

CLIX: Cross-Lingual Explanations of Idiomatic Expressions

Aaron Gluck¹ Katharina von der Wense^{1,2} Maria Leonor Pacheco¹

¹University of Colorado Boulder ²Johannes Gutenberg University Mainz
{aaron.gluck, katharina.kann, maria.pacheco}@colorado.edu

Abstract

Automated definition generation systems have been proposed to support vocabulary expansion for language learners. The main barrier to the success of these systems is that learners often struggle to understand definitions due to the presence of potentially unfamiliar words and grammar, particularly when non-standard language is involved. To address these challenges, we propose CLIX, the task of Cross-Lingual explanations of Idiomatic eXpressions. We explore the capabilities of current NLP models for this task and observe that, while it remains challenging, large language models show promise. Finally, we perform a detailed error analysis to highlight the key challenges that need to be addressed before we can reliably incorporate these systems into educational tools.

1 Introduction

The use of technology has become an important part of the language learning process (Golonka et al., 2014; Zhang and Zou, 2022). One of the main areas of interest for the proponents of technology-assisted language learning is vocabulary expansion, where recent studies have demonstrated a significant impact in student engagement and increased vocabulary knowledge (Fisher, 2016; Guaqueta and Castro-Gárces, 2018; Tao Hao and Ardasheva, 2021). To support the development of these technologies, considerable work has been devoted to the study of automated definition generation (Ni and Wang, 2017; Gadetsky et al., 2018; Ishiwatari et al., 2019; Bevilacqua et al., 2020).

The effective use of definition generation systems in second language vocabulary expansion is complicated by phenomena such as the coexistence of many possible meanings for a particular word or phrase (Enayati and Gilakjani, 2020; Tyler and Evans, 2003), and the potential presence of unfamiliar words and grammar in the generated definitions (Zhang, 2011; Kong et al., 2022). While we make

no assertion that any particular word in a definition will be unfamiliar to a language learner, the fact remains that there *might* be one or more such words. However, providing the language learner with an explanation in their native language may help remove this concern in many cases. In addition, a lot of these systems ignore the complexities of non-standard language such as idiomatic expressions, which are generally composed of words that combine to represent a meaning that does not equal the sum of its parts (Nunberg et al., 1994; Kövecses and Szabó, 1996; Adelnia and Dastjerdi, 2011).

Motivated by these challenges, we propose the task of Cross-Lingual explanation of Idiomatic eXpressions (CLIX). Given a source language idiomatic expression, we aim to **generate a natural language explanation in a specified target language**. As opposed to *definition*, we use the term *explanation* to allow for more flexibility in tailoring explanations to certain situations as well as potential inclusions of context. Rather than expecting a static, one-to-one mapping from idiom to definition, we allow models to approximate a one-to-many function. For example, explanations can be considered correct if they include elaborations, examples of use, and etymological information. The core concern of this research is the efficacy of the explanation and the final outcome for a potential student.

We additionally hypothesize that obtaining explanations of idiomatic expressions in the learner’s first language removes many of the barriers to understanding introduced by traditional definition generation systems. We choose to focus on idiomatic expressions as they are an important element of language learning that is particularly challenging for learners and automated systems alike. Consider the utterance, *he and I don’t see eye to eye on a variety of topics*. The idiomatic expression contained within this sentence is not composed of particularly challenging words, yet in this instance

it denotes a negative stance relative to the referent’s opinions and its meaning has nothing to do with the general definitions of the words *see* and *eye*. This example helps to display the fact that idiomatic expressions often are of unpredictable form, appear in limited numbers of syntactic frames, include various types of figuration, indicate an affective stance, and can describe situations of social interest (Nunberg et al., 1994).

To explore the capabilities and limitations of modern NLP systems for generating cross-lingual explanations of idiomatic expressions, we construct two datasets of English idioms and their explanations in English, Spanish, and German: EPIE-ME, an open dataset composed of 628 unique idioms with multi-lingual explanations built from the EPIE corpus (Saxena and Paul, 2020), and Oxford-ME, composed of 6,218 unique idioms taken from the Oxford Dictionary of Idioms 4th Edition (Ayto, 2020). We explore performance of both pre-trained sequence-to-sequence models and Large Language Models (LLMs), and show that our task configuration is much more challenging than previous definition generation tasks. We further provide some discussion on why naïve translation approaches do not suffice for this task.

The explicit goal of this paper is to motivate the need for and inspire approaches to designing systems that assist language learners. A summarization of our contributions is as follows:

1. We introduce a new cross-lingual idiomatic expression explanation task.
2. We release, to our knowledge, the first cross-lingual dataset for this task containing English idiomatic expressions as well as explanations of their meanings in English, Spanish, and German.¹
3. We evaluate a set of state-of-the-art models for the task and show that, while automated metrics paint a less-favorable picture, the outcome of few-shot LLMs is judged positively by native speakers of the target language.
4. We perform a detailed error analysis and highlight key areas that need improvement before we can confidently use automated cross-lingual explanations for non-standard language in educational applications.

¹Data and code available at: <https://github.com/blast-cu/CLIX>. We are not able to publish the Oxford-ME portion due to copyright, but can provide the data upon request.

2 Related Work

Disambiguation of Non-Standard Language

Non-standard language takes many different forms ranging from slang to metaphor to idiom. To disambiguate some more general forms of non-standard language, Ni and Wang (2017) proposed a hybrid word-character sequence-to-sequence model that can disambiguate language found on social media. Kurfah and Östling (2020) show that sufficiently different contextual embeddings can be found for both literal and idiomatic instances of a multi-word expression, thus allowing for disambiguation of the expressions in context. Qiang et al. (2023), on the other hand, try to directly paraphrase Chinese idioms such that they can replace the idiom using an infill strategy similar to the cloze task. These approaches are all tested in monolingual settings. The cross-lingual settings for which a solution might benefit language learners remain less explored.

The cross-lingual version of idiomatic expression disambiguation has traditionally been approached as a translation task. Often, sentences containing idiomatic expressions are mapped to equivalent sentences without such an expression, but models may try to translate them word-by-word (Baziotis et al., 2023). Fadaee et al. (2018) explore sentence-level translation to generate an English–German idiom translation dataset, additionally using attention-based model to translate sentences containing idiomatic expressions, requiring in-context disambiguation. Zaninello and Birch (2020) show that data augmentation and external linguistic resources can sizably improve translations of multi-word expressions. More recently, Li et al. (2024) take a very similar prompting-based approach to us. However, the authors’ primary goal appears to be the creation of an ontology, they perform their experiments with languages we do not consider, and they additionally remain focused on a translation-based perspective. We, however, formulate the task more closely to definition generation and explore the capabilities of NLP systems for this purpose in addition to applications in high-stakes scenarios such as education. We believe that this reformulation opens up interesting research avenues, such as investigating whether end-users benefit more from explanations, examples usage, or a combination. We also see new datasets specifically intended for determining whether an idiomatic expression is being used literally or figuratively given some context – now at the cross-

lingual level (Sentsova et al., 2025). To the best of our knowledge, however, we are the first to release a cross-lingual dataset specifically for the purpose of explaining these kinds of expressions.

Definition Generation Since Noraset et al. (2017) introduced this task, new techniques have continued to be introduced to improve performance. Gadetsky et al. (2018) created embeddings from dictionary definitions, while taking into account the issue of polysemy by allowing their model to recognize that there could potentially be an infinite number of senses for a given word. In contrast, using a novel encoding scheme, Bevilacqua et al. (2020) are able to autoregressively generate definitions as well as score them in a discriminative manner. Zhang et al. (2023) generate cross-lingual definitions of words in English, Chinese, and French using prompt learning. They apply a contrastive learning objective to correct pre-trained models’ tendency to ignore this task and/or mix languages in the output. In this paper, we focus specifically on idiomatic expressions and their explanations, which allows us to be more flexible and we hypothesize is more challenging for NLP systems.

3 CLIX

In this section, we define the CLIX task and describe the process to create a dataset of idiomatic expressions in English with their corresponding explanations in Spanish and German. We present statistics for the resulting dataset and describe the models that we will use to generate predictions.

3.1 Task Definition

Our goal is to provide an explanation of the meaning of an idiomatic expression in their native language of a language learner with the assumption that it would be more understandable than an explanation in the language they are attempting to learn. Following prior work on definition generation, we formalize our task as a text-to-text generation task. Given an idiomatic expression $I = \{i_1, i_2, i_3, \dots\}$ and some optional context $C = \{c_1, c_2, c_3, \dots\}$ in the source language L_S , our goal is to produce an explanation $E = \{e_1, e_2, e_3, \dots\}$ in the target language L_T . While this formulation is essentially identical to that of definition generation, we opt to call our generated text explanations. While minute, the difference we wish to highlight is that we feel a definition is generally more static as opposed to an explanation which may be tailored to an individ-

EPIE-ME	Train	Development	Test
Idioms	278	150	200
I Avg Len	3.04	3.38	3.16
EN Avg Len	18.23	17.16	17.16
ES Avg Len	19.30	17.98	9.63
DE Avg Len	18.52	17.41	9.10
Oxford-ME	Train	Development	Test
Idioms	4352	933	933
I Avg Len	3.74	3.77	3.68
EN Avg Len	9.04	9.08	9.08
ES Avg Len	9.34	9.47	9.11
DE Avg Len	9.21	9.21	9.06

Table 1: Dataset Statistics. I Refers to Idiom, EN, ES, and DE Refer to English, Spanish, and German Explanations

ual or particular context. Explanations might also include additional contextual information that can be realized in either speech or writing. While we do not require any of these factors to be present in our explanations, because our focus is on language learners and educational contexts, we wish to allow flexibility in how explanations are presented to a user in current and future exploration of this task. Furthermore, EPIE (Saxena and Paul, 2020) and our modified version EPIE-ME do not only contain static definitions. We see samples such as *cheek and jowl* which are instead given an explanation: *Positioned very close together. The cheek and the jowl — the lower part of the jaw — are in close proximity to each other on the face.*

3.2 Dataset

We construct – to the best of our knowledge – the first two datasets for cross-lingual explanation generation of idiomatic expressions, EPIE-ME and Oxford-ME (EPIE and Oxford with Multilingual Explanations), each containing explanations from languages belonging to two separate branches of the Indo-European family, Spanish and German. We build on the EPIE corpus (Saxena and Paul, 2020), composed of 25,206 sentences containing instances of 717 formal and static idiomatic expressions with definitions, and example usages.² We select this dataset because it was in-part originally meant for detection of idiomatic expressions. As a result, it also contains expressions that are situationally idiomatic, a useful quality for judging explanation ability. We filter out expressions for which there are no gold explanations or in-context

²Formal idiomatic expressions can undergo lexical changes, while static ones cannot.

examples available, leaving us with 628. We additionally use the Oxford Dictionary of Idioms, 4th Edition (Ayto, 2020), comprising of 6,218 English idiomatic expressions and their explanations. This is generally considered the most comprehensive resource for idiomatic expressions in English. The Oxford data only comes with sentence-level context for 1,780 of the expressions (28.6%), so we prompt Llama 3.1 (Dubey et al., 2024) to generate sentences using missing expressions in the manner found in Tab. 14. For each unique expression in the corpus, we include an explanation in English, Spanish, and German. We obtain non-English explanations by feeding the gold English explanations to Google Translate. In Tab. 1 we show the train/dev/test splits and the average word length of the idioms and explanations for both datasets. Splits are generated by random sample.

Manual Corrections For the test sets, the English explanations and the automated translations are closely inspected and manually corrected by native speakers of all three languages to ensure they are entirely correct and not just the output of the automated translation. Any confusion about the meaning of a particular expression was referred to the native English speaker for review. These annotators are all computational linguists fluent in English and working in English speaking countries.

Idiom Categorization For EPIE-ME, we further enhance our idioms with thematic categories inspired by findings suggesting that classifying idioms facilitates the learning process (Boers, 2000). To do this, we build on the taxonomy developed by Rafatbakhsh and Ahmadi (2019), which suggests a list of 81 themes to classify English idioms (See full list with frequencies in App. C). For any idioms already assigned themes by Rafatbakhsh and Ahmadi (2019) (a little over 20%), we use the assigned theme. To tag the remainder of our dataset using this taxonomy, we prompt GPT-3.5 Turbo (OpenAI, 2023) to provide us with the three most appropriate categories over three runs where the categories are provided in random order. We then eliminate duplicates in the resulting nine predictions and correct them by hand. The manual corrections were done by two English speaking annotators working together to resolve conflicts.

Estimating Noise in Non-Corrected Examples Finally, we examine the amount of noise in the explanations without manual correction. We do so by

calculating the edit distance (Levenshtein, 1965) between the original output of Google Translate and the corrected version produced by our annotators over our verified test set. We further normalize by length following Tashima et al. (2018). We see that the amount of correction required is generally low, the largest amount being required on average for the German EPIE-ME explanations with a normalized edit distance of 0.283, while Spanish had an edit distance of 0.159. More information can be found in Appendix E.

4 Learning and Inference Strategies

In this section, we outline the different learning and inference strategies that we experiment with to generate cross-lingual explanations of idiomatic expressions. We experiment with fine-tuned and few-shot generation models. Strategies include: whether performing the task end-to-end or decomposing it into explanation and translation, the type of contextual information used to enhance the input prompts, and the number of demonstrations used for LLM-based solutions. All prompt templates can be found in App. D

4.1 Explanation Strategies

Direct Explanation To obtain an explanation for the benefit of a potential language learner, we fine-tune and/or prompt models for an explanation in what we assume to be the L1 of the learner. All inputs, however, are in English. Our training pairs in the fine-tuning setup consist of an input $X_i = [T_{i,L_T}; I_i]$ and label $Y_i = [E_{i,L_T}]$, where T_{i,L_T} is set of tokens informing the model that its task is to produce an explanation in a target language L_T , I_i is an idiomatic expression, and E_{i,L_T} is a gold explanation of the idiomatic expression. This model can then be formulated as:

$$P(Y_k|X_k) = \prod_t (y_k | y_{0:k-1}, X_k; \theta) \quad (1)$$

Pipeline In this configuration, we aim to generate our final explanation in two stages. The first stage is to generate an explanation in English. In our fine-tuning setup, we use T5 (Raffel et al., 2020) for this step. Where mT5 (Xue et al., 2021) was explicitly designed to learn many languages, T5 was only trained to translate between a small set, and most of its training data was English. Thus, we feel T5 is more likely to perform well on this task. The model definition remains the same as in Eq. 1.

Theme	Spanish	German
Language, speech, and conversation	71.66	72.07
Happiness, pleasure, and enjoyment	76.88	79.41
Health and illness	70.20	71.46
Argument and conflict	78.97	77.62
Cooperation	77.61	79.01
Experience	71.99	68.65
Duty and responsibility	68.37	69.14
Success	77.92	78.95
Haste and speed	72.77	76.57
Anger and annoyance	87.98	73.12

Table 2: Performance of Best Model as Measured by Sentence Similarity on EPIE-ME Data per Theme for Top 10 Themes

However, the input pair is now $X_i = [T_i; I_i]$ and $Y_i = [E_{i,L_S}]$. In this case, our task sequence T_i no longer provides any language information, and our explanation E_{i,L_S} remains in the source language (in our case, English).

The second stage of our pipeline requires a translation of the output of the previous step. Because T5 was trained for the task of translation to/from German but not to/from Spanish, any comparison between results would be unfair. While we considered the use of mT5 in this second stage, its performance on translation tasks is already well understood (Xue et al., 2021). As a result, and to additionally be consistent with our choice in the curation of our dataset, we use Google Translate for the second step of our pipeline in the fine-tuning scenarios. In the experiments performed with GPT and Llama, however, the models are asked to translate their own output.

To validate our assumptions we provide results for a full-mT5 pipeline where we fine-tune versions for each of the explanation and translation tasks.

4.2 Contextual Enhancements

We explore different ways to add contextual information into the input prompt. This is done only for LLM models: Llama and GPT.

Sentence-Level Context Because idiomatic expressions encode information that is potentially unrelated to their linguistic form, many researchers have used sentence-level context to assist with identification, disambiguation, and translation of idiomatic expressions (Fadaee et al., 2018; Kurfali and Östling, 2020; Saxena and Paul, 2020; Liu et al., 2023; Li et al., 2024). We experiment with including example sentences (one usage example per instance) containing the idiomatic expression

to be explained at the end of our prompts in an effort to provide more information to the model.

Categorical Enhancements In this setting we explore how the addition of small hints can help these models. We add to our prompt a category that an idiomatic expression might belong to. All categories we use come from the Rafatbakhsh and Ahmadi (2019) taxonomy.

4.3 Few-Shot Prompting

We explore the effects of few-shot prompting with both Llama and GPT. We implement scenarios where the models are provided with 1, 3, or 5 examples from which they can learn in-context. While there is some expectation that more demonstrations will increase performance, we also aim to determine whether or not the same categorical information used for context enhancement could be used to strategically select examples and further increase performance. We compare random selection to a scenario in which we induce the category label from the LLM and one in which we assume that the category label is known.

To induce a label, we ask the model to choose the most suitable label from our list of categories. If the induced label does not exactly match an existing label, we add it to our list and treat it as an existing label for only the instance in which we observe this idiomatic expression, and from this point forward, the procedure is the same as if we assume the category is known. Once we have a known category label (either given or acquired by induction), meaning we are given some idiomatic expression I_i with category C_i , we form an initial selection of $2 * k$ idiomatic expressions with category label C_j where $C_i = C_j$ and k is the number of shots. If there are too few expressions to form this selection, we draw from the most similar categories until we obtain $2 * k$ examples. Similarity between categories is calculated as the similarity between the labels by the same SBERT (Reimers and Gurevych, 2020) model used for our evaluation in Sec. 5. Once these $2 * k$ examples are obtained, we take a final sample of k from this set. This sampling strategy is necessitated by the small size of our dataset, where we may not have even a single example in our training set with the same category label as one in our validation or test set. The first step of selecting $2 * k$ initial examples helps reduce variance between instances where there are enough examples where the categories match and instances

Model	EPIE-ME			Oxford-ME		
	Overall	Spanish	German	Overall	Spanish	German
mT5 Direct	38.06	36.32	39.80	43.21	43.56	42.86
mT5 Pipeline	37.13	37.13	37.81	42.58	42.40	42.77
T5 Pipeline	43.54	43.19	43.89	46.09	45.96	46.22
Llama Direct Zero-Shot	59.39	60.13	58.66	55.32	56.01	54.62
Llama Pipeline Zero-Shot	60.16	60.48	59.83	55.01	55.10	54.91
GPT Direct Zero-Shot	65.06	65.96	64.16	61.03	62.02	60.03
GPT Pipeline Zero-Shot	69.60	70.12	69.07	66.10	66.28	65.91
GPT Direct 5-Shot	71.15	71.42	70.87	66.13	66.9	65.36
GPT Pipeline 5-Shot	71.84	72.04	71.63	68.54	68.87	68.20

Table 3: Best Model Results by Sentence Similarity for Direct and Pipeline Strategies. LLM Results are Averaged Across 3 Runs on the Test Set. Scores for English Explanations in the First Half of the Pipeline can be Found in Tab. 8

Enhancement	GPT	Llama
Direct	63.91	60.45
+ SL	61.36	59.73
+ Cat	61.96	59.44
+ SL + Cat	61.61	58.98
Pipeline	66.98	61.32
+ SL	66.22	64.00
+ Cat	65.64	60.53
+ SL + Cat	66.36	63.04

Table 4: Zero-Shot Results on EPIE-ME by Contextual Enhancement. SL: Sentence-Level, Cat: Categorical Information.

where there are too few.

5 Experiments

We compare results of several models across a range of strategies outlined in Sec. 4. We aim to investigate what makes certain models better at understanding idiomatic expressions and what types of supplemental information are beneficial for in-context learning within this domain. We discuss the addition of contextual information outlined in Sec. 4.2. We also more closely examine methods of selecting good demonstrations in few-shot prompting scenarios when levels of information for the newly-observed idiomatic expression vary, as described in Sec. 4.3.

Models and Metrics In our experiments, we benchmark the performance of two sequence-to-sequence models, T5 (Raffel et al., 2020) and mT5 (Xue et al., 2021), as well as two LLMs, Llama3.1³ (Dubey et al., 2024) and GPT-3.5 Turbo⁴ (OpenAI, 2023). To evaluate our approaches we primarily

³Llama3.1 8B Instruct

⁴Model Version: gpt-3.5-turbo-0125

follow Zhang et al. (2023). To determine the quality of the generated explanations, we calculate their similarity to the gold explanations in the English using a multilingual version of SBERT (Reimers and Gurevych, 2020). We use this metric as it is more likely to assess the meaning of the expressions. Finally, in Tab. 6 we include BLEU and ROUGE scores calculated with our gold translated explanations for the best models using SacreBLEU (Post, 2018) and ROUGE-L.

5.1 Results and Discussion

Our main results are described in Tab. 3. In the case of mT5 and T5, the quality of the results is generally unimpressive, demonstrated by the semantic similarity scores hovering around 40. However, we find that a T5-based pipeline strategy increased performance by more than 10%. Upon manually inspecting the outputs, in both cases the output is solely in the indicated language, but we did not see a significant difference in the fluency between the different strategies. There were, however, cases where the model would continuously repeat a key word or phrase in its response. Examples include *Una persona típicamente **atractiva**, **atractiva** o **atractiva*** and *Un estado de **desorientación** o **desorientación***. In the case of the models trained with Oxford-ME data, we only see a 10% increase in performance, indicating that models trained on the smaller EPIE-ME dataset remain competitive with those trained on 10 times the amount of data.

In our experiments using LLMs, we see significant improvements in performance, reaching a similarity of nearly 70 with the best zero-shot strategy. This result is unsurprising, given that GPT and Llama have been exposed to significantly larger

amounts of data, likely including many in context occurrences of the idioms in our dataset. Consistent with the results from the sequence to sequence models, pipeline strategies seem to be most effective. While the final difference in score between the 5-shot direct and pipeline approaches is very small, we see a larger gap in the zero-shot setting. We believe that the increased gains in the direct setup are due to the direct scenario requiring more internal reasoning steps, where the demonstrations appear to help make up the difference. In contrast, the pipeline method appears to approach an upper bound in similarity scores sooner. In the case of the similar scores when using Llama in a zero-shot setting, we cannot find an exact reason. However, the fact that the pipeline approach universally performs better indicates that it is most effective. For the Oxford-ME dataset, we observe that the improvement of LLMs over fine-tuned models is smaller. We attribute this to the gold explanations in the Oxford-ME dataset being about 33% shorter, penalizing the often verbose responses of the LLMs.

In the final results of our experiments, we found that while additional context benefited Llama in certain cases, it did not help GPT. While we are unable to pinpoint why it was unhelpful, we hypothesize that it is a result of the model both being more capable and the lack of a direct link between the contextual enhancement and the gold explanation. This absence is particularly noticeable in relation to experiments where the model is provided categorical information. Dev set performance can be found in Tab. 4. Lastly, we see consistently higher performance by GPT when compared to Llama across all metrics. This is likely due to our use of the – small for an LLM – 8 billion parameter version of Llama, and we do not necessarily expect to see the same trend when using a larger version.

When comparing our zero and few-shot experiments, we can clearly see an increase in performance as the number of shots increases. However, we find that differing selection strategies for our examples in the few-shot scenarios do not help whatsoever. The model that performed the best used the random selection strategy. This indicates to us that either the model already represents the idiomatic expressions very well, or the relatedness of examples chosen for in-context learning are largely unimportant for this task. Results calculated on our dev set can be found in Tab. 9.

Finally, we perform an evaluation of the performance per theme (See Tab. 2) using the best

performing model on our test set. We can see that for the majority of our most common themes, this model is performing above expectations with respect to the overall results in Tab. 3. Intuitively, the model is able to better explain idiomatic expressions from more common domains.

5.2 A Note on Translation

Although there has been success translating idiomatic expressions in the past (Fadaee et al., 2018; Zaninello and Birch, 2020), translation alone is not sufficient in educational scenarios. In many past approaches, the task assumes that a model will be provided with at least sentence-level context, which helps models perform well on translation-based tasks. However, there is no guarantee that a live user will give a model an entire sentence or more when seeking to understand an idiomatic expression. Our task tests the robustness of these models to a lack of context. In Appendix F, we outline an experiment where data from human evaluators indicates that 42% of Spanish and 48.5% of German translations remain unnatural, while good translations will be word-for-word in roughly 50% of cases and thus cannot be used as a viable method for this task unless for an idiomatic expression in L_S there is an equivalent expression in L_T .

5.3 Data Contamination

There is always potential for data contamination in LLM experiments. While we cannot make claims about what data the LLMs may have seen in training, we attempt to correlate the explanation generation performance with the number of times an idiom appears online and the likelihood of generating a particular expression from in context examples using an LLM. We find there is **no correlation** (Spearman’s ρ between -0.03 and -0.3) between either proxy measure and model performance. We include a comprehensive report in Appendix G.

6 Human Evaluation

Our main goal is to motivate the need for and inspire design approaches for tools that can assist language learners. As a result, we perform manual evaluation to determine how helpful the best model’s explanations might be to real people. Following Zhang et al. (2023), we measure two qualities of this model’s generated explanations, Fluency and Accuracy. For the highest performing model’s Spanish generations, we ask native speakers to rate both aspects on a Likert scale ranging from 1-5.

Error Type	Examples
Word repetition and redundant language (Fluency)	<ul style="list-style-type: none"> • producir resultados positivos o resultados. [ENG: produce positive results or results] • abandonar a alguien o algo repentinamente y sin previo aviso, a menudo abandonándolos. [ENG: to abandon someone or something suddenly and without notice, usually by abandoning them] • Algo o alguien que es un placer y una vista agradable de ver. [ENG: something or someone who is a pleasure and a (<i>literal translation of idiom</i>) sight worth seeing]
Odd grammar (Fluency)	<ul style="list-style-type: none"> • Aquellos que lleguen o apliquen primero serán los primeros en recibir o ser atendidos. [ENG: Those who arrive or apply first will be the firsts to receive or be seen] – <i>uses verb to receive with no direct object.</i> • Refiriéndose a algo que es el más avanzado o sofisticado de su tipo en un momento dado. [ENG: Referring to something which is the most advanced or sophisticated of its type in a given time.] – <i>wrong co-referent article for ‘algo’, uses masculine form ‘el’ instead of neutral ‘lo’. In English they both translate to ‘the’.</i>
Meaning mismatch (Accuracy)	<ul style="list-style-type: none"> • One’s flesh and blood: Se refiere al cuerpo físico, pero a menudo se usa para enfatizar el aspecto humano de alguien o algo, [ENG: About the physical body, but often used to emphasize the human aspect of something or someone.] – <i>identifies that the body is used as a metaphor, but ties it to humanity rather than family.</i>
Literal meaning (Accuracy)	<ul style="list-style-type: none"> • Dice with death: Participar en una actividad arriesgada o peligrosa que podría resultar potencialmente en la muerte. [ENG: Participate in an activity that could cause death.] – <i>Makes explicit mention to death, rather than seeing it as language for danger.</i>

Table 5: EPIE-ME Common Error Types

The accuracy measure captures the congruence between the generated explanation and the meaning of the idiom. Fluency, however, captures the level of familiarity with the language indicated by the model’s response. There can often be a challenge in describing fluency of a text, especially when short. Thus, we create annotation guidelines inspired by the American Council on the Teaching of Foreign Languages proficiency guidelines (ACTFL, 2024), but modified to suit our specific task. Detailed rubrics can be found in App. B. The generations were rated by two annotators independently. Both annotators are computational linguists and native Spanish speakers from different Spanish-speaking regions: Europe and South America. Both annotators are also fluent English speakers.

Our annotators found that the model performed extremely well in terms of both Fluency and Accuracy, with average scores of 4.70 and 4.78 respectively. To measure inter-annotator agreement, we calculate Krippendorff’s α . These scores range from -1 to 1, where 1 is perfect agreement, 0 is random agreement and -1 is inverse agreement. We found $\alpha = 0.642$ for fluency and $\alpha = 0.417$ for accuracy. We found two main reasons for disagreement. First, in some cases, an explanation may begin with a literal interpretation and end with a figurative interpretation. In this case, one of our annotators appeared to penalize for the inclusion

of the literal interpretation where the other did not. Second, our raters’ native variants of Spanish are different, resulting in differences in understood meaning or connotation of certain words and phrases. However, we do not consider difference in variant a significant issue as it pertains to our overall results. Since the generative models examined in this paper are trained on the same generalized Spanish data that the embedding model is trained on (all ignoring the existence of any variants), the scores given by our annotators ought to be lower than an automated evaluation rather than higher (assuming the automated evaluation is infallible).

6.1 Error Analysis

During human evaluation, our raters pointed out common errors affecting both fluency and accuracy. Issues with fluency come from the model’s tendency to repeat words that may be key to the meaning of the idiomatic expression or make grammatical mistakes. Problems with accuracy arise in cases where the model is either too literal or misidentifies the specific type of entity or domain that an expression refers to. Our annotators also observed that the model might introduce information not necessarily related to the meaning of the expression. Rather, it might indicate the origin of the idiomatic expression, or an indicated stance of the speaker. We see this as a positive addition to

the explanation in educational contexts, but there is inconsistency and we neither ask for nor measure this phenomenon in our experiments. In the context of our educational goals, these findings present a few key issues that we would not wish to see arise in classroom settings. Despite our high fluency and accuracy scores reported in Sec. 6, there is clearly still room to grow before these models are given real responsibility in educational domains.

7 Conclusion and Future Work

This work represents a first step towards generating cross-lingual explanations of idiomatic expressions. We introduced two datasets, one consisting of 628 English idioms and another containing 6218 in English idioms, each with their explanations in English, Spanish and German, and showed that state of the art systems perform well when generating cross-lingual explanations, but not well enough in certain aspects for our ideal target domain of education. For our learned models, we wish to devise more effective strategies that incorporate some notion of directly unwinding figuration and explicit focus on a complete explanation that delivers information so that a learner attains some requisite understanding.

We also showed a disconnect in automated measurements of performance and human judgements with relation to our task. Because the model only generated small samples of text, there is no perfect way to evaluate its overall fluency within a specific language. Although sentence similarity is one way to evaluate accuracy, it does not necessarily focus on whether the core of the idiomatic expression is being represented in the explanation, only how the model represents certain pieces of text.

Finally, although outside the scope of this paper, we can see that some idiomatic expressions are more unconventional than others, but we have no obvious way of measuring or directly coping with this phenomenon. Given the examples *building bridges*, *break a leg*, and *kick the bucket*, clearly the first is more clearly linked to its meaning by metaphor than the latter two, and might even be more easily disambiguated with less context because of the kind of words required to surround it. This is something we would like to measure in the future. Specifically, assuming idiomatic expressions live on this gradient of metaphor, are expressions on one end of the spectrum easier for models to disambiguate than those on the other?

8 Limitations

Given that we cannot make Oxford-ME public, the primary limitation of our paper is the small size of our publishable dataset. However, we are able to provide this data to other researchers upon request. Additionally, we are limited in the number of languages we support. We have only one source language of idiomatic expressions, English. There is a reasonable chance that the models we have experimented with might perform worse in experiments conducted with languages that have smaller amounts of available data. Furthermore, we only attempt cross-lingual explanation into two languages, Spanish and German, which are similarly well supported with data. Finally, the metrics we use in this paper are simply not able to fully capture the phenomenon we wish to measure. We hope to motivate and/or participate in the development of more informative metrics that best suit a task such as ours.

Acknowledgments

We thank members of the BLAST and NALA labs for their feedback. This work utilized the Blanca condo computing resource at the University of Colorado Boulder. Blanca is jointly funded by computing users and the University of Colorado Boulder.

References

- ACTFL. 2024. [Actfl proficiency guidelines 2024](#).
- Amineh Adelnia and Hossein Vahid Dastjerdi. 2011. Translation of idioms: a hard task for the translator. *Theory and practice in language studies*, 1(7):879–883.
- John Ayto. 2020. *The Oxford Dictionary of Idioms*. Oxford University Press.
- Christos Baziotis, Prashant Mathur, and Eva Hasler. 2023. [Automatic evaluation and analysis of idioms in neural machine translation](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3682–3700, Dubrovnik, Croatia. Association for Computational Linguistics.
- Michele Bevilacqua, Marco Maru, and Roberto Navigli. 2020. [Generatory or “how we went beyond word sense inventories and learned to gloss”](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7207–7221, Online. Association for Computational Linguistics.

- F Boers. 2000. [Metaphor awareness and vocabulary retention](#). *Applied Linguistics*, 21(4):553–571.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, and 516 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Fatemeh Enayati and Abbas Pourhosein Gilakjani. 2020. The Impact of Computer Assisted Language Learning (CALL) on Improving Intermediate EFL Learners’ Vocabulary Learning. *International Journal of Language Education*, pages 96–112.
- Marzieh Fadaee, Arianna Bisazza, and Christof Monz. 2018. Examining the tip of the iceberg: A data set for idiom translation. *arXiv preprint arXiv:1802.04681*.
- St. John Fisher. 2016. [The use of technology to support vocabulary development of english language learners](#).
- Artyom Gadetsky, Ilya Yakubovskiy, and Dmitry Vetrov. 2018. [Conditional Generators of Words Definitions](#). *arXiv preprint*. ArXiv:1806.10090.
- Ewa M. Golonka, Anita R. Bowles, Victor M. Frank, Dorna L. Richardson, and Suzanne Freynik. 2014. [Technologies for foreign language learning: a review of technology types and their effectiveness](#). *Computer Assisted Language Learning*, 27(1):70–105.
- Cesar A. Guaqueta and Angela Yicely Castro-Gárces. 2018. [The use of language learning apps as a didactic tool for efl vocabulary building](#). *English Language Teaching*, 11:61–71.
- Shonosuke Ishiwatari, Hiroaki Hayashi, Naoki Yoshinaga, Graham Neubig, Shoetsu Sato, Masashi Toyoda, and Masaru Kitsuregawa. 2019. [Learning to describe unknown phrases with local and global contexts](#). 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 3467–3476, Minneapolis, Minnesota. Association for Computational Linguistics.
- Cunliang Kong, Yun Chen, Hengyuan Zhang, Liner Yang, and Erhong Yang. 2022. [Multitasking framework for unsupervised simple definition generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5934–5943, Dublin, Ireland. Association for Computational Linguistics.
- Murathan Kurfali and Robert Östling. 2020. [Disambiguation of potentially idiomatic expressions with contextual embeddings](#). In *Proceedings of the Joint Workshop on Multiword Expressions and Electronic Lexicons*, pages 85–94, online. Association for Computational Linguistics.
- Murathan Kurfali and Robert Östling. 2020. [Disambiguation of potentially idiomatic expressions with contextual embeddings](#). In *Joint Workshop on Multiword Expressions and Electronic Lexicons, Barcelona, Spain (Online), December 13, 2020*, pages 85–94.
- Zoltán Kövecses and Péter Szabó. 1996. [Idioms: A view from cognitive semantics](#). *Applied Linguistics*, 17.
- Vladimir I. Levenshtein. 1965. [Binary codes capable of correcting deletions, insertions, and reversals](#). *Soviet physics. Doklady*, 10:707–710.
- Shuang Li, Jiangjie Chen, Siyu Yuan, Xinyi Wu, Hao Yang, Shimin Tao, and Yanghua Xiao. 2024. [Translate meanings, not just words: Idiomkb’s role in optimizing idiomatic translation with language models](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(17):18554–18563.
- Emmy Liu, Aditi Chaudhary, and Graham Neubig. 2023. [Crossing the threshold: Idiomatic machine translation through retrieval augmentation and loss weighting](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15095–15111, Singapore. Association for Computational Linguistics.
- Ke Ni and William Yang Wang. 2017. [Learning to Explain Non-Standard English Words and Phrases](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 413–417, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Thanapon Noraset, Chen Liang, Larry Birnbaum, and Doug Downey. 2017. [Definition modeling: Learning to define word embeddings in natural language](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI’17*, page 3259–3266. AAAI Press.
- Geoffrey Nunberg, Ivan A. Sag, and Thomas Wasow. 1994. [Idioms](#). In Stephen Everson, editor, *Language*, pages 491–538. Cambridge University Press.
- OpenAI. 2023. GPT-3.5 Turbo. <https://platform.openai.com/docs/models/gpt-3-5-turbo>. [Online; accessed December 2023 to time of publication].
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Jipeng Qiang, Yang Li, Chaowei Zhang, Yun Li, Yi Zhu, Yunhao Yuan, and Xindong Wu. 2023. [Chinese idiom paraphrasing](#). *Transactions of the Association for Computational Linguistics*, 11:740–754.

- Elaheh Rafatbakhsh and Alireza Ahmadi. 2019. A thematic corpus-based study of idioms in the corpus of contemporary american english. *Asian-Pacific journal of second and foreign language education*, 4(1):1–21.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Nils Reimers and Iryna Gurevych. 2020. [Making monolingual sentence embeddings multilingual using knowledge distillation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Prateek Saxena and Soma Paul. 2020. [EPIE Dataset: A Corpus For Possible Idiomatic Expressions](#). *arXiv preprint*. ArXiv:2006.09479.
- Uliana Sentsova, Debora Ciminari, Josef Van Genabith, and Cristina España-Bonet. 2025. [MultiCoPIE: A multilingual corpus of potentially idiomatic expressions for cross-lingual PIE disambiguation](#). In *Proceedings of the 21st Workshop on Multiword Expressions (MWE 2025)*, pages 67–81, Albuquerque, New Mexico, U.S.A. Association for Computational Linguistics.
- Zhe Wang Tao Hao and Yuliya Ardasheva. 2021. [Technology-assisted vocabulary learning for efl learners: A meta-analysis](#). *Journal of Research on Educational Effectiveness*, 14(3):645–667.
- Keiichiro Tashima, Hirohisa Aman, Sousuke Amasaki, Tomoyuki Yokogawa, and Minoru Kawahara. 2018. [Fault-prone java method analysis focusing on pair of local variables with confusing names](#). In *2018 44th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*, pages 154–158.
- Andrea Tyler and Vyvyan Evans. 2003. [Reconsidering prepositional polysemy networks: the case of over*](#). De Gruyter Mouton, Berlin, New York.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Andrea Zaninello and Alexandra Birch. 2020. [Multiword expression aware neural machine translation](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 3816–3825, Marseille, France. European Language Resources Association.
- Hengyuan Zhang, Dawei Li, Yanran Li, Chenming Shang, Chufan Shi, and Yong Jiang. 2023. [Assisting language learners: Automated trans-lingual definition generation via contrastive prompt learning](#). In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 260–274, Toronto, Canada. Association for Computational Linguistics.
- Ruofei Zhang and Di Zou. 2022. [Types, purposes, and effectiveness of state-of-the-art technologies for second and foreign language learning](#). *Computer Assisted Language Learning*, 35(4):696–742.
- Yihua Zhang. 2011. Discussion on the definitions in chinese learner’s dictionaries: Comparative study of domestic and foreign learner dictionaries (translated from chinese). *Chinese Teaching in the World*.

A Hyperparameters and Other Tables

For T5 and mT5, we use a learning rate of $1e-5$, the AdamW optimizer, a batch size of 32, and run for 1000 epochs with early stopping. For GPT and Llama, we also explore two different configurations to constrain the length of our explanation to a maximum of one or two sentences. Additionally, for GPT 3.5 Turbo we set the temperature to 0.5 for all scenarios, and Llama was left with its default temperature. All hyperparameters were chosen using the development set.

Model	Overall		Spanish		German	
	BLEU	ROUGE	BLEU	ROUGE	BLEU	ROUGE
mT5 Direct	0.37	9.94	0.34	10.59	0.47	9.34
mT5 Pipeline	0.09	4.90	0.18	9.71	0.29	7.80
T5 Pipeline	0.89	11.95	1.03	13.09	0.74	10.73
Llama Direct Zero-Shot	1.00	12.31	1.12	14.19	0.87	10.48
Llama Pipeline Zero-Shot	1.40	14.69	1.67	16.54	1.08	12.88
GPT Direct Zero-Shot	3.58	22.1	3.66	24.39	3.50	19.87
GPT Pipeline Zero-Shot	4.59	26.41	5.56	28.21	3.57	24.63
GPT Direct 5-Shot	4.04	23.23	4.22	25.09	3.86	21.34
GPT Pipeline 5-Shot	4.19	24.76	4.88	26.38	3.47	23.13

Table 6: Best Model Results on EPIE-ME Using BLEU and ROUGE for Direct and Pipeline Strategies. LLM Results are Averaged Across 3 Runs on the Test Set.

Language	T5	mT5
Overall	48.4	41.60
Spanish	48.23	41.49
German	48.57	42.02
Pre-Translation		
English	44.83	41.14

Table 7: EPIE-ME T5 & mT5 Dev Set Results

Model	Sim	BLEU	ROUGE
mT5	37.45	0.31	11.82
T5	42.04	0.46	13.58
Llama Zero-Shot	57.66	1.82	18.11
GPT Zero-Shot	71.05	5.59	29.87
GPT 5-Shot	72.57	4.7	28.1

Table 8: EPIE-ME Pipeline Model Performance Before Translation (EN output, Test Set)

GPT						
	Direct			Pipeline		
Shots	1	3	5	1	3	5
Known	65.86	67.50	67.75	68.26	68.31	68.62
Induced	65.13	66.16	67.52	67.77	68.09	68.57
Random	65.25	67.08	67.45	67.81	68.59	68.79

Table 9: EPIE-ME Few-Shot Results by Demo Selection Strategy

B Human Evaluation Guide

We provide the rubrics given to our annotators for rating both fluency and accuracy in Tabs. 10 and 11. The fluency rubric is inspired in-part by ACTFL (2024).

Rating	Description
1	Little to No Understanding: <ul style="list-style-type: none"> Has a severely limited vocabulary, acquiring only a handful of words in the language. No understanding of grammatical constructs. Is not in the correct language.
2	Beginner Language Learner: <ul style="list-style-type: none"> Still includes some words from their native language. Basic vocabulary - only uses words a child would understand. Sentence structure that mimics the native language.
3	Intermediate Language Learner: <ul style="list-style-type: none"> Understanding of standard grammatical structures. Inefficient communication of ideas - vocabulary is clearly more sizable, but word choice is somewhat odd or ineffective. Perhaps gets to their point in a roundabout fashion.
4	Advanced Speaker: <ul style="list-style-type: none"> Strong vocabulary. Strong understanding of tense morphology and usage. Grammar is near perfect - capable of more complex constructs.
5	Fluent/Native Speaker: <ul style="list-style-type: none"> Language production as expected of someone with a complete mastery of the language. Communicates efficiently and precisely (inclusion of examples to provide clarity should not be seen as inefficient communication). Produces language as if communicating with another fluent speaker.

Table 10: Fluency Rubric

Rating	Description
1	Completely Wrong: <ul style="list-style-type: none"> Language may also be too incoherent to determine.
2	Literal Interpretation: <ul style="list-style-type: none"> Does not understand the figurative nature of the expression, and tries to reference the lexical form when inappropriate (i.e., mentioning chickens in reference to "putting all of your eggs in one basket").
3	Domain Matches but Not Meaning: <ul style="list-style-type: none"> Identifies the general context in which the expression is to be used, but the meaning is incorrect (i.e., has some sense of related topics, but cannot identify sentiment or value).
4	Good Sense of Meaning: <ul style="list-style-type: none"> Near-complete understanding of the expression or the value which it communicates, still referencing the main ideas found in the gold explanation.
5	Perfectly Captures Meaning: <ul style="list-style-type: none"> Very high congruence with the gold explanation.

Table 11: Accuracy Rubric

C Themes

We include the full list of themes from Rafatbakhsh and Ahmadi (2019) and their frequencies within our dataset (in descending order) in Tabs. 12 and 13.

Theme	Num. Idioms
Language, speech, and conversation	29
Happiness, pleasure, and enjoyment	22
Health and illness	22
Argument and conflict	20
Cooperation	19
Experience	19
Duty and responsibility	18
Success	16
Haste and speed	15
Anger and annoyance	14
Hope and optimism	14
Caution	13
Expense	13
Futility	13
Honesty	13
Power	13
Time	13
Anxiety and worry	12
Crime and punishment	12
Work and employment	12
Appearance	11
Misfortune and adversity	11
Thoroughness	11
Change	10
Chaos and disorder	10
Love	10
Preparation and readiness	10

Table 12: EPIE-ME Idiomatic Expression Count (≥ 10) per Theme

Theme	Num. Idioms
Action	9
Deception and lying	9
Excess and extravagance	9
Poverty	9
Reputation and fame	9
Travel and transport	9
Danger	8
Age	7
Certainty	7
Doubt and uncertainty	7
Intelligence and knowledge	7
Secrecy	7
Beauty	6
Critics and criticism	6
Foresight and the future	6
Forgiveness and reconciliation	6
Opportunity	6
Self-interest	6
Courage	5
Crisis	5
Family	5
Fools and foolishness	5
Surprise	5
Unhappiness and disappointment	5
Bribery, corruption, and extortion	4
Indecision and prevarication	4
Laziness	4
Youth	4
Death	3
Equality	3
Fate and chance	3
Justice	3
Marriage	3
Mistakes	3
Money, wealth, and prosperity	3
Revenge and retribution	3
Strength	3
Violence	3
Ambition	2
Boastfulness and conceit	2
Class	2
Embarrassment, shame, and humiliation	2
Food	2
Gossip and rumor	2
Madness	2
Traitors and treachery	2
Warfare	2
Weakness	2
Weather	2
Jealousy and envy	1
Pregnancy	1
Clothes	0
Debt	0
Hypocrisy	0

Table 13: Cont’: EPIE-ME Idiomatic Expression Count (< 10) per Theme

D Example Prompts

We provide all prompts used for our LLM experiments in Tab. D. The prefixes were used when exploring configurations. The contextual enhancement prompts are shown in zero-shot format, but were also used in the k -shot scenario. We assume

the reader can infer the k -shot form based on the other k -shot templates.

Prompt Type	Text
Length control prefix	<ul style="list-style-type: none"> • Single sentence: Provide responses no longer than one sentence. • Two sentences max: Provide responses no longer than two sentences.
Zero-Shot	<ul style="list-style-type: none"> • Direct: Explain in {lang} the meaning of the idiomatic expression '{idiom}'. • Pipeline-Explain: Explain the meaning of the idiomatic expression '{idiom}'. • Pipeline-Translate: Translate the following into {lang}: '{explanation}'. Respond with nothing but the translation.
Contextual Enhancements	<ul style="list-style-type: none"> • Categorical information: Explain in {lang} the meaning of the idiomatic expression '{idiom}' in the context of '{category}'. • Sentence-level context: Explain the meaning of the idiomatic expression '{idiom}' given the sentence '{sentence}'. • Both: Explain in {lang} the meaning of the idiomatic expression '{idiom}' in the context of '{category}' and given the sentence '{sentence}'.
k -Shot	<p>All are repeated k times with a $k + 1$th occurrence having the idiom to be explained followed by A:</p> <ul style="list-style-type: none"> • Direct: Q: Explain in {lang} the meaning of the idiomatic expression '{idiom}'. A: {explanation}. • Pipeline-Explain: Q: Explain the meaning of the idiomatic expression '{idiom}'. A: {explanation}. • Pipeline-Translate: Q: Translate the following into {lang}: '{explanation}'. A: {explanation}.
Fill-in-the-Blank	<ul style="list-style-type: none"> • System: Fill in the appropriate idiomatic expression that completes the sentence below. • User: {context} Response: [blank] can be replaced with:
Oxford-ME Context Generation	<ul style="list-style-type: none"> • Use the idiom [idiom] meaning [explanation] in a sentence. Respond with nothing but the sentence.

Table 14: LLM Prompts

E Noise in Translations

Where EPIE-ME needed a small amount of correction, Oxford-ME required significantly less modification.⁵ We believe this is primarily due to the difference in English explanation length between the two sources of data. Examining the quartiles of normalized edit distance, we see that 50% of the data required roughly less than a 14% modification to the Spanish explanations in EPIE-ME, less than a 27% modification for German in EPIE-ME, less than an 8% modification for Spanish explanations in Oxford-ME, and less than a 4% modification for German in Oxford-ME. Overall, quite low. The high standard deviations can be explained by the lack of samples near the mean, rather a large portion of samples are in the lower quartiles. These measurements can be found in Tab. 15. Graphs of the same information can be found in Fig. 1.

F Translation as a Method

To test our assumption that translation is not a wholesale solution, we ask GPT-3.5 Turbo to provide only a translation of the idiomatic expressions in our test set. When comparing the semantic similarity of the generated translation to the gold explanation, we see that this approach appears to perform better than that of the fine-tuned models used in our experiments in Tab 3, earning scores of 45.80 and 44.21 for Spanish and German respectively. However, we asked our annotators to take a closer look at the outputs to make note of 3

⁵For Oxford-ME, we verify only the Spanish translations for this analysis.




Figure 1: Histograms of Normalized Edit Distance vs Percentage of Data with Bins of Size 0.05 Over Both Languages in EPIE-ME and Oxford-ME.

Statistic	EPIE-ME		Oxford-ME	
	Spanish	German	Spanish	German
Mean	0.159	0.283	0.131	0.183
Min	0	0.092	0	0
Max	0.781	0.857	0.8	0.889
Stdev	0.162	0.216	0.165	0.241
Q1	0	0.092	0	0
Q2	0.131	0.269	0.071	0.043
Q3	0.249	0.440	0.205	0.307
Q4	0.781	0.857	0.839	.89

Table 15: Noise in the Data Measured by Normalized Edit Distance with Noise by Quartile in the Lower Portion.

things for each sample: *Is the translated expression good/natural? Is the translation done word-for-word? Is this translation an instance of an idiom in your native language?*

Using these 3 factors, we analyze how often a sample is above a threshold of 40⁶ while still being judged as a poor/unnatural expression, which was the case for 14.5% and 18.5% of our test set in

⁶We found the scores had incredibly high standard deviations of 20.52 and 19.84 respectively, so we set this threshold at a rough 10% under the mean.

Spanish and German respectively. While semantic similarity is not meant to represent the actual quality of a sample, this further indicates that the similarity score we obtain cannot be the only metric used for evaluating models on this task. Of the total number of expressions that were good/natural, the proportion that were word-to-word translations were 48.3% and 56.3%, where 71.4% and 93.1% of those subsets were a case where the idiom was also present in the target language. This tells us that translation of just an idiomatic expression is not likely to generalize, as one will often need the same idiom to be present in the target language. These findings also hint at the potential that models trained to capture multilingual sentence similarity in their embedding spaces are more strongly capturing word-to-word similarity than meaning similarity in the case of idiomatic expressions.

G Data Contamination

Some may be skeptical of our experiments and the potential for data contamination in our LLMs. Specifically, that they may have seen the idiom-explanation pairs within our dataset. While we do

Correlation with Similarity Scores		
Model / Proxy	ρ	τ -b
GPT vs SR	-0.181	-0.123
Llama vs SR	-0.293	-0.205
Llama vs GP	-0.030	-0.021
Correlation with Human Evaluation		
Metric / Proxy	τ -b	
Fluency vs SR	-0.141	
Accuracy vs SR	-0.067	
Fluency vs GP	-0.159	
Accuracy vs GP	0.072	

Table 16: EPIE-ME Model Similarity vs Search Results (SR) and Generation Probabilities (GP)

not have access to the training data of either GPT or Llama, we will assume that they have seen almost the entire internet for the purposes of determining the level of contamination. Under this assumption, if the model is performing well on certain idiomatic expressions, we would expect those to be the most used across the internet. For both GPT and Llama, we attempt to correlate both the similarity score and human evaluation results per expression with the number of search results returned by the Google API when using the expression as input – not accounting for morphological shift or interjecting words. For Llama, we additionally extract the probability of generating an expression given 10 in-context examples via fill-in-the-blank prompting.⁷ Our similarity score results are correlated with the search results and generation probabilities using Spearman’s rank correlation coefficient (ρ) and Kendall’s rank correlation coefficient, variant b (τ -b). Because of the many ties in our human evaluation results, we correlate only with τ -b. In every single case, we see a score very close to 0, indicating no association whatsoever (Tab. 16). We further examine our results by quartile of search result hits. There is no clear upward trend as search results increase. In fact, the Llama results show an inverse relationship, the Fluency results decrease until the final quartile, and the GPT and Accuracy results show a bell curve peaking in the second quartile (Tab. 17).

Metric	Q1	Q2	Q3	Q4
GPT Similarity	71.76	72.60	69.71	67.62
Llama Similarity	63.40	61.99	60.85	55.81
Fluency	4.85	4.70	4.58	4.67
Accuracy	4.78	4.86	4.80	4.69

Table 17: EPIE-ME Results by Quartile of Search Results

⁷Prompts found in Tab. 14