

# IMPLI : Investigating NLI Models’ Performance on Figurative Language

Anonymous ACL submission

## Abstract

Natural language inference (NLI) has been widely used as a task to train and evaluate models for language understanding. However, the ability of NLI models to perform inferences requiring understanding of figurative language such as idioms and metaphors remains understudied. We introduce the **IMPLI** (**I**diomatic and **M**etaphoric **P**aired **L**anguage **I**nference) dataset, an English dataset consisting of paired sentences spanning idioms and metaphors. We develop novel methods to generate 24k semi-automatic pairs as well as manually creating 1.8k gold pairs. We use **IMPLI** to evaluate NLI models based on RoBERTa fine-tuned on the MNLI dataset, and show that while they can reliably detect entailment relationship between figurative phrases with their literal counterparts, they perform poorly on examples where pairs are designed to be non-entailing. This suggests the limits of current NLI models with regard to understanding figurative language and this dataset serves as a benchmark for future improvements in this direction.<sup>1</sup>

## 1 Introduction

Understanding figurative language (i.e., that in which the intended meaning of the utterance differs from the literal compositional meaning) is a particularly difficult area in NLP (Shutova, 2011; Veale et al., 2016), but is essential for proper natural language understanding. We consider here two types of figurative language: idioms and metaphors. Idioms can be viewed as non-compositional multiword expressions (Jochim et al., 2018), and have been historically difficult for NLP systems. For instance, sentiment systems struggle with multiword expressions in which individual words do not directly contribute to the sentiment (Sag et al., 2002). Metaphors involve linking conceptual properties

<sup>1</sup>Dataset and all related resources are included in the supplementary materials.

Idioms	Jamie was <i>pissed off</i> this afternoon. → Jamie was <i>irritated this afternoon</i>
	There’s a marina down <i>in the docks</i> . ↯ There’s a marina down <i>under scrutiny</i> .
Metaphors	The <i>hearts of men</i> were softened. → The <i>men</i> were made kinder and gentler.
	The gun <i>kicked</i> into my shoulder. ↯ The <i>mule</i> kicked into my shoulder.

Table 1: Examples of entailment (→) and non-entailment pairs (↯) from the **IMPLI** dataset.

of two or more domains, and are known to be pervasive in everyday language (Lakoff and Johnson, 1980; Stefanowitsch and Gries, 2008; Steen et al., 2010). Recent work has shown that these types of figurative language are impactful across a broad array of NLP tasks (see §2.1).

Deep pretraining and transformer-based architectures have yielded increasingly powerful language models (Vaswani et al., 2017; Devlin et al., 2019; Liu et al., 2019). However, relatively little work has explored these models’ representations of figurative and creative language. NLI datasets have widely been used for evaluating the performance of language models (Dagan et al., 2006; Bowman et al., 2015a; Williams et al., 2018; Nie et al., 2020), but there are insufficient figurative language datasets in which a literal sentence is linked to a corresponding figurative counterpart that are large enough to be suitable for evaluating NLI. Due to the creative nature of human language, creating a dataset of diverse, high-quality literal/figurative pairs is time-consuming and difficult.

To address this gap, we build a new English dataset of paired expressions designed to be leveraged to explore model performance via NLI. Our dataset, **IMPLI** (**I**diomatic/**M**etaphoric **P**aired **L**anguage **I**nference), is comprised of both silver pairs, which are built using semi-automated methods (§3.1), as well as hand-written gold pairs

(§3.4), crafted to reflect both entailment and non-entailment scenarios. Each pair consists of a sentence containing a figurative expression (idioms/metaphors) and a literal counterpart, designed to be either entailed or non-entailed by the figurative expression (Table 1 shows some examples).

Our contribution thus consists of three key parts:

- We create a new **IMPLI** dataset consisting of 24,029 silver and 1,831 gold sentence pairs consisting of idiomatic and metaphoric phrases that result in both entailment and non-entailment relationship (see Table 2).
- We evaluate language models in an NLI setup, showing that metaphoric language is surprisingly easy, while non-entailing idiomatic relationships remain extremely difficult.
- We evaluate model performance in a number of experiments, showing that incorporating idiomatic expressions into the training data is less helpful than expected, and that idioms that can occur more in more flexible syntactic contexts tend to be easier to classify.

## 2 Background

### 2.1 Figurative Language and NLP

Figurative language includes idioms, metaphors, metonymy, hyperbole, and more. Critically, figurative language is that in which speaker meaning (what the speaker intends to accomplish through an utterance) differs from the literal meaning of that utterance. This leads to problems in NLP systems if they are trained mostly on literal data, as their representations for particular words and/or phrases will not reflect their figurative intended meanings.

Figurative language has a significant impact on many NLP tasks. Metaphoric understanding has been shown to be necessary for proper machine translation (Mao et al., 2018; Mohammad et al., 2016). Sentiment analysis also relies critically on figurative language: irony and sarcasm can reverse the polarity of a sentence, while metaphors and idioms may make more subtle changes in the speaker meaning (Ghosh et al., 2015). Political discourse tasks including bias, misinformation, and political framing detection benefit from joint learning with metaphoricality (Huguet Cabot et al., 2020). Figurative language engendered by creativity on social media also poses difficulty for many NLP tasks including identifying depression symptoms (Yadav et al., 2020; Iyer et al., 2019) and hate speech detection (Lemmens et al., 2021).

We are here focused on idioms and metaphors. There is currently a gap in diagnostic datasets for idioms, and our work fills this gap. There exist some relevant metaphoric resources (see §2.2); metaphors are known to be extremely common and important to understanding figurative language, our resource serves to build upon this work.

### 2.2 NLI and related challenges

Natural language inference is the task of predicting, given two fragments of text, whether the meaning of one (*premise*) entails the other (*hypothesis*) (Dagan et al., 2006). The task is formulated as a 3-way classification problem, in which the premise and hypothesis pairs are labeled as *entailment*, *contradiction*, or *neutral*, if their relationship could not be directly inferred (Bowman et al., 2015b). NLI has been widely used as an evaluation task for language understanding, and there have been a large number of challenging datasets, which have been used to further our understanding of the capabilities of language models (Wang et al., 2018, 2019).

Paired data for figurative language is relatively sparse, and there is a gap in the diagnostic datasets used for NLI in this area. Previous work includes the literal/metaphoric paraphrases of Mohammad et al. (2016) and Bizzoni and Lappin (2018), although both contain only hundreds of samples, insufficient for proper model training and evaluation. With regard to NLI, early work proposed the task of textual entailment as a way of understanding metaphor processing capabilities (Agerri et al., 2008; Agerri, 2008). Poliak et al. (2018) build a dataset for diverse NLI, which includes some creative language such as puns, albeit making no claims with regard to figurativeness.

Zhou et al. (2021) build a dataset consisting of paired idiomatic and literal expressions. They begin with a set of 823 idiomatic expressions yielding 5,170 sentences, and had annotators manually rewrite sentences containing these idioms as literal expressions. We expand on this methodology by having annotators only correct definitions for the idioms themselves, and use these definitions to automatically generate the literal interpretations of the idioms by replacing them into appropriate contexts: this allows us to scale up to over 24k silver sentences. We also expand beyond paraphrasing by incorporating both entailment and non-entailment pairs to enable NLI-based evaluation.

Similar to this work, Chakrabarty et al. (2021a)

Fig. Type	Ent.	Gold/silver	Description	Count
Idioms	→	Silver	Replace idiom used in figurative context with definition	16652
	↯	Silver	Replace idiom used in literal context with definition	886
	↯	Silver	Replace idiom used in figurative context with adversarial definition	6116
	→	Gold	Hand written literal definition of idiom	532
	↯	Gold	Manual replacement of key words in definition w/ antonyms	375
	↯	Gold	Hand written non-entailed sentence	254
Metaphors	→	Silver	Replace metaphoric construction with literal construction	375
	→	Gold	Hand written literal paraphrase of metaphor	388
	↯	Gold	Hand written non-entailed sentence	282

Table 2: **Dataset Summary:** Overview of entailments/non-entailments in **IMPLI**. → denotes entailments, ↯ non-entailments. Note the descriptions are simplified: some intermediate steps are omitted (see §3.1).

build a dataset for NLI based on figurative language. Their dataset consists of figurative/literal pairs recast from previously developed simile and metaphor datasets, along with a parallel dataset between ironic and non-ironic rephrasing. This sets the groundwork for figurative NLI, but the dataset is relatively small outside of the irony domain, and the non-entailments are generated purely by replacing words with their antonyms, restricting the novelty of the hypotheses. Their dataset is relatively easy for NLI models; here we show that figurative language can be challenging, particularly with regard to non-entailments.

Zhou et al. (2021) and Chakrabarty et al. (2021a) provide invaluable resources for figurative NLI; our work aims to cover gaps in a number of areas. First, we generate a large number of both entailment and non-entailment pairs, allowing for better evaluation of adversarial non-entailing examples. Second, our silver methods allow for rapid development of larger scale data, allowing for model training and evaluation. We show that while entailment pairs are relatively easy (accuracy scores ranging from .86 to .89), the non-entailment pairs are exceedingly challenging, with the `roberta-large` model achieving accuracy scores ranging from .311 to .539.

### 3 Building a Dataset

Our **IMPLI** dataset is built from idiomatic and metaphoric sentences paired with entailing and non-entailing counterparts, from both silver pairs (§3.1) and manually written sentences (§3.4). For our purposes, we follow McCoy et al. (2019) in conflating the neutral and contradiction categories into a non-entailment label. We then label every pair as either entailment (→) or non-entailment (↯).

Due to the difficult nature of the task and to avoid issues with crowdsourcing (Bowman et al., 2020),

we employed expert annotators. We used two fluent English speakers, both graduate students in linguistics with strong knowledge in figurative language. For each method below, we ran pilot studies, incorporated annotator feedback and iteratively assessed the viability of identifying and generating appropriate expressions. As the annotators were working on generating new expressions, agreement was not calculated: we instead assessed the quality of the resulting expressions (see Section 3.3). Table 2 contains an overview of the different entailment and non-entailment types collected (Detail examples are also provided in Appendix D).

#### 3.1 Silver pairs

First, we explore a method for generating silver pairs using annotators to create phrase definitions which can be inserted automatically into relevant contexts, yielding a large number of possible entailment and non-entailment pairs that differ only with regard to the relevant phrase. Our procedure hinges on a key assumption: for any given figurative phrase, we can generate a contextually independent literal paraphrase. We then replace the original expression with the literal paraphrase, following the assumption that the figurative expression necessarily entails its literal paraphrase:

He’s stuck in bed, which is his *hard cheese*. → He’s stuck in bed, which is his *bad luck*.

Conversely, in contexts where the original phrase is used literally, replacing it with the literal paraphrase should yield a non-entailment relation.

Switzerland is famous for six cheeses, sometimes referred to as *hard cheeses*. ↯ Switzerland is famous for six cheeses, sometimes referred to as *bad luck*.




Figure 1: **Idiomatic definition replacement.** Pairs are generated using corrected dictionary definitions, substituted into figurative (left) and literal (center) sentences.

### 3.1.1 Idioms

To build idiomatic pairs, we use three corpora that contain sentences with idiomatic expressions (IEs) labelled as either figurative or literal.<sup>2</sup> These are the MAGPIE Corpus (Haagsma et al., 2020), the PIE Corpus (Adewumi et al., 2021), and the SemEval 2013 Task 5 (Korkontzelos et al., 2013). We collect the total set of IEs that are present in these corpora. We then extract definitions for these using freely available online idiom dictionaries.<sup>3</sup>

These definitions are often faulty, incomplete, or improperly formatted. We employed annotators to make manual corrections. The annotators were given the original IE as well as the definition extracted from the dictionary. The annotators were asked to ensure that the dictionary definition given was (1) a correct literal interpretation and (2) fit syntactically in the same environments as the original IE. If the definition met both of these criteria, the IE can be replaced by its definition to yield an entailment pair. If either criteria wasn't met, annotators were asked to minimally update the definition so that it satisfied the requirements.

In total this process yielded 697 IE definitions. We then used the above corpora, replacing these definitions into the original sentences (see Figure 1). We use the figurative/literal labels from the original corpora: replacing them into figurative contexts yields entailment relations, while replac-

<sup>2</sup>We here use "idiomatic expression" or "IE" to refer to the specific idiom in question (ie. "kick the bucket", "spill the beans"), as opposed to the sentence/context containing it.

<sup>3</sup>[www.theidioms.com](http://www.theidioms.com), [www.wiktionary.org](http://www.wiktionary.org)

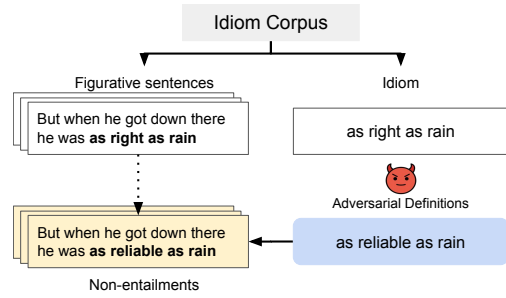


Figure 2: **Adversarial Pair Generation.** Non-entailing pairs are generated by replacing adversarial definitions into figurative contexts.

Original IE	Adversarial Definition
man of the cloth	tailor
heart of gold	cold, mean heart
come clean	bathe
turn a trick	do a magic trick

Table 3: Sample hand-written adversarial definitions.

ing them into contexts where the phrase is meant literally then yields non-entailments.

### 3.1.2 Adversarial Definitions

As a second method for generating non-entailment pairs, we asked annotators to write novel, adversarial definitions for IEs. Given a particular phrase, they were instructed to invent a new meaning for the IE that was not entailed by the true meaning, but which seemed reasonable presuming they had never heard the original IE. Some examples of this process are shown in Table 3.

We then replace these adversarial definitions into figurative sentences from the corpora. This yields pairs where the premise is an idiom used figuratively, and the hypothesis is a sentence that attempts to rephrase the idiom literally, but does so incorrectly, thus yielding non-entailments (Figure 2).

### 3.1.3 Metaphors

Metaphors are handled in a similar way: we start with a collection of minimal metaphoric expressions (MEs). These are subject-verb-object and adjective-noun constructions from Tsvetkov et al. (2014). Each is annotated as being either literal or metaphoric, along with an example sentence. We passed these MEs directly to annotators, who were then instructed to replace a word in the ME so that it would be considered literal in a neutral context.

1. *drop* prices → *reduce* prices
2. **hard** truth → **unpleasant** truth
3. *hairy* problem → *difficult* problem

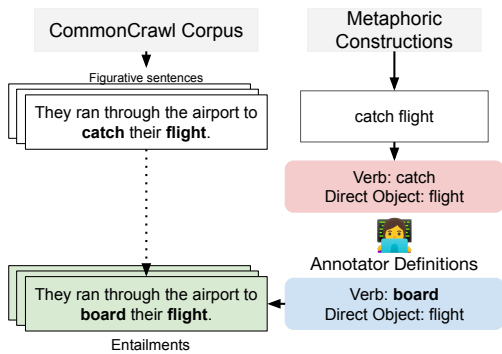


Figure 3: **Metaphor entailment generation.** Pairs are generated using annotator-defined literal translations substituted into metaphoric contexts.

These can then be replaced in a similar fashion: we start with the original figurative sentence, replace the ME with the literal replacements, and the result is an entailing pair with the metaphoric sentence entailing the literal.

We apply this procedure to the dataset of Tsvetkov et al. (2014), yielding 100 metaphoric/literal NLI entailment pairs. We then take a portion of the Common Crawl dataset<sup>4</sup>, and identify sentences that contain these original MEs. We identify sentences that contain the words from the metaphoric phrase, and replace the metaphoric word itself with its literal counterpart. This yields 645 additional silver pairs.

### 3.2 Postprocessing

For all silver methods, we also employ syntactic postprocessing to overcome a number of hurdles. First, phrases used idiomatically often follow different syntactic patterns than when used literally.

**Original:** These point *out of this world*, but where to is not made clear.

**Replaced:** \*These point *wonderful*, but where to is not made clear.

This phrase in literal contexts functions syntactically as a prepositional phrase, while idiomatically it is used as an adjective. When replaced with the definition "wonderful" in a literal context, we get a grammatically incoherent sentence. Second, phrases in their literal usage often don't form full constituents, due to the string-matching approach of the original datasets. Many literal usages of these phrases are thus incompatible with the defined replacement.

- I think [this one has *to die*] for the other one to live.

<sup>4</sup><https://commoncrawl.org/>

- Turn *in* [*the raw edges*] of both seam allowances towards each other and match the folded edges.

To avoid these issues, we ran syntactic parsing on the definition and the expression within each context, requiring that the expression in context begins with the same part of speech as the definition and that it does not end inside of another phrase.

Additionally, for each replacement, we ensured that the verb conjugation matched the context. For this, we identified the conjugation in the context, and used a de-lemmatization script to conjugate the replacement verb to match the original.

### 3.2.1 Additional Issues

In implementing and analyzing this procedure, we noted a number of practical issues. First, a large number of the MEs provided are actually idiomatic or proverbial: the focus word doesn't actually contribute to the metaphor, but rather the entire expression is necessary. Relatedly, we found that replacing individual parts of MEs is often insufficient to fully remove the metaphoric meaning. We iterated over possible solutions to circumvent these issues; we eventually decided to simply skip instances for which a replacement does not yield a feasible literal interpretation.

### 3.3 Evaluating Pair Quality

In order for these automatically created pairs to be useful for NLI-based evaluation, they need to be of sufficiently high quality. As the annotators were generating novel definitions and pairs, rather than inter-annotator agreement, we instead evaluate the quality of the resulting pairs by testing whether the automatically generated pairs contained the appropriate entailment relation. For this task, each annotator was given 100 samples for each general category of silver generations (idiomatic entailments, idiomatic non-entailments, and metaphoric entailments). They were asked if the entailment relation between the two sentences was as expected. An expert then adjudicated disagreements to determine the final percentage of valid pairs.

To evaluate the syntactic validity of the generated pairs, we additionally ran the Stanford PCFG dependency parser (Klein and Manning, 2003) on the pairs. Per previous work in NLI (Williams et al., 2018), we evaluate the proportion of sentences for which the root node is S.

Table 4 shows the results. The semi-supervised examples evoked the correct entailment relation

	→ Idioms	↯ Idioms	→ Met.
Correct Entailments	%88	%90	%97
Premise S root	%89	%90	%82
Hypothesis S root	%90	%90	%82

Table 4: **Valid pairs.** Percentage of valid pairs, syntactically and with regard to the intended entailments, of automatic data generation.

between %88 and %97 of the time: while there is still noise present, this indicates the effectiveness of the proposed methods. With regard to syntax, we see S node roots for between 82% and %90 of the sentences: within the range of the SNLI performance (74%-88%), and slightly behind the MNLI (91%-98%). We find that the generated hypotheses aren't significantly different in quality than the premises: this indicates that the method for generation preserves the original syntax.

These methods allow us to quickly generate a substantial number of high-quality pairs to evaluate NLI systems on figurative language. However, they may introduce additional bias as we employ a number of restrictions in order to ensure syntactic and semantic compatibility, and we lack full non-entailment pairs for metaphoric data. We therefore expand our dataset with manually generated pairs.

### 3.4 Manual Creation of Gold Pairs

To create gold pairs, annotators were given a figurative sentence along with the focus of the figurative expression: for idioms, this is the IE; for metaphors, the focus word of the metaphor. For idioms, we used the MAGPIE dataset to collect contextually figurative expressions. For metaphors, we collected metaphoric sentences from the VUA Metaphor Corpus (Steen et al., 2010), the metaphor dataset of (Mohammad et al., 2016), and instances from the Gutenberg poetry corpus (Jacobs, 2018) annotated for metaphoricality (Chakrabarty et al., 2021b; Stowe et al., 2021). Annotators were instructed to rewrite the sentence literally. This was done by removing or rephrasing the figurative component of the sentence. This yields gold standard paraphrases for idiomatic and metaphoric contexts.

We then asked annotators to write non-entailed hypotheses for each premise. They were encouraged to keep as much of the original utterance as possible, ensuring high lexical overlap, while removing the main figurative element of the sentence. For idioms, this comes from adding or adjusting words to force a literal reading of the idiom:

- The old girl finally kicked the bucket. ↯ The girl kicked the bucket on the right.

For metaphors, this typically involves keeping the same phrasing while adapting the sentence to have a different, non-metaphoric meaning.

- You must adhere to the rules. ↯ You must adhere the rules to the wall.

### 3.5 Antonyms

Previous work in NLI has employed the technique of replacing words in the literal sentences with their antonyms to yield non-entailing pairs (Chakrabarty et al., 2021a). We replicate this process for idioms: for the manually elicited definitions, we replace key words as determined by annotators with their antonyms. This yields sentences which negate the original figurative meaning and are thus suitable non-entailment pairs. Previous work found this antonym replacement for figurative language remains relatively easy for NLI systems; we can additionally explore this with regard to idioms.

These manual annotations provide a number of concrete benefits. First, they aren't restricted to individual words or phrases (excluding antonyms): the figurative components can be rewritten freely, allowing for diverse, interesting pairs. Second, they are written by experts, ensuring higher quality than the automatic annotations, which may be noisy.

## 4 Experiments / Results

Using the IMPLI dataset, we aim to answer a series of questions via NLI pertaining to language models' ability to understand and represent figurative language accurately. These questions are:

- R1: *How well do pretrained models perform on figurative entailments and non-entailments?*
- R2: *Does adding idiomatic pairs into the training data affect model performance?*
- R3: *Does the flexibility of idiomatic expressions affect model performance?*

Our dataset provides unique advantages in addressing these research questions that cover gaps in previous work: it contains a large number of both entailments and non-entailments, and is large enough to explore model training.

Model	MNLI	MNLI-MM	Idioms						Metaphors		
			→ S	↯ S <sup>l</sup>	↯ S <sup>d</sup>	→ G	↯ G <sup>a</sup>	↯ G	→ S	→ G	↯ G
roberta-base	.878	.876	.848	.539	.409	.890	.771	.311	.947	.818	.818
roberta-large	.899	.899	.866	.536	.418	.889	.777	.348	.936	.871	.840

Table 5: **R1: Model accuracy** Accuracy on MNLI and **IMPLI** pairs, divided into silver (S) and gold (G) datasets. S<sup>l</sup> Silver non-entailment based on replacement in literal contexts, S<sup>d</sup> Silver non-entailment based on adversarial definitions, G<sup>a</sup> Gold non-entailment based on antonyms.

## 473 R1: Pretrained Model Performance

474 We obtain baseline NLI models by fine-tuning  
475 roberta-base and roberta-large models  
476 on the MNLI dataset (Williams et al., 2018), with  
477 entailments as the positive class and all others as  
478 the negative, and evaluate them on their original  
479 test sets as well as **IMPLI**.<sup>5</sup> Due to variance in  
480 neural model performance (Reimers and Gurevych,  
481 2017), we take the mean score over 5 runs using  
482 different seeds.

483 We report results in Table 5. We observe that  
484 idiomatic entailments are relatively easy to classify,  
485 with accuracy scores over .84. Non-entailments  
486 were much more challenging. Silver pairs gener-  
487 ated through adversarial definitions were especially  
488 difficult: the pairs contain high lexical overlap, and  
489 in many cases the premise and hypotheses are se-  
490 mantically similar. The replacement into literal  
491 samples were easier, as the idiomatic definition  
492 clashes more starkly with the original premise,  
493 making non-entailment predictions more likely.  
494 Consistent with Chakrabarty et al. (2021a)’s work  
495 in metaphors, non-entailment through antonym re-  
496 placement is easiest for idioms: the antonymic  
497 relationship can be a marker for non-entailment,  
498 despite the high word overlap.

499 With regard to metaphors, silver entailment pairs  
500 are relatively easy. Manual pairs are more challeng-  
501 ing, but are still much easier than idioms. This  
502 is supported by the fact that metaphors are com-  
503 mon in everyday language: these models have  
504 likely seen the same (or similar) metaphors in train-  
505 ing. Our findings show that in fact metaphoric-  
506 ity may not be particularly challenging for deep  
507 pretrained models, as they are able to effec-  
508 tively capture the metaphoric entailment relations.  
509 The roberta-large model performs better  
510 for metaphoric expressions than roberta-base,  
511 but the difference on other partitions is relatively  
512 small.<sup>6</sup> We also find that lexical overlap plays a

<sup>5</sup>Model hyperparameters found in Appendix A.

<sup>6</sup>We found minimal differences between these models across R1-R3.

513 significant role here as noted by previous work (Mc-  
514 Coy et al., 2019): sentences with high overlap tend  
515 to be classified as entailments regardless of the true  
516 label (for more, see Appendix B).

517 We note that the manual pairs tend to be more  
518 difficult for both idioms and metaphors: these pairs  
519 can be more flexible and creative, whereas the sil-  
520 ver pairs are restricted to more regular patterns.

521 **R2: Incorporating Idioms into Training** To  
522 evaluate incorporating idioms into training, we  
523 then split the idiom data by idiomatic phrase types,  
524 keeping a set of IEs separate as test data to as-  
525 sess whether the model can learn to correctly han-  
526 dle novel, unseen phrases. Our goal is to assess  
527 whether poor performance is due to models’ not  
528 containing these expressions in training, or be-  
529 cause their ability to represent figurative language  
530 inherently limited. We hypothesize that the non-  
531 compositional nature of these types of figuration  
532 should lead to poor performance on unseen phrases,  
533 even if the model is trained on other idiomatic data.

534 For each task, we split the data into 10 folds by  
535 IE and incrementally incorporate these folds into  
536 the original MNLI for training, leaving one fold out  
537 for testing. We experiment with incorporating all  
538 training data for both labels, as well as using only  
539 entailment or non-entailment samples. We then  
540 evaluate our results on the entire test set, as well as  
541 the entailment and non-entailment partitions.

542 Figure 4 shows the results, highlighting that ad-  
543 ditional training data yields only small improve-  
544 ments. Pairs with non-entailment relations remain  
545 exceedingly difficult, with performance capping  
546 out at only slightly better than chance. As hypothe-  
547 sized, additional training data is only somewhat  
548 effective in improving language models’ idiomatic  
549 capabilities; this is not sufficient to overcome diffi-  
550 culties from literal usages of idiomatic phrases and  
551 adversarial definitions, indicating that idiomatic  
552 language remains difficult for pretrained language  
553 models to learn to represent.

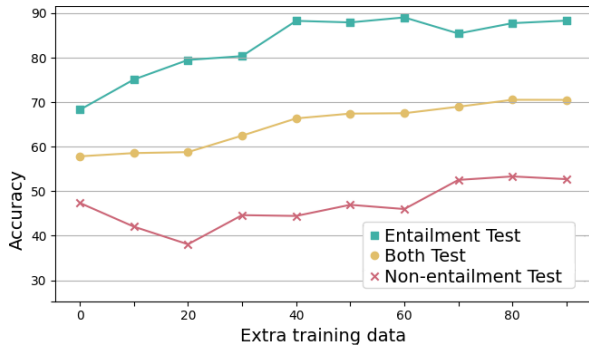


Figure 4: **R2: Training.** Performance of the roberta-base model as idiom types are added to the training data.

**R3: Syntactic Flexibility** Finally, we assess models’ representation of idiomatic compositionality. Nunberg et al. (1994) indicate that there are two general types of idioms: "idiomatic phrases", which exhibit limited flexibility and generally occur only in a single surface form, and "idiomatically combining expressions" or ICEs, in which the constituent elements of the idiom carry semantic meaning which can influence their syntactic properties, allowing them to be more syntactically flexible.

For example, in the idiom *spill the beans*, we can map the spilling activity to divulging of information, and the beans to the information. Because this expression has semantic mappings to figurative meaning for its syntactic constituents, Nunberg et al. (1994) argue that it can be more syntactically flexible, allowing for expressions like *the beans that were spilled by Martha* to maintain idiomatic meaning. For fixed expressions such as *kick the bucket*, no syntactic constituents map directly to the figurative meaning ("die"). We then expect less syntactic flexibility, and thus *the bucket that was kicked by John* loses its idiomatic meaning.

We hypothesize that model performance will be correlated with the degree to which a given idiom type is flexible: more fixed expressions may be easier, as they are seen in regular, fixed patterns that the models can memorize, while more flexible ICEs will be more difficult, as they can appear in different patterns, cases, and word order, often even mixing in with other constituents. To test this, we define an ICE score as the percentage of times a phrase occurs in our test data in a form that doesn’t match its original base form. Higher percentages mean the phrase occurs more frequently in a non-standard form, acting as a measure for the syntactic flexibility of the expression. We assessed the performance of the roberta-base model for each

idiom type, evaluating Spearman correlations between performance and idioms’ ICE scores.

We found no correlation between ICE scores and performance for entailments, nor for adversarial definition non-entailments ( $r = .004/.45$ ,  $p = .921/.399$ , see Appendix C). However, we do see a weak but significant correlation ( $r = .188$ ,  $p = 0.016$ ) with non-entailments from literal contexts: the model performs better when the phrases are more flexible, contrary to our initial hypothesis.

One possible explanation is that the model memorizes a specific figurative meanings for each fixed expression, disregarding the possibility of these words being used literally. When the expression is used in a literal context, the model then still assumes the figurative meaning, resulting in errors on non-entailment samples. The ICEs are more fluid, and thus the model is less likely to have a concrete representation for the given phrase: it is better able to reason about the context and interacting words within the expression, making it easier to distinguish the entailing and non-entailing samples.

## 5 Conclusions and Future Work

In this work we introduce the IMPLI dataset, which we then use to evaluate NLI models’ capabilities on figurative language. We show that while standard NLI models handle entailment admirably, and metaphoric expressions are relatively easy, non-entailment idiomatic relationships are more difficult. Additionally, adding idiom-specific training data fails to alleviate poor performance for non-entailing pairs. This highlights how currently language models are inherently limited in representing some figurative phenomena, and can provide a target for future model improvements.

For future work, we aim to expand our data collection processes to new data sources. Our dataset creation procedure relies on annotated samples and definitions: as more idiomatic and metaphoric resources become available, this process is broadly extendable to create new figurative/literal pairs. Additionally, we only explore this data for evaluating NLI systems: this data could also be used for other parallel data tasks such as figurative language interpretation (Shutova, 2013; Su et al., 2017) and figurative paraphrase generation. As natural language generation often relies on training or fine-tuning models with paired sentences, this data could be a valuable resource for figurative NLG systems.



641  
642  
643  
644  
645  
646  
647  
648  
649  
650  
651  
652  
653  
654  
655  
656  
657  
658  
659  
660  
661  
662  
663  
664  
665  
666  
667  
668  
669  
670  
671  
672  
673  
674  
675  
676  
677  
678  
679  
680  
681  
682  
683  
684  
685  
686  
687  
688  
689  
690  
691  
692  
693  
694  
695  
696  
697

## References

Tosin P. Adewumi, Saleha Javed, Roshanak Vadoodi, Aparajita Tripathy, Konstantina Nikolaidou, Foteini Liwicki, and Marcus Liwicki. 2021. [Potential idiomatic expression \(pie\)-english: Corpus for classes of idioms](#). *arXiv preprint arXiv:2105.03280*.

Rodrigo Agerri. 2008. [Metaphor in textual entailment](#). In *Coling 2008: Companion volume: Posters*, pages 3–6, Manchester, UK. Coling 2008 Organizing Committee.

Rodrigo Agerri, John Barnden, Mark Lee, and Alan Wallington. 2008. [Textual entailment as an evaluation framework for metaphor resolution: A proposal](#). In *Semantics in Text Processing. STEP 2008 Conference Proceedings*, pages 357–363. College Publications.

Yuri Bizzoni and Shalom Lappin. 2018. [Predicting human metaphor paraphrase judgments with deep neural networks](#). In *Proceedings of the Workshop on Figurative Language Processing*, pages 45–55, New Orleans, Louisiana. Association for Computational Linguistics.

Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015a. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.

Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015b. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.

Samuel R. Bowman, Jennimaria Palomaki, Livio Baldini Soares, and Emily Pitler. 2020. [New protocols and negative results for textual entailment data collection](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8203–8214, Online. Association for Computational Linguistics.

Tuhin Chakrabarty, Debanjan Ghosh, Adam Poliak, and Smaranda Muresan. 2021a. [Figurative language in recognizing textual entailment](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3354–3361, Online. Association for Computational Linguistics.

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

Ido Dagan, Oren Glickman, and Bernardo Magnini. 2006. [The pascal recognising textual entailment challenge](#). In *Machine Learning Challenges. Evaluating Predictive Uncertainty, Visual Object Classification, and Recognising Tectual Entailment*, pages 177–190, Berlin, Heidelberg. Springer Berlin Heidelberg.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Aniruddha Ghosh, Guofu Li, Tony Veale, Paolo Rosso, Ekaterina Shutova, John Barnden, and Antonio Reyes. 2015. [SemEval-2015 task 11: Sentiment analysis of figurative language in Twitter](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 470–478, Denver, Colorado. Association for Computational Linguistics.

Hessel Haagsma, Johan Bos, and Malvina Nissim. 2020. [MAGPIE: A large corpus of potentially idiomatic expressions](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 279–287, Marseille, France. European Language Resources Association.

Pere-Lluís Hugué Cabot, Verna Dankers, David Abadi, Agneta Fischer, and Ekaterina Shutova. 2020. [The Pragmatics behind Politics: Modelling Metaphor, Framing and Emotion in Political Discourse](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4479–4488, Online. Association for Computational Linguistics.

Adithy Iyer, Aditya Joshi, Sarvnaz Karimi, Ross Sparks, and Cecile Paris. 2019. [Figurative usage detection of symptom words to improve personal health mention detection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1142–1147, Florence, Italy. Association for Computational Linguistics.

Arthur M Jacobs. 2018. [The Gutenberg English poetry corpus: exemplary quantitative narrative analyses](#). *Frontiers in Digital Humanities*, 5:5.

Charles Jochim, Francesca Bonin, Roy Bar-Haim, and Noam Slonim. 2018. [SLIDE - a sentiment lexicon of common idioms](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association.

Dan Klein and Christopher D. Manning. 2003. [Accurate unlexicalized parsing](#). In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, pages 423–430, Sapporo, Japan. Association for Computational Linguistics.

756	Ioannis Korkontzelos, Torsten Zesch, Fabio Massimo Zanzotto, and Chris Biemann. 2013. <a href="#">SemEval-2013 task 5: Evaluating phrasal semantics</a> . In <i>Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)</i> , pages 39–47, Atlanta, Georgia, USA. Association for Computational Linguistics.	812
757		813
758		
759		814
760		815
761		
762		
763		
764	George Lakoff and Mark Johnson. 1980. <i>Metaphors We Live By</i> . University of Chicago Press, Chicago and London.	
765		
766		
767	Jens Lemmens, Iliia Markov, and Walter Daelemans. 2021. <a href="#">Improving hate speech type and target detection with hateful metaphor features</a> . In <i>Proceedings of the Fourth Workshop on NLP for Internet Freedom: Censorship, Disinformation, and Propaganda</i> , pages 7–16, Online. Association for Computational Linguistics.	
768		
769		
770		
771		
772		
773		
774	Vladimir Iosifovich Levenshtein. 1965. <a href="#">Binary codes capable of correcting deletions, insertions, and reversals</a> . In <i>Doklady Akademii Nauk</i> , volume 163, pages 845–848. Russian Academy of Sciences.	
775		
776		
777		
778	Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. <a href="#">RoBERTa: A robustly optimized bert pretraining approach</a> . <i>arXiv preprint arXiv:1907.11692</i> .	
779		
780		
781		
782		
783	Rui Mao, Chenghua Lin, and Frank Guerin. 2018. <a href="#">Word embedding and WordNet based metaphor identification and interpretation</a> . In <i>Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 1222–1231, Melbourne, Australia. Association for Computational Linguistics.	
784		
785		
786		
787		
788		
789		
790	Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. <a href="#">Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference</a> . In <i>Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics</i> , pages 3428–3448, Florence, Italy. Association for Computational Linguistics.	
791		
792		
793		
794		
795		
796		
797	Saif Mohammad, Ekaterina Shutova, and Peter Turney. 2016. <a href="#">Metaphor as a medium for emotion: An empirical study</a> . In <i>Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics</i> , pages 23–33, Berlin, Germany. Association for Computational Linguistics.	
798		
799		
800		
801		
802		
803	Yixin Nie, Yicheng Wang, and Mohit Bansal. 2019. <a href="#">Analyzing compositionality-sensitivity of nli models</a> . <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , 33(01):6867–6874.	
804		
805		
806		
807	Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. <a href="#">Adversarial NLI: A new benchmark for natural language understanding</a> . In <i>Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics</i> , pages 4885–4901, Online. Association for Computational Linguistics.	812
808		813
809		
810		814
811		815
	Geoffrey Nunberg, Ivan A. Sag, and Thomas Wasow. 1994. <a href="#">Idioms</a> . <i>Language</i> , (3):491–538.	
	Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. <a href="#">Hypothesis only baselines in natural language inference</a> . In <i>Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics</i> , pages 180–191, New Orleans, Louisiana. Association for Computational Linguistics.	816
		817
		818
		819
		820
		821
		822
	Nils Reimers and Iryna Gurevych. 2017. <a href="#">Reporting score distributions makes a difference: Performance study of LSTM-networks for sequence tagging</a> . In <i>Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing</i> , pages 338–348, Copenhagen, Denmark. Association for Computational Linguistics.	823
		824
		825
		826
		827
		828
		829
	Ivan A. Sag, Timothy Baldwin, Francis Bond, Ann Copestake, and Dan Flickinger. 2002. <a href="#">Multiword expressions: A pain in the neck for nlp</a> . In <i>Computational Linguistics and Intelligent Text Processing</i> , pages 1–15, Berlin, Heidelberg. Springer Berlin Heidelberg.	830
		831
		832
		833
		834
		835
	Ekaterina Shutova. 2013. <a href="#">Metaphor identification as interpretation</a> . In <i>Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 1: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity</i> , pages 276–285, Atlanta, Georgia, USA. Association for Computational Linguistics.	836
		837
		838
		839
		840
		841
		842
	Ekaterina V Shutova. 2011. <a href="#">Computational approaches to figurative language</a> . Technical report, Cambridge University.	843
		844
		845
	Gerard J. Steen, Aletta G. Dorst, J. Berenike Herrmann, Anna Kaal, Tina Krennmayr, and Trijntje Pasma. 2010. <a href="#">A Method for Linguistic Metaphor Identification: From MIP to MIPVU</a> . John Benjamins.	846
		847
		848
		849
	Anatol Stefanowitsch and Stefan Th. Gries. 2008. <a href="#">Corpus-Based Approaches to Metaphor and Metonymy</a> . De Gruyter Mouton.	850
		851
		852
	Kevin Stowe, Tuhin Chakrabarty, Nanyun Peng, Smaranda Muresan, and Iryna Gurevych. 2021. <a href="#">Metaphor generation with conceptual mappings</a> . In <i>Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)</i> , pages 6724–6736, Online. Association for Computational Linguistics.	853
		854
		855
		856
		857
		858
		859
		860
		861
	Chang Su, Shuman Huang, and Yijiang Chen. 2017. <a href="#">Automatic detection and interpretation of nominal metaphor based on the theory of meaning</a> . <i>Neurocomputing</i> , 219:300–311.	862
		863
		864
		865

866	Yulia Tsvetkov, Leonid Boytsov, Anatole Gershman,	<b>A Model Hyperparameters</b>	917
867	Eric Nyberg, and Chris Dyer. 2014. <a href="#">Metaphor detection with cross-lingual model transfer</a> . In <i>Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 248–258, Baltimore, Maryland. Association for Computational Linguistics.	We use a fixed set of hyperparameters for all NLI fine-tuning experiments: learning rate of $1e^{-5}$ , batch size 32, and maximum input length of 128 tokens. The models are trained for 3 epochs.	918
868			919
869			920
870			921
871			
872			
873	Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob	<b>B Lexical Overlap</b>	922
874	Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz	Previous research shows that NLI systems exploit cues based on lexical overlap, predicting entailment for overlapping sentences (McCoy et al., 2019; Nie et al., 2019). Our dataset consists mostly of pairs with high overlap: this could explain why the non-entailment sections are more difficult. We thus evaluate system predictions for our datasets as a function of lexical overlap. Figure 5 shows density-based histograms of the results, comparing overlap via Levenshtein distance (Levenshtein, 1965) for correctly and incorrectly classified pairs.	923
875	Kaiser, and Illia Polosukhin. 2017. <a href="#">Attention is all you need</a> . In <i>Advances in Neural Information Processing Systems</i> , volume 30. Curran Associates, Inc.		924
876			925
877			926
878	Tony Veale, Ekatarina Shutova, and Beata Beigman		927
879	Klebanov. 2016. <i>Metaphor: A computational perspective</i> . Morgan and Claypool.		928
880			929
881	Alex Wang, Yada Pruksachatkun, Nikita Nangia,		930
882	Amanpreet Singh, Julian Michael, Felix Hill, Omer		931
883	Levy, and Samuel R. Bowman. 2019. SuperGLUE: A stickier benchmark for general-purpose language understanding systems. <i>arXiv preprint 1905.00537</i> .		932
884			933
885			934
886	Alex Wang, Amanpreet Singh, Julian Michael, Felix		935
887	Hill, Omer Levy, and Samuel Bowman. 2018. <a href="#">GLUE: A multi-task benchmark and analysis platform for natural language understanding</a> . In <i>Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP</i> , pages 353–355, Brussels, Belgium. Association for Computational Linguistics.		936
888			937
889			938
890			939
891			940
892			941
893			942
894	Adina Williams, Nikita Nangia, and Samuel Bowman.		943
895	2018. <a href="#">A broad-coverage challenge corpus for sentence understanding through inference</a> . In <i>Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)</i> , pages 1112–1122. Association for Computational Linguistics.		944
896			945
897			946
898			947
899			948
900			
901			
902	Shweta Yadav, Jainish Chauhan, Joy Prakash Sain,	<b>C Syntactic Flexibility Correlations</b>	949
903	Krishnaprasad Thirunarayan, Amit Sheth, and	Figure 6 shows correlations between ICE scores (determined by frequency of occurrences of a given IE outside of its normal form) and roberta-base model performance on that IE.	950
904	Jeremiah Schumm. 2020. <a href="#">Identifying depressive symptoms from tweets: Figurative language enabled multitask learning framework</a> . In <i>Proceedings of the 28th International Conference on Computational Linguistics</i> , pages 696–709, Barcelona, Spain (Online). International Committee on Computational Linguistics.		951
905			
906			
907			
908			
909			
910			
911	Jianing Zhou, Hongyu Gong, and Suma Bhat. 2021.	<b>D Dataset Examples</b>	950
912	<a href="#">PIE: A parallel idiomatic expression corpus for idiomatic sentence generation and paraphrasing</a> . In <i>Proceedings of the 17th Workshop on Multiword Expressions (MWE 2021)</i> , pages 33–48, Online. Association for Computational Linguistics.	Table 6 shows examples from each type of pair generation.	951
913			
914			
915			
916			




Figure 5: **R2: Lexical Overlap.** Classification performance by lexical overlap. The x axis shows Levenshtein distance; the y axis shows stacked density of correctly and incorrectly tagged pairs. The **IMPLI** non-entailments contain extremely high overlap, and are thus frequently misclassified as entailment.




Figure 6: **R3: Syntactic Flexibility.** Performance of idiom types compared to their syntactic flexibility (based on ICE score defined in **R3**), with Spearman coefficient correlations  $r$  and significance values  $p$ . The middle figure is non-entailments based on replacement in literal context; the right is those based on adversarial definitions. Further right on the x-axis indicates greater flexibility.

### Idioms

(→ S) Replace idiom used in figurative context with definition

---

BITTER BLOW: Beer sales are *feeling the pinch*. → BITTER BLOW: Beer sales are *suffering a hardship*.  
I must *have a word* with them. → I must *speak privately* with them.  
I've been *knocked out cold*. → I've been *knocked unconscious*.

---

(↯ S<sup>l</sup>) Replace idiom used in literal context with definition

---

It would be good to roll *in hot water* all over. ↯ It would be good to roll *in a difficult situation* all over.  
Pour *in the soup*. ↯ Pour *in trouble*.  
There's a marina down *in the docks*. ↯ There's a marina down *under scrutiny*.

---

(↯ S<sup>d</sup>) Replace idiom used in figurative context with adversarial definition

---

After *taking a bow*, the cast met Margaret backstage. ↯ After *apologizing*, the cast met Margaret backstage.  
I've been *knocked out cold* ↯ I've been *knocked out into the cold air*.  
It worked *like a charm!* ↯ It worked *poorly!*

---

(→ G) Hand written literal definition of idiom

---

How have you *weathered the storm?* → How have you *succeeded in getting through the difficult situation?*  
It *breaks my heart* that his career has been ruined. → It *overwhelms me* that his career has been ruined.  
Jamie rushed out *pissed off* and upset this afternoon. → Jamie rushed out *irritated* and upset this afternoon.

---

(↯ G<sup>a</sup>) Manual replacement of key words in definition w/ antonyms

---

Alison *makes the grade* for Scotland ↯ Alison *fails* for Scotland.  
I'll *catch a cold* ↯ I'll become *healthy*  
It's very much *swings and roundabouts* ↯ It's very much *one-sided*.

---

(↯ G) Hand written non-entailed sentence

---

How have you *weathered the storm?* ↯ How have you *calmed the storm?*  
Now Paul will *think twice*. ↯ Now Paul will *score twice*.  
They *went to ground* somewhere in the area. ↯ They *went to party* somewhere in the area.

---

### Metaphors

(→ S) Replace metaphoric construction with literal construction

---

Don't go and *blow your paycheck*. → Don't go and *waste your paycheck*.  
My computer *battery died*. → My computer *battery lost all power*.  
Competition is *dropping prices*. → Competition is *reducing prices*.

---

(→ G) Hand written literal paraphrase of metaphor

---

He *absorbed* the knowledge or beliefs of his tribe. → He *mentally assimilated* the knowledge or beliefs of his tribe.  
Avon *treads warily*. → Avon *proceeds warily*.  
All the *hearts of men were softened*. → All the *men were made kinder and gentler*.

---

(↯ G) Hand written non-entailed sentence

---

The gun kicked back into my shoulder. ↯ The mule kicked back into my shoulder.  
This was conveniently *encapsulated* on the first try. ↯ This was conveniently *encapsulated* in the first battle.  
On their tracks his eyes were *fastened*. ↯ On their tracks his *hands* were fastened.

---

Table 6: *Dataset Summary*: Overview of each entailment/non-entailment category in the **IMPLI** dataset.