, the initial assumption that the addition and intersection models contain the same word vectors is contradicted, which confirms that A does not necessarily equal I .

References

- [Brysbaert et al.2013] Marc Brysbaert, Amy Beth Wariner, and Victor Kuperman. 2013. Concreteness ratings for 40 thousand generally known english words and lemmas. In *Behavior Research Methods*, pages 1–8.
- [Empson2004] William Empson. 2004. *Seven types of ambiguity*, volume 645. Random House.
- [Fauconnier and Turner2008] Gilles Fauconnier and Mark Turner. 2008. *The way we think: Conceptual blending and the mind's hidden complexities*. Basic Books.
- [Firth1957] John Rupert Firth. 1957. A synopsis of linguistic theory, 1930-1955. *Studies in linguistic analysis (Special volume of the Philological Society)*, pages 1–31. Reprinted in: Frank R. Palmer (ed.) *Selected papers of J. R. Firth 1952-59*, Longmans, Green and Co Ltd, London and Harlow, UK, 168-205; citation on page 179.
- [Harmon2015] Sarah Harmon. 2015. Figure8: A novel system for generating and evaluating figurative language. In *Proceedings of the Sixth International Conference on Computational Creativity June*, page 71.
- [Jurgens et al.2012] David Jurgens, Saif Mohammad, Peter Turney, and Keith Holyoak. 2012. Semeval-2012 task 2: Measuring degrees of relational similarity. In *In *SEM 2012: The First Joint Conference on Lexical and Computational Semantics (SemEval 2012)*, page 356364. Association for Computational Linguistics.
- [Kao and Jurafsky2015] Justine Kao and Dan Jurafsky. 2015. A computational analysis of poetic style: Imagism and its influence on modern professional and amateur poetry. *Linguistic Issues in Language Technology*, 12(3).
- [Lakoff and Turner1989] George Lakoff and Mark Turner. 1989. *More than cool reason: A field guide to poetic metaphor*. University of Chicago Press.
- [Levy and Goldberg2014] Omer Levy and Yoav Goldberg. 2014. Linguistic regularities in sparse and explicit word representations. In *CoNLL*.
- [Mikolov et al.2013a] Tomas Mikolov, Kai Chen, Greg Corrado, Jeffrey Dean, and Dan Jurafsky. 2013a. Efficient estimation of word representations in vector space. In *arXiv preprint arXiv:1301.3781*.
- [Mikolov et al.2013b] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *arXiv:1310.4546 [cs.CL]*.
- [Mikolov et al.2013c] Tomas Mikolov, Wen-tau Yih, Greg Corrado, and Jeffrey Dean. 2013c. Linguistic regularities in continuous space word representations. In *HLT-NAACL*.
- [Řehůřek and Sojka2010] Radim Řehůřek and Petr Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta, Malta, May. ELRA.
- [Rosch and Mervis1975] Eleanor Rosch and Carolyn B Mervis. 1975. Family resemblances: Studies in the internal structure of categories. *Cognitive psychology*, 7(4):573–605.
- [Schütze1993] Hinrich Schütze. 1993. Word space. In *Advances in Neural Information Processing Systems* 5.
- [Turney2012] Peter Turney. 2012. Domain and function: A dual-space model of semantic relations and compositions. In *Journal of Artificial Intelligence Research*, pages 533–585.
- [Veale and Hao2007] Tony Veale and Yanfen Hao. 2007. Comprehending and generating apt metaphors: a web-driven, case-based approach to figurative language. In *AAAI*, volume 2007, pages 1471–1476.
- [Veale and Hao2008] Tony Veale and Yanfen Hao. 2008. A fluid knowledge representation for understanding and generating creative metaphors. In *Proceedings of the 22nd International Conference on Computational Linguistics-Volume 1*, pages 945–952. Association for Computational Linguistics.
- [Veale et al.2000] Tony Veale, Diarmuid O Donoghue, and Mark T Keane. 2000. Computation and blending. *Cognitive Linguistics*, 11(3/4):253–282.
- [Veale2013] Tony Veale. 2013. Less rhyme, more reason: Knowledge-based poetry generation with feeling, insight and wit. In *Proceedings of the International Conference on Computational Creativity*, pages 152–159.
- [WordNet2010] WordNet. 2010. Princeton university. <http://wordnet.princeton.edu>.