

PolyFrame at MWE-2026 AdMIRe 2: When Words Are Not Enough: Multimodal Idiom Disambiguation

Nina Hosseini-Kivanani
University of Luxembourg & RTL
Esch-sur-Alzette, Luxembourg
nina.hosseinikivanani@uni.lu

Abstract

Multimodal models struggle with idiomatic expressions due to their non-compositional meanings, a challenge amplified in multilingual settings. We introduced PolyFrame, our system for the MWE-2026 AdMIRe 2 shared task on multimodal idiom disambiguation, featuring a unified pipeline for both image+text ranking (Subtask A) and text-only caption ranking (Subtask B). All model variants retain frozen CLIP-style vision–language encoders and the multilingual BGE M3 encoder, training only lightweight modules: a logistic regression and LLM-based sentence-type predictor, idiom synonym substitution, distractor-aware scoring, and Borda rank fusion. Starting from a CLIP baseline (26.7% Top-1 on English dev, 6.7% on English test), adding idiom-aware paraphrasing and explicit sentence-type classification increased performance to 60.0% Top-1 on English, and 60.0% Top-1 (0.822 NDCG@5) in zero-shot transfer to Portuguese. On the multilingual blind test, our systems achieved average Top-1/NDCG scores of 0.35/0.73 for Subtask A and 0.32/0.71 for Subtask B across 15 languages. Ablation results highlight idiom-aware rewriting as the main contributor to performance, while sentence-type prediction and multimodal fusion enhance robustness. These findings suggest that effective idiom disambiguation is feasible without fine-tuning large multimodal encoders.

1 Introduction

Idiomatic expressions are a long-standing problem for computational semantics because their figurative meanings are not compositionally predictable from surface form (Flor et al., 2025; Zeng and Bhat, 2021). Models trained primarily on literal data often interpret idioms word by word, which leads to semantic drift in downstream tasks such as translation, retrieval, and captioning (Pickard

et al., 2025a). This effect is amplified in multilingual settings, where idioms differ in lexicalisation and transparency across languages (Domhan et al., 2022; Cap et al., 2015). Recent work treats idiomaticity as a structured prediction problem, for example, by classifying usages as literal or idiomatic or by ranking contextually appropriate interpretations (Endaliev et al., 2023).

AdMIRe 2.0 provides sentences containing potentially idiomatic nominal compounds together with five candidate images and parallel caption-based variants (Arslan et al., 2026). Subtask A asks systems to rank images according to how well they reflect the intended meaning of the compound in context. Subtask B replaces images with captions and prompts to probe text-only reasoning under the same schema. Both subtasks implicitly require robust sentence-type judgements (idiomatic vs. literal) and fine-grained alignment of the chosen sense with visual or textual cues.

This paper presents our system description for AdMIRe 2.0 (Arslan et al., 2026). We build a shared data and evaluation pipeline and instantiate three variants within a single architecture: (i) a CLIP-based multimodal baseline for Subtask A; (ii) an improved multimodal ranker with supervised sentence-type classification (logistic regression plus a literal-first LLM classifier), idiom synonym replacement, distractor-aware scoring, and fusion of visual and caption-based scores; and (iii) a text-only counterpart of the improved ranker for Subtask B (Figure 1). All systems use frozen CLIP-style vision–language encoders and the multilingual BGE M3 sentence encoder (Chen et al., 2024), with learning restricted to lightweight idiomaticity classifiers, scalar fusion weights, and rank-level Borda ensembles over multiple CLIP and LLM backends. Our experiments examine how sentence-type prediction, idiom-aware rewriting, and multimodal fusion jointly improve idiom-sensitive image and caption ranking across languages.

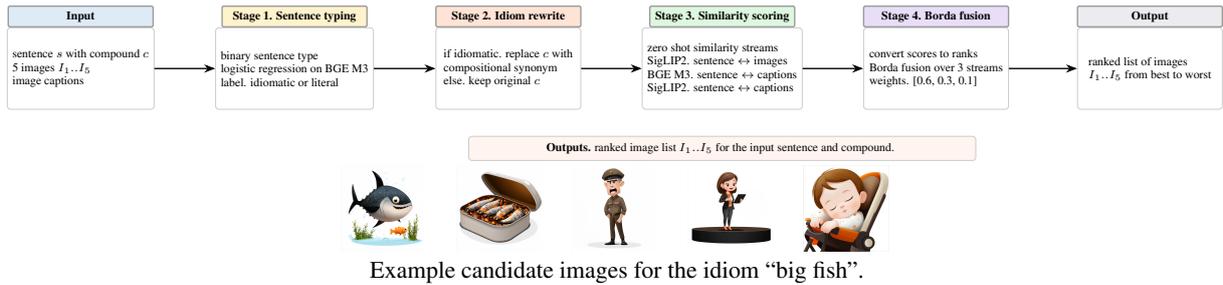


Figure 1: Overview of the final POLYFRAME pipeline. Sentence typing via logistic regression, idiom replacement for idiomatic cases, three zero shot similarity streams with SigLIP2 and BGE M3, and Borda fusion

2 Related Work

Idiom processing and idiomaticity detection.

Text-based idiom processing has been studied extensively through shared tasks and benchmarks that ask models to decide whether a target expression is used literally or idiomatically in context (e.g., (Jakhotiya et al., 2022; Boisson et al., 2022)). SemEval 2022 Task 2 introduced multilingual idiomaticity detection for verbal multiword expressions and showed that glosses, translations, and lexical knowledge bases can improve cross-lingual disambiguation (Phelps et al., 2022). Subsequent work has evaluated large language models on idiom detection and translation, reporting that they still over-predict idiomatic readings in literal contexts and require careful prompting to approach the performance of supervised baselines (Phelps et al., 2024). Other studies have proposed targeted metrics for evaluating idiom translation in neural machine translation and documented a strong tendency toward literal renderings (Khan and Akter, 2025). The AdMIRE shared tasks extend this line of work by framing idiomaticity as a multimodal ranking problem and by providing graded relevance annotations for sentence type and image suitability (Torunoğlu-Selamet et al., 2026; Pickard et al., 2025b).

Multimodal idiom understanding and vision-language models. Multimodal encoders such as CLIP learn aligned image–text representations via large-scale contrastive learning and have become standard backbones for zero-shot classification and retrieval. However, their pretraining data and objectives are not tailored to figurative language, and idiomatic usages are often under-represented. Recent approaches adapt CLIP to idiom interpretation by enriching prompts with natural-language definitions, few-shot examples, or idiomatic paraphrases, and by combining visual similarity with textual sen-

tence embeddings (Markchom et al., 2025; Wang et al., 2025). AdMIRE 2.0 offers a test bed for such strategies by providing both image-based and caption-based variants of the same task. Our systems follow this direction: we keep OpenCLIP-style encoders frozen, pair them with BGE M3 for multilingual sentence embeddings, and focus learning capacity on lightweight idiomaticity classifiers and fusion mechanisms that combine visual, caption-based, and definition-based scores within a single architecture.

3 Methodology¹

We describe our systems for the two AdMIRE 2.0 subtasks, which jointly probe sentence-level idiomaticity and multimodal disambiguation (see Figure 1 for an overview). Subtask A presents a context sentence containing a potentially idiomatic nominal compound and five candidate images; the system must rank the images by appropriateness. Subtask B reuses the same TSV schema but replaces images with captions and prompts, yielding a text-only ranking task. Sentence-type prediction is evaluated by accuracy, and ranking by Top-1 accuracy and NDCG@5.

3.1 Task setup and system variants

For Subtask A, the organisers provide supervised splits for English (EN) and Portuguese (PT) with train, development, test, and cross-evaluation partitions. Each instance consists of a sentence, the target compound, five image file names organised under idiom-specific directories, and optionally image captions and prompts. For evaluation, we run our system on the official blind-test bundles covering 15 languages, resolving image paths via a language-aware global root and tracking the number of processed instances per language (approximately 50–360) to ensure full coverage.

¹<https://github.com/NinaKivanani/PolyFrame>

For Subtask B, we reuse the TSV structure but operate in text-only mode: image identifiers are retained for compatibility, while ranking uses captions and prompts only. Within a shared architecture, we instantiate three system variants to study the impact of modality, sentence-type modelling, and idiom replacement: (i) a CLIP-based multimodal baseline for Subtask A, (ii) an improved multimodal ranker with explicit sentence-type classification, idiom replacement, and distractor-aware fusion, and (iii) a text-only variant of the improved ranker for Subtask B.

3.2 Representations and encoders

Each TSV row is parsed into a context sentence x , a target compound c , five image identifiers $\{i_1, \dots, i_5\}$, and, when available, image captions $\{c_k\}_{k=1}^5$. Supervised splits additionally provide a gold sentence type $y_{\text{sent}} \in \{\text{idiomatic}, \text{literal}\}$ and a gold ranking $\pi^* = [\pi_1^*, \dots, \pi_5^*]$.

All systems use vision-language encoders for multimodal understanding. The baseline uses a multilingual CLIP variant (xlm-roberta-large-ViT-H-14) with frozen LAION-5B weights. The improved systems employ a dual-encoder architecture: SigLIP2 (ViT-SO400M-14-SigLIP2-378) (Tschannen et al., 2025) as the primary vision-language encoder for image-text similarity, combined with BGE-M3 as a separate multilingual text encoder for text-only ranking via caption matching. Rankings from both encoders are fused using Borda-count aggregation. Additional CLIP or SigLIP models can be added in ensemble mode with weighted fusion.

For sentence and caption representations we additionally use BGE M3, a multilingual dense retrieval encoder. Given a sentence-compound pair (x, c) , we concatenate them and obtain an embedding $g(x, c) \in \mathbb{R}^d$, which provides text-only similarity scores for caption-based ranking and serves as the feature space for a lightweight binary sentence-type classifier.

3.3 Sentence-type prediction and ranking

Sentence-type prediction. Sentence-type prediction combines a supervised classifier and an LLM-based component. We train a two-way logistic regression model on EN data using BGE M3 features $g(x, c)$ to predict $y_{\text{sent}} \in \{\text{literal}, \text{idiomatic}\}$. This serves as the primary sentence-type signal when available.

As a complementary mechanism and cross-lingual fallback, we use literal-first prompting with

external LLMs (GPT-4o, Qwen3-32B, Llama 3.1-70B, DeepSeek-v3, Mistral). For each compound-language pair, the system first asks the LLM to generate a small set of clearly literal example sentences, which are cached. The final prompt presents these examples together with the target sentence and instructions to decide between LITERAL and IDIOMATIC. We allow brief reasoning but constrain the output to a single label.

Baseline multimodal ranker. The baseline first predicts sentence type heuristically, then constructs prompts, and finally performs visual and textual ranking. Given the predicted type, we build a small set of text queries that combine the sentence, the compound, short natural-language definitions from a hand-curated idiom lexicon, and few-shot examples. These queries are encoded with the CLIP text tower and averaged to obtain a single query embedding q .

For visual ranking, each image i_k is mapped to an embedding $v_k = f_{\text{img}}(i_k)$, and similarity scores are computed via cosine similarity between q and v_k , optionally scaled by a temperature τ to sharpen the distribution over the five candidates; images are then ranked by descending similarity. Caption-based ranking is analogous, using captions instead of images. Visual and caption-based rankings are fused via a Borda-style scheme that converts rank positions into scores and interpolates them, typically with a higher weight on visual information.

Improved multimodal ranker. The improved ranker extends the baseline with supervised sentence-type classification, idiom synonym replacement, and distractor-aware fusion. For sentences classified as idiomatic, we rewrite the context by replacing occurrences of c with a compositional paraphrase that makes the figurative meaning explicit; the modified sentence is then used to construct CLIP prompts, while literal sentences are left unchanged.

The final ranking fuses three score streams: (1) SigLIP2 vision-language similarity (sentence \leftrightarrow images), (2) BGE-M3 text similarity (sentence \leftrightarrow captions), and (3) SigLIP2 text similarity (sentence \leftrightarrow captions via SigLIP2’s text encoder). In our submitted configuration we use weighted Borda-count aggregation with fixed weights of [0.6, 0.3, 0.1] for vision-language similarity, BGE-M3 text similarity, and SigLIP2 caption similarity, respectively, in image+text mode. For text-only mode, weights are adjusted to [0.0, 0.7, 0.3] to focus on text components while ignoring vision. A confidence measure

derived from the gap between the top two scores in each stream can adjust these weights. For non-English languages, an optional cross-lingual mode combines scores computed on the original sentence with scores computed on an English translation obtained from the same LLM used for literal-first classification.

Text-only ranker. The text-only system reuses the same architecture in caption-only mode. Image embeddings are set to zero and receive zero weight in the fusion stage; sentence-type prediction, idiom synonym replacement, and caption-based scoring with BGE M3 and SigLIP2 remain unchanged. This design ensures that improvements over random ranking are attributable purely to text modelling of captions and that ablations of the classifier, literal-first prompting, or idiom replacement are directly comparable across subtasks.

3.4 Ensembles, transfer strategies

Model ensembles and Borda fusion. To improve robustness and support zero-shot transfer, we employ ensembles at both encoder and classifier level with a unified rank-level fusion scheme. On the vision side, we optionally use an ensemble of CLIP variants. Our best-performing configuration uses a single SigLIP2 model (ViT-SO400M-14-SigLIP2-378) rather than an ensemble, as this provided optimal performance while maintaining efficiency. The system supports ensemble configurations where multiple CLIP or SigLIP variants can be combined using weighted Borda fusion, but the submitted system operated in single-model mode.

For sentence-type prediction, we evaluate several LLMs under the same literal-first prompting strategy; in the final configuration, we combine Qwen3-32B and GPT-4o with weights $[0.6, 0.4]$, selected based on development performance, while other LLMs act as fallbacks. In all cases, we apply weighted Borda fusion: each candidate’s rank in each list is converted into a score (higher for better ranks), multiplied by the corresponding fusion weight, and summed. Candidates are then re-ranked by their aggregated Borda scores, providing a consistent mechanism for combining both modalities and model variants.

Zero-shot and few-shot transfer. To handle idioms and languages with limited training data, we support both zero-shot and few-shot strategies. When few-shot prompting is enabled, CLIP queries are augmented with hand-curated examples that illustrate literal and idiomatic usage (e.g., idiomatic

big fish in a company vs. literal *big fish* in a fishing context), which are prepended to CLIP text queries to guide the vision–language encoder. For languages without training data, we rely on the multilingual capabilities of SigLIP2 and BGE M3: the system applies zero-shot classification directly to non-English sentences using the multilingual encoders, and when this fails, the system falls back to simple heuristics based on compound frequency in captions and basic lexical markers. This zero-shot setup yields reasonable performance (around 60% accuracy on PT) without language-specific training.

4 Results & Discussion

We evaluate our systems using the official AdMIRE 2.0 metrics, reporting Top-1 accuracy and NDCG@5. Experiments cover both subtasks and include baseline comparisons, strategy-specific variants, ablation studies, and the official blind-test evaluation on Codabench. We used a temperature of $\tau = 0.7$ for similarity scaling and Borda-count-based fusion as the default for final evaluation.

The CLIP-based baseline, which relies on frozen vision–language embeddings and heuristic sentence typing, achieves 26.7% Top-1 accuracy on the EN development set and 6.7% on the test set (see Table 1). This large drop highlights the difficulty of the task and confirms that naive multimodal retrieval is strongly biased toward literal interpretations when idiomatic compounds are present.

System	EN dev		EN test		PT dev (zero-shot)	
	Top-1 (%)	NDCG@5	Top-1 (%)	NDCG@5	Top-1 (%)	NDCG@5
CLIP baseline	26.7	0.655	6.7	0.607	–	–
+ Idiom replacement	60.0	0.800	–	–	–	–
+ LLM sentence typing	40.0	0.739	–	–	–	–
+ All improvements	60.0	0.797	–	–	60.0	0.822

Table 1: System performance across evaluation sets.

Introducing idiom synonym replacement yields the largest single improvement, raising Top-1 accuracy on the EN development set to 60.0% (Table 1). This confirms observations from AdMIRE 2024 that rewriting idiomatic expressions into compositional paraphrases substantially reduces literal bias in vision–language models. Literal-first sentence classification using an external LLM improves performance to 40.0% when used in isolation, indicating that sentence-type awareness is beneficial but insufficient without explicit paraphrasing. The best submission didn’t use LLM classification. Combining sentence-type prediction, idiom replacement, and multimodal fusion in the improved ranker

System	Top ₁ A	Top ₁ B	ST Acc.
CLIP baseline (XLM-RoBERTa-ViT-H-14)	26.7	6.7	53.3
GPT-4o	32.5	28.0	58.0
Qwen3-32B	31.0	26.5	56.5
Llama3.1-70B	29.5	25.0	55.0
DeepSeek-v3	28.0	23.5	53.5
Mistral-large	30.0	25.5	54.5
Ensemble (Qwen3+GPT-4o)	33.5	29.0	59.0
Best submission	35.0	32.0	-

Table 2: Multilingual blind test results (15 languages). Best submission used SigLIP2 + idiom replacement with zero-shot classification.

maintains 60.0% Top-1 accuracy while providing greater robustness across prompt conditions and idiom types. Zero-shot transfer to PT also reaches 60.0% Top-1 and 0.822 NDCG@5, suggesting that the core strategies generalise across related languages.

Multilingual blind-test results. We submitted our systems to the official Codabench evaluation for both subtasks. On the multilingual blind test covering 15 languages, our submission achieved an average Top-1 accuracy of 0.35 and an NDCG of 0.73 for Subtask A, and 0.32 Top-1 accuracy with an NDCG of 0.71 for Subtask B. These scores correspond to a mid-range ranking among participating systems at the time of evaluation. Performance was consistent across languages, with comparable scores on Chinese and non-Chinese subsets, indicating stable cross-lingual behaviour in a fully zero-shot setting (Table 2).

For the blind-test phase, we apply the improved system unchanged to all 15 evaluation languages, using the official directory structure and file formats. The submitted configuration combines SigLIP2 with idiom replacement and zero-shot sentence-type classification. We verify full coverage for each language, with sample counts ranging from 48 to 363 instances (see Table 3 in the Appendix).

Ablation studies. Ablation experiments on the EN development set show that idiom replacement accounts for most of the performance gains, while sentence-type prediction mainly improves stability and reduces variance across idioms. Multimodal fusion consistently outperforms unimodal ranking, and Borda-style aggregation is more robust than alternative rank-combination strategies across hyperparameter settings. Temperature and fusion-weight sweeps show limited sensitivity around the chosen

defaults, indicating that the improvements are not driven by fragile tuning.

Taken together, the results suggest that most of the attainable gains come from better framing of the vision–language matching problem rather than from stronger back-end language models (Radford et al., 2021). Idiom rewriting and multimodal fusion close much of the gap between the CLIP baseline and our best systems (Gao et al., 2024), while LLM-based sentence-type prediction yields smaller incremental gains at the cost of additional latency, API dependence, and prompt sensitivity (Jin et al., 2025). This trade-off is reflected in the final Codabench submission, which uses SigLIP2 with idiom replacement and zero-shot sentence typing, but omits LLM classification despite its slightly higher scores on the EN development set. The relatively modest drop from Subtask A to Subtask B and the stable NDCG values across languages indicate that caption-based reasoning can approximate image-based disambiguation when captions are informative, but also highlight persistent weaknesses on low-resource and culturally marked varieties such as Spanish Ecuador and Uzbek (Tschannen et al., 2025). Overall, the pattern of ablation results supports our design choice to prioritise frozen encoders and lightweight, transparent adaptations over heavy fine-tuning or tightly coupled LLM components (Xing et al., 2024).

Limitations and Future Work

A key limitation of our approach is that the idiom synonym replacement database is manually curated and currently covers only around 50 common English idioms, with limited coverage for the 14 non-English languages in the blind test. This restricts the effectiveness of idiom replacement for less frequent expressions and for many non-English instances. Future work will explore learned fusion weights as an alternative to our current fixed Borda-based scheme. Since our experiments so far rely mainly on decoder-style large language models, we also plan to incorporate encoder-based models (e.g., BERT) and encoder–decoder architectures (e.g., T5, mBART). These architectures may capture idiomaticity in different ways and provide complementary perspectives on representation learning for figurative language.

References

- Doğukan Arslan, Rodrigo Wilkens, Wei He, Dilara Torunoğlu-Selamet, Thomas Pickard, Aline Villavicencio, Adriana S. Pagano, and Gülşen Eryiğit. 2026. MWE-2026 Shared Task 2: AdMIRE 2 - Advancing Multimodal Idiomaticity Representation. In *Proceedings of the 22nd Workshop on Multiword Expressions (MWE 2026)*, Rabat, Morocco. Association for Computational Linguistics.
- Joanne Boisson, Jose Camacho-Collados, and Luis Espinosa Anke. 2022. Cardiffnp-metaphor at semeval-2022 task 2: Targeted fine-tuning of transformer-based language models for idiomaticity detection. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 169–177.
- Fabienne Cap, Manju Nirmal, Marion Weller, and Sabine Schulte Im Walde. 2015. How to account for idiomatic german support verb constructions in statistical machine translation. In *Proceedings of the 11th workshop on multiword expressions*, pages 19–28.
- Jianlyu Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. **M3-embedding: Multi-linguality, multi-functionality, multi-granularity text embeddings through self-knowledge distillation**. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 2318–2335, Bangkok, Thailand. Association for Computational Linguistics.
- Tobias Domhan, Eva Hasler, Ke M Tran, Sony Trenous, Bill Byrne, and Felix Hieber. 2022. The devil is in the details: On the pitfalls of vocabulary selection in neural machine translation. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1861–1874.
- Demeke Endalie, Getamesay Haile, and Wondmagegn Taye. 2023. Deep learning-based idiomatic expression recognition for the amharic language. *PLoS One*, 18(12):e0295339.
- Michael Flor, Xinyi Liu, and Anna Feldman. 2025. A survey of idiom datasets for psycholinguistic and computational research. In *Proceedings of the 21st Conference on Natural Language Processing (KONVENS 2025): Long and Short Papers*, pages 90–100.
- Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. 2024. Clip-adapter: Better vision-language models with feature adapters. *International Journal of Computer Vision*, 132(2):581–595.
- Yash Jakhotiya, Vaibhav Kumar, Ashwin Pathak, and Raj Shah. 2022. Jarvix at semeval-2022 task 2: It takes one to know one? idiomaticity detection using zero and one-shot learning. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 165–168.
- Yizhang Jin, Jian Li, Tianjun Gu, Yexin Liu, Bo Zhao, Jinxiang Lai, Zhenye Gan, Yabiao Wang, Chengjie Wang, Xin Tan, and 1 others. 2025. Efficient multimodal large language models: A survey. *Visual Intelligence*, 3(1):27.
- Muhammad Farnal Khan and Mousumi Akter. 2025. Evaluating large language models on urdu idiom translation. *arXiv preprint arXiv:2510.17460*.
- Thanet Markchom, Tong Wu, Liting Huang, and Huizhi Liang. 2025. Uor-ncl at semeval-2025 task 1: Using generative llms and clip models for multilingual multimodal idiomaticity representation. *arXiv preprint arXiv:2502.20984*.
- Dylan Phelps, Xuan-Rui Fan, Edward Gow-Smith, Harish Tayyar Madabushi, Carolina Scarton, and Aline Villavicencio. 2022. Sample efficient approaches for idiomaticity detection. *MWE 2022*, page 105.
- Dylan Phelps, Thomas MR Pickard, Maggie Mi, Edward Gow-Smith, and Aline Villavicencio. 2024. Sign of the times: Evaluating the use of large language models for idiomaticity detection. In *Proceedings of the Joint Workshop on Multiword Expressions and Universal Dependencies (MWE-UD)@ LREC-COLING 2024*, pages 178–187.
- Thomas Pickard, Aline Villavicencio, Maggie Mi, Wei He, Dylan Phelps, and Marco Idiart. 2025a. **SemEval-2025 task 1: AdMIRE - advancing multimodal idiomaticity representation**. In *Proceedings of the 19th International Workshop on Semantic Evaluation (SemEval-2025)*, pages 2597–2609, Vienna, Austria. Association for Computational Linguistics.
- Thomas Pickard, Aline Villavicencio, Maggie Mi, Wei He, Dylan Phelps, and Marco Idiart. 2025b. Semeval-2025 task 1: Admire—advancing multimodal idiomaticity representation. *arXiv preprint arXiv:2503.15358*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, and 1 others. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR.
- Dilara Torunoğlu-Selamet, Dogukan Arslan, Rodrigo Wilkens, Wei He, Doruk Eryiğit, Thomas Pickard, Adriana S. Pagano, Aline Villavicencio, Gülşen Eryiğit, Ágnes Abuczki, Aida Cardoso, Alesia Lazarenka, Dina Almassova, Amalia Mendes, Anna Kanellopoulou, Antoni Brosa-Rodríguez, Baiba Saulite, Beata Wojtowicz, Bolette Pedersen, and 59 others. 2026. **A parallel cross-lingual benchmark for multimodal idiomaticity understanding**. *Preprint*, arXiv:2601.08645.
- Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer,

Ye Xia, Basil Mustafa, and 1 others. 2025. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv preprint arXiv:2502.14786*.

Tongguan Wang, Mingmin Wu, Guixin Su, Dongyu Su, Yuxue Hu, Zhongqiang Huang, and Ying Sha. 2025. Mchirc: A multimodal benchmark for chinese idiom reading comprehension. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 25398–25406.

Jialu Xing, Jianping Liu, Jian Wang, Lulu Sun, Xi Chen, Xunxun Gu, and Yingfei Wang. 2024. A survey of efficient fine-tuning methods for vision-language models—prompt and adapter. *Computers & Graphics*, 119:103885.

Ziheng Zeng and Suma Bhat. 2021. Idiomatic expression identification using semantic compatibility. *Transactions of the Association for Computational Linguistics*, 9:1546–1562.

A Appendix

Table 3 summarises the official Codabench blind test results for our submission nikoniko. For Subtask A (image+text), our system achieved an average Top₁ accuracy of 0.35 and an NDCG of 0.73 across 15 languages. For Subtask B (text only), the corresponding scores were 0.32 Top₁ and 0.71 NDCG. In both subtasks, our run ranked fifth on the Codabench leaderboard at evaluation time. The per language breakdown shows clear variation across the multilingual blind test. For Subtask A, the best performing languages are Portuguese Brazil (PT-BR) (0.46), PT (0.42), and Norwegian (0.42), while Spanish Ecuador (ES-EC) (0.17), Uzbek (UZ) (0.29), and Georgian (KA) (0.27) emerge as the most challenging. Despite these differences in Top₁ accuracy, NDCG scores remain consistently high in the 0.69 to 0.76 range, with an overall average of 0.73.

Language	Subtask A (img+txt)		Subtask B (text)	
	Top ₁ (%)	NDCG@5	Top ₁ (%)	NDCG@5
Chinese (ZH)	35.0	0.72	28.0	0.69
Georgian (KA)	27.0	0.69	28.0	0.68
Greek (EL)	36.0	0.72	28.0	0.71
Igbo (IG)	33.0	0.73	29.0	0.69
Kazakh (KK)	33.0	0.74	28.0	0.72
Norwegian (NO)	42.0	0.75	41.0	0.75
Portuguese BR (PT-BR)	46.0	0.76	37.0	0.75
Portuguese (PT)	42.0	0.75	37.0	0.73
Russian (RU)	40.0	0.73	39.0	0.74
Serbian (SR)	39.0	0.74	31.0	0.71
Slovak (SK)	38.0	0.73	35.0	0.73
Slovenian (SL)	40.0	0.75	37.0	0.74
Spanish EC (ES-EC)	17.0	0.66	19.0	0.67
Turkish (TR)	34.0	0.71	32.0	0.70
Uzbek (UZ)	29.0	0.71	32.0	0.71
Macro average	35.4	0.72	32.1	0.72

Table 3: Per language performance on the multilingual blind test. Subtask A uses images and captions; Subtask B is text only.