

Scalable Performance Analysis for Vision-Language Models

Santiago Castro* Oana Ignat* Rada Mihalcea
University of Michigan – Ann Arbor, USA
{sacastro, oignat, mihalcea}@umich.edu

Abstract

Joint vision-language models have shown great performance over a diverse set of tasks. However, little is known about their limitations, as the high dimensional space learned by these models makes it difficult to identify semantic errors. Recent work has addressed this problem by designing highly controlled probing task benchmarks. Our paper introduces a more scalable solution that relies on already annotated benchmarks. Our method consists of extracting a large set of diverse features from a vision-language benchmark and measuring their correlation with the output of the target model. We confirm previous findings that CLIP behaves like a bag of words model and performs better with nouns and verbs; we also uncover novel insights such as CLIP getting confused by concrete words. Our framework is available at <https://github.com/MichiganNLP/Scalable-VLM-Probing> and can be used with other multimodal models and benchmarks.

1 Introduction

Recent years have witnessed an explosion of vision-language models (Lu et al., 2019; Li et al., 2019; Zhang et al., 2021; Radford et al., 2021; Singh et al., 2022). These models have shown great performance in a variety of tasks, such as image/video classification and text-image/video retrieval (Radford et al., 2021; Luo et al., 2022), even without leveraging task-specific or in-domain training. In addition, these models have shown to be practical when leveraged as underlying models for text-to-image generation such as DALL-E 2 (Ramesh et al., 2022) and image captioning such as Clip-Cap (Mokady et al., 2021).

Little is however known about the limitations of these models. Recent work, such

*Equal contribution.

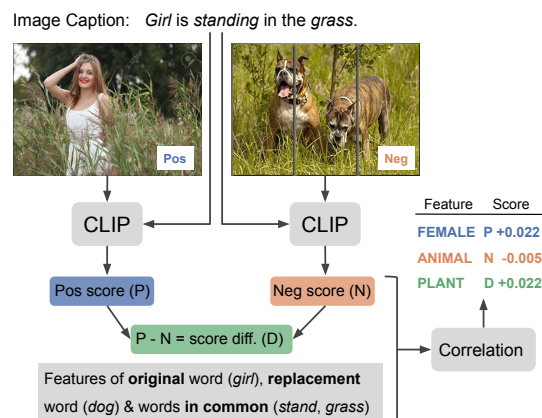


Figure 1: We propose a simple framework to analyze CLIP performance on SVO-Probes data. We test CLIP on the benchmark, extract a diverse set of semantic features from the data, and measure the correlation between each feature and the CLIP score (P , N , or D). Features with positive correlation (e.g., *Female*, *Plant*) impact positively the model performance, while features with negative correlation (e.g., *Animal*) impact negatively the model performance.

as Winoground (Thrush et al., 2022), SVO-Probes (Hendricks and Nematzadeh, 2021), or VALSE (Parcalabescu et al., 2022), have designed benchmark probing tasks by annotating data to follow specific properties (i.e., object color, location, size, swapping word order, replacing words). This line of research led to valuable insights into the limitations of current state-of-the-art multi-modal models such as CLIP (Radford et al., 2021) and ViLBERT (Lu et al., 2019).

An important limitation of current work is the reliance on time-consuming data annotation procedures, making it unscalable and limited in scope. As a complementary solution, we propose a method to probe vision-language models by relying on existing data, without requiring extra annotations. The method consists of extracting a large set of candidate features from a vision-language benchmark and testing their correlation with respect to

the output of the target models on the given benchmark.

By applying our method on CLIP (Radford et al., 2021), a widely used state-of-the-art multi-modal model, using the SVO-Probes (Hendricks and Nematzadeh, 2021) benchmark, we confirm the findings of Thrush et al. (2022) of CLIP behaving like a bag of words model and that of Parcalabescu et al. (2022) of CLIP performing better with nouns and verbs. We also find that CLIP gets confused by concrete words and that it surprisingly improves in performance for more ambiguous words while noting little change from the word frequencies. To the best of our knowledge, we are the first to conduct an in-depth analysis of how language semantic properties influence CLIP’s performance.

We summarize our contributions as follows. First, we propose a scalable way of measuring the limitations of vision-language models. Second, we test our method using a state-of-the-art vision-language model (CLIP) and a popular benchmark (SVO-Probes), validate known challenges, and uncover new ones. Third, our work opens up avenues for future models to focus on solving the newly discovered challenges.

2 Related Work

Recently, an increasing number of benchmarks have been created for the evaluation of vision-language model abilities to perform various multi-modal tasks.

Hendricks and Nematzadeh (2021) evaluate state-of-the-art vision-language models by building SVO-Probes, a probing benchmark focused on verb understanding. They show that image–language transformers fail to distinguish fine-grained differences between images and find that they are worse at verb understanding compared to subjects or objects. In our work, we continue their proposed future work direction by analyzing model performance on fine-grained verb categories.

Other work focuses on testing more precise capabilities of vision-language models using other probing techniques. In VALSE, Parcalabescu et al. (2022) demonstrate that vision-language models have difficulty in counting objects and in correctly classifying spatial relations between objects. Salin et al. (2022); Zhao et al. (2022) show that, although state-of-the-art vision-language models can grasp color, they do not fully understand more difficult concepts such as object size and position in the

image.

In Winoground, Thrush et al. (2022) designed adversarial examples that require differentiating between a similar image and text, where the text pairs only differ in their word order. Their results show that state-of-the-art vision-language models lack compositional reasoning abilities. Several other works build benchmarks on probing vision-language on compositional reasoning (Akula et al., 2020; Ma et al., 2023; Liu et al., 2023; Park et al., 2022; Yuksekogonul et al., 2023) find that they behave like a bag-of-words model – i.e., have poor relational understanding and a severe lack of word order sensitivity.

In contrast, our work focuses not on creating new probing tasks for vision-language models, but on using current benchmarks to learn additional, more fine-grained features that can be discovered using simple correlation methods. To the best of our knowledge, we are the first to analyze the performance of CLIP on a diverse set of semantic features and use correlation methods to draw insights about what concepts are challenging for the model.

3 Methodology to Probe CLIP

Given a benchmark, we measure how a vision-language model performs on a variety of semantic concepts. Our aim is to quantify which concepts are the most and the least challenging for the model. Our setting is illustrated in Figure 1, and can be described in three main steps.

First, we use CLIP (Radford et al., 2021) to compute scores for instances from the SVO-Probes (Hendricks and Nematzadeh, 2021) dataset and obtain two corresponding alignment scores for each sentence and its corresponding *positive* and *negative* image. Next, we extract and process a diverse set of semantic features from SVO-Probes. Finally, we compute the correlation coefficients between each feature and the CLIP score. The features with the highest coefficients will represent concepts that CLIP performs well on, while features with the lowest coefficients will represent challenging concepts for CLIP.

3.1 Dataset

We choose the SVO-Probes (Hendricks and Nematzadeh, 2021) dataset due to its design and large scale size (421 verbs and over 48,000 image-sentence pairs). SVO-Probes was designed for

probing image-text models for their understanding of subject, verb, object triplets. Each instance from the dataset consists of a text caption, a *positive* image that matches the caption, and a controlled (adversarial) *negative* image that shares two out of three aspects (subject, verb, and object) from the sentence but does not match the other one, as shown in Figure 1. These controlled examples enable one to probe models for their understanding of verbs as well as subjects and objects. The instances also include information about the negative image, such as a (hidden) associated negative caption which we leverage in this paper.

We propose to use this dataset to evaluate the CLIP (Radford et al., 2021) model. We choose to test CLIP, as opposed to other language-vision models, due to its widely-spread use and impressive zero-shot performance on a variety of vision-language tasks (e.g., text-to-image retrieval, image question answering, human action segmentation, image-sentence alignment – Cafagna et al. 2021). Furthermore, Hendricks and Nematzadeh (2021) test only ViLBERT-based (Lu et al., 2019) models, which are known to perform worse than CLIP (Cafagna et al., 2021).

3.2 Model Output

As depicted in Figure 1, we obtain three CLIP scores for each pair of *positive* and *negative* images: a *positive* score (P), computed between the caption and the *positive* image; a *negative* score (N), computed between the caption and the *negative* image; and the *difference* between these scores ($D = P - N$).

Because the text and the positive image are aligned, P represents an absolute alignment score. In the case of the text and the negative image, even though the negative image is similar in some ways to the text (because of how SVO-Probes was designed), they do not correspond to each other. Thus, N represents an absolute misalignment score. D represents a relative alignment score. Ideally, CLIP should have a high P score and a low N score, and a high difference between them (a high D). We propose to pay special attention to D given that CLIP is generally used in relative comparisons, such as when using it for classification (choosing the class text that maximizes the alignment score, given an image) or when using it for retrieval (finding the text/image that maximizes the alignment score given an image/text).

3.3 Feature Extraction

For each given sentence and corresponding image in the benchmark, we extract features from the words marked in the SVO-Probes benchmark (i.e., subject, verb, and object).

If the corresponding image is *positive*, all the extracted features are from words *in common*, i.e., that appear both in the image and the text. Otherwise, if the corresponding image is *negative*, in addition to words *in common*, we also extract features from words present in the sentence and not in the image (*original* word) and words present in the image but not in the text (*replacement* word). As an example, in Figure 1 the words *in common* are “sit” and “grass”, the *original* word is “girl” and the *replacement* word is “dogs”. The *original* and *replacement* words represent what is different between the image and the text, while the words *in common*, as the name suggests, represent what is common between the image and the text.

We extract the following **semantic** textual features: Levin (1993) verb classes, LIWC psycholinguistic markers (Pennebaker et al., 2007, 2015), General Inquirer (Stone et al., 1967) semantic classes, WordNet hypernyms (Miller, 1995), word presence, semantic similarity, ambiguity, frequency, sentence length, and concreteness (Brybaert et al., 2014).

Levin verb classes. Levin (1993) groups verbs according to their semantic content and also according to their participation in argument alternations.

Levin’s semantic content-based taxonomy provides a classification of 3,024 verbs into 48 broad classes and 192 fine-grained classes.¹ A verb can belong to one or more classes. Some examples of verb classes are: (1) broad *change of state* (e.g., clean, divide, soak), *manner of motion* (e.g., climb, drop, run) or *social interaction* (e.g., marry, meet, hug); (2) fine-grained: “roll” verbs (e.g., bounce, coil, drift), “run” verbs (e.g., amble, bolt, race) or “hug” verbs (e.g., cover, encircle, touch)

LIWC psycholinguistic markers. Linguistic Inquiry and Word Count (LIWC) (Pennebaker et al., 2007, 2015) is a widely used word-counting software that includes dictionaries of English words related to human cognitive processes. Specifically, we use the LIWC2015 dictionary, which contains 6,400 words and word stems. Each word or word

¹<https://websites.umich.edu/~jlawler/levin.verbs>

stem defines one or more categories: e.g., the word “mother” is assigned the categories: *female*, *family*, *social*.

General Inquirer classes. General Inquirer (Stone et al., 1967) is a resource for automatic content analysis. More specifically, it categorizes words into emotional and cognitive states, as well as into diverse semantic categories outlined in the Lasswell dictionary (Namenwirth and Weber, 1987, pg. 46–53).

WordNet classes. WordNet (Miller, 1995) is a large lexical database of English words that are grouped into sets of cognitive synonyms, known as synsets. The synsets are interlinked by semantic and lexical relations. The most frequent relation among synsets is the super-subordinate relation, also called *hyperonymy*. It links more general synsets to specific ones: e.g., “building” is a *hypernym* of “house” and “school”. For each given word, we collect all the hypernyms of the most common word synset.

Word presence. For each given word, we use a marker to indicate if the word is present or not in the sentence. Note that studying the effect of specific words does not imply that they have no dependencies with other words. Their role may change depending on the context; however, we study them in aggregate.

Sentence length. We measure the length of each sentence as the number of words in the sentence.

Semantic similarity. In the case of *negative* images, we compute the cosine similarity score between the *original* words and the corresponding *replacement* words. The word representations are computed using Sentence-Transformers (Reimers and Gurevych, 2019), with the model `all-MiniLM-L6-v2`, which is based on MiniLM (Wang et al., 2020).

Concreteness score. For measuring the concreteness of words, we use a dataset of words with associated concreteness scores from Brysbaert et al. (2014). Each word is labeled by a human annotator with a value between 1 (very abstract) and 5 (very concrete). Abstract words (e.g., “beauty”, “sadness”) denote ideas, feelings, or other intangible concepts while concrete words (e.g., “table”, “write”) refer to objects and actions.

Ambiguity. We measure the ambiguity of a given word by counting the number of synsets in WordNet (Miller, 1995).

Frequency. We measure the word frequency in a subset (~13M image captions) of LAION (Schuhmann et al., 2021), a dataset representative of CLIP’s training data.

3.4 Feature Representation

The **binary** features, i.e., Levin, LIWC, General Inquirer, WordNet classes, and word presence, are represented as binary vectors, while the **numerical** features i.e., sentence length, concreteness, similarity, ambiguity, and frequency are standardized. All the features are then concatenated together.

3.5 Feature Selection

We measure the degree of correlation between each feature and the model performance. For each of the **binary** features, we compute a two-sample two-tailed t-test (Student, 1908) along with the model output score. This test evaluates if the means of the populations coming from each feature value (true or false) are significantly different. If so, we compute the difference of means as a reference value. In the case of **numerical** features, we compute the Pearson’s correlation coefficient (Benesty et al., 2009) between each feature and the model performance score.

Next, we employ a one-sample, two-tailed t-test to determine if the coefficient is significantly different from zero, i.e., if there is any correlation according to this metric. We chose a p-value threshold of 0.05 (a confidence level of 95%) to filter out the features.²

3.6 Experimental Details

We use an OpenAI pre-trained CLIP (Radford et al., 2021) ViT-L/14 (Dosovitskiy et al., 2021) model.

4 Results

Our main observations and takeaways from this evaluation are the following:

(1) CLIP behaves like a bag-of-words model.

As shown in Figure 2, the distributions of P and N highly overlap. This is explained partly by the negative image being adversarial; it contains elements

²See the obtained scores and p-values in the web page from this paper: <https://github.com/MichiganNLP/Scalable-VLM-Probing>.




Figure 2: Histogram plot of the distribution of CLIP scores between the text with the positive image, and the text with the negative image. A kernel density estimation curve is included to aid this visualization.

in common with the text. This finding is coherent with that of Thrush et al. (2022), that CLIP performs like a bag-of-words model.

This finding is also supported by the fact that many features from words *in common* contribute to increasing both the positive (P) and the negative scores (N): e.g., `hypernym_food.n.02` increases P by 0.042 and N by 0.050; LIWC “money” increases P by 0.036, and N by 0.032. As described in Section 3.5, we measure the importance of each feature as the difference of means between the CLIP scores when the feature is present and when is not. We observed that many of the features for the words *in common* appeared to influence similarly both P and N , confirming this hypothesis.

(2) CLIP performs better with nouns than with verbs. When computing the number of times CLIP assigns a higher score to the similarity between the text and the *positive* image as compared to the similarity between the text and the *negative* image, the verbs obtain 81.45% accuracy while the subjects get 86.87% and the objects 88.78%. The number obtained for verbs is relatively close to that of a similar setting experimented by the VALSE benchmark (Parcalabescu et al., 2022), in which they reported 75.6% accuracy (also considering that we could not determine which pre-trained CLIP variant the authors evaluated). At the same time, the noun (objects and subjects) replacement numbers are consistent with those reported by the same authors (88.8%), obtained from FOIL it! (Shekhar et al., 2017).




Figure 3: Linear regression plot of the average concreteness for the words in the sentence that are common to both images vs. the CLIP score. The shadowed areas are 95%-confidence intervals for the expected value.

(3) CLIP gets confused by concrete words. Figure 3 shows both the *positive* and *negative* CLIP scores improve the more concrete a word is (words from the caption represented in both the positive and the negative images). As seen in this figure, however, the *negative* score increases faster. This implies that, in an image classification or image-to-text retrieval setting, CLIP will more likely consider an incorrect text as correct if it has more concrete words than the actual correct text.

(4) CLIP prefers average-length sentences. We present in Figure 4 how the score is affected by the caption sentence word length. CLIP presents a low performance when the sentences are very short (around 3 words long), improving when the sentences are longer since the difference between the *positive* and *negative* scores (D) gets larger with the sentence length.

Figure 5 shows how the CLIP scores are distributed for the different number of words, showing for example that there is a great overlap between the similarity scores between texts of length 6 and a *negative* image, and the similarity scores between texts of length 3 and a *positive* image. This implies CLIP is more likely to select the wrong text when comparing an image with a short correct text and one with long incorrect text.

(5) CLIP is affected by word frequency. Figure 6 studies the frequency effect on the score for the words that represent concepts that appear in both the *positive* and *negative* images. The more frequent a word is, the higher the CLIP score. Still, the difference in scores is barely affected.




Figure 4: Line plot of the number of words in the caption sentence vs. the CLIP score. The shadowed areas are 95%-confidence intervals for the expected value.




Figure 5: Box plot for the number of words in the caption sentence vs. the CLIP score. Unlike Figure 4 that shows the expected values, this plot shows the distributions.




Figure 6: Linear regression plot of the average frequency for the words in the sentence that are common to both images vs. the CLIP score. The shadowed areas are 95%-confidence intervals for the expected value.

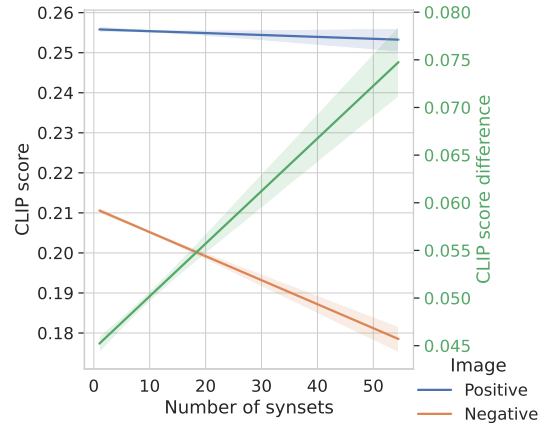


Figure 7: Linear regression plot of the average synset count for the words in the sentence that are common to both images vs. the CLIP score. The shadowed areas are 95%-confidence intervals for the expected value.

(6) The score improves for more ambiguous words. Surprisingly, there is a larger gap in the score difference (D) when the words have more meanings associated with them (for the words that represent concepts in both the *positive* and *negative images*), as shown in Figure 7. The positive score seems to remain almost constant while the negative score drops, widening the difference. The word frequency seems not to be a confounding factor based on (5).

(7) Similar situations confuse CLIP. Unsurprisingly, the higher the similarity between the caption and the negative image caption, the higher the *negative* CLIP score, as depicted by Figure 8.

We also studied the influence of the similarity between the *original* word (from the caption) and the *replacement* word (from the text associated with the negative image) in Figure 9. The effect of the word change seems to be smaller than that of the whole sentence change.

(8) CLIP performs relatively better on *nature-related* and *personal care* concepts and relatively worse on *furniture*, *transportation*, *herbivores*, *sports*, *academia*. As mentioned in Section 3.2, score D measures the relative CLIP performance, which is more relevant for retrieval models like CLIP. Therefore, we measure the importance of each feature with respect to D . Specifically, we compute the mean differences of the D scores when the binary feature is present and when is not. We show the CLIP performance analysis on **binary** features in Table 1. Following the example of

Topic	Feature	Mean diff.	Example Words
CLIP PERFORMS BETTER ON			
Natural Phenomenon	Hypernym physical_phenomenon.n.01 (original)	0.038	snow, fog, rain, mist
	Hypernym physical_phenomenon.n.01 (replacement)	0.022	snow, rain, cloud, fog, mist
Waterfront Infrastructure	Hypernym platform.n.01 (original)	0.038	pier, deck, podium
	Hypernym horizontal_surface.n.01 (original)	0.032	pier, pavement, quay
Landscapes	Hypernym community.n.06 (original)	0.038	meadow, desert, grassland
	Hypernym natural_elevation.n.01 (original)	0.035	dune, sandbar, reef
	Hypernym geological_formation.n.01 (original)	0.027	beach, shore, cliff
	Hypernym plant.n.02 (original)	0.025	grass, tree, flower
	Hypernym natural_elevation.n.01 (replacement)	0.020	mountain, hill
Grooming	Presence of word “wash” (original)	0.035	wash
	Levin “floss verbs” (original)	0.030	wash, brush, shave
	Levin “wipe verbs”(original)	0.022	wear, sweep, trim, rub
	Levin “dress verbs” (original)	0.027	exercise, bathe, dress
Domestic Animals	Hypernym young.n.01 (original)	0.033	puppy, kitten, foal
	Hypernym domestic_animal.n.01 (original)	0.032	puppy, retriever, pug
	General Inquirer “animal” (replacement)	0.023	dog, animal, cat, goat
	Hypernym canine.n.02 (replacement)	0.021	puppy, retriever, pug
CLIP PERFORMS WORSE ON			
Furniture	Presence of word “sofa” (in common)	-0.032	sofa
	Hypernym bedroom_furniture.n.01 (in common)	-0.026	bed, sofa
	Hypernym furniture.n.01 (in common)	-0.017	couch, bed, sofa, chair, bench
	LIWC “home” (in common)	-0.015	bed, window, sofa, room
Transportation	Presence of word “ride” (original)	-0.027	ride
	Hypernym vessel.n.02 (in common)	-0.019	boat, ship, yacht
	Levin “pedal” verbs (original)	-0.018	ride, drive, fly, sail, cruise
	Hypernym craft.n.02 (in common)	-0.018	boat, balloon, ship, scooter, kayak
Herbivores	Hypernym ungulate.n.01 (in common)	-0.021	horse, cow, camel, goat, deer
	Presence of word “horse” (in common)	-0.019	horse
Sports	Hypernym happening.n.01 (in common)	-0.021	wave, win, tap, slam
	Hypernym contestant.n.01 (in common)	-0.020	footballer, golfer, goalkeeper, cricketer, tackle
Academia	Levin “admire” verbs (original)	-0.017	stand, enjoy, admire, support
	General Inquirer “academia” (in common)	-0.020	student, classroom, library, teacher, book, computer, conference
	Presence of word “student” (in common)	-0.020	student

Table 1: CLIP relative performance analysis on a subset of binary features: the top-5 **easier** topics are *Natural Phenomenon*, *Waterfront Infrastructure*, *Landscapes*, *Grooming* and *Domestic Animals*, while the top-5 **harder** topics are *Furniture*, *Transportation*, *Herbivores*, *Sports* and *Academia*.




Figure 8: Linear regression plot of the similarity between the text caption and the negative image text caption vs. the CLIP score for the negative image. The shadowed areas are 95%-confidence intervals for the expected value. The unimodal distributions are also shown.




Figure 9: Linear regression plot of the similarity between the originally replaced word from the text caption and new word from the negative image text caption vs. the CLIP score for the negative image. The shadowed areas are 95%-confidence intervals for the expected value. The unimodal distributions are also shown.

SEAL (Rajani et al., 2022), we use ChatGPT to cluster the features under a broad topic automatically.³

We find that CLIP performs relatively **better** on topics related to nature: *Natural Phenomenon*, *Waterfront Infrastructure*, *Landscapes*, *Domestic Animals*, and personal care: *Grooming*, and **worse** on topics like *Furniture*, *Transportation*, *Herbivores*, *Sports* and *Academia*.

5 Conclusion

In this work, we proposed a simple and effective method to probe vision-language models. Our method is scalable, as it does not require data annotation and makes use of existing datasets. With our method, we analyzed the performance of CLIP, a popular state-of-the-art multi-modal model, on the SVO-Probes benchmark. We confirmed the recent findings of Thrush et al. (2022) of CLIP behaving like a bag of words model and that of Parcalabescu et al. (2022) of CLIP performing better with nouns and verbs. We also uncovered novel findings, for instance, that CLIP gets confused by concrete words, surprisingly improves performance for more ambiguous terms, or that the frequency of words does not significantly change the behavior of CLIP.

We hope our work contributes to ongoing efforts to discover the limitations of multi-modal models and help build more robust and reliable systems. Our framework can be easily used to analyze other benchmarks, features, and multi-modal models, and it is publicly available at <https://github.com/MichiganNLP/Scalable-VLM-Probing>.

Limitations

SVO-Probes dataset is not balanced. For example, “person”, “man”, and “woman” are considerably more frequent than other words. Future work can address this limitation by aggregating data from multiple datasets and balancing it out. At the same time, the target dataset should reflect the phenomenon one wants to study. For example, LAION (Schuhmann et al., 2021) could be employed to study how VLMs perform with everyday human actions. Still, it may be too centered around objects (as opposed to actions) and overly noisy – future work can consider using subsets instead. A

³We use the following prompt: "Name a topic for the following words: ..."

smaller yet cleaner alternative is Conceptual Captions (Sharma et al., 2018).

Another limitation is not considering the polysemy when using LIWC or Levin dictionaries. This may lead to incorrect word categorization and influence the error analysis. Future work can mediate this limitation by linking semantic dictionaries such as Levin or LIWC with their WordNet synsets.

Acknowledgements

We want to thank the anonymous reviewers for their helpful comments. We also thank Artem Abzaliev, Fabian Caba, Mohamed El Banani, and Karan De-sai for the productive discussions. This material is partly based on work supported by the Automotive Research Center (“ARC”). Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of ARC or any other related entity.

References

- Arjun Akula, Spandana Gella, Yaser Al-Onaizan, Song-Chun Zhu, and Siva Reddy. 2020. [Words aren’t enough, their order matters: On the robustness of grounding visual referring expressions](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6555–6565, Online. Association for Computational Linguistics.
- Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. 2009. [Pearson correlation coefficient](#). In *Noise reduction in speech processing*, pages 37–40. Springer.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit*. O’Reilly Media, Inc.
- Marc Brysbaert, Amy Beth Warriner, and Victor Kuperman. 2014. [Concreteness ratings for 40 thousand generally known english word lemmas](#). *Behavior research methods*, 46(3):904–911.
- Michele Cafagna, Kees van Deemter, and Albert Gatt. 2021. [What vision-language models ‘see’ when they see scenes](#). *ArXiv*, abs/2109.07301.
- Casper da Costa-Luis, Stephen Karl Larroque, Kyle Al-tendorf, Hadrien Mary, richardsheridan, Mikhail Korobov, Noam Raphael, Ivan Ivanov, Marcel Bargull, Nishant Rodrigues, Guangshuo Chen, Antony Lee, Charles Newey, CrazyPython, JC, Martin Zugnoni, Matthew D. Pagel, mjestevens777, Mikhail Dektyarev, Alex Rothberg, Alexander Plavin, Daniel Panteleit, Fabian Dill, FichteFoll, Gregor Sturm, HeoHeo, Hugo van Kemenade, Jack McCracken, MapleCCC, and Max Nordlund. 2023. [tqdm: A fast, Extensible Progress Bar for Python and CLI](#).
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). In *International Conference on Learning Representations*.
- Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. 2020. [Array programming with NumPy](#). *Nature*, 585(7825):357–362.
- Lisa Anne Hendricks and Aida Nematzadeh. 2021. [Probing image-language transformers for verb understanding](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3635–3644, Online. Association for Computational Linguistics.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- John D Hunter. 2007. [Matplotlib: A 2D graphics environment](#). *Computing in science & engineering*, 9(03):90–95.
- Beth Levin. 1993. *English verb classes and alternations: A preliminary investigation*. University of Chicago press.
- Quentin Lhoest, Albert Villanova del Moral, Patrick von Platen, Thomas Wolf, Mario Šaško, Yacine Jernite, Abhishek Thakur, Lewis Tunstall, Suraj Patil, Mariama Drame, Julien Chaumond, Julien Plu, Joe Davison, Simon Brandeis, Victor Sanh, Teven Le Scao, Kevin Canwen Xu, Nicolas Patry, Steven Liu, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Nathan Raw, Sylvain Lesage, Anton Lozhkov, Matthew Carrigan, Théo Matussière, Leandro von Werra, Lysandre Debut, Stas Bekman, and Clément Delangue. 2021. [Datasets: A Community Library for Natural Language Processing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184. Association for Computational Linguistics.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. [VisualBERT: A simple and performant baseline for vision and language](#). *ArXiv*, abs/1908.03557.
- Fangyu Liu, Guy Edward Toh Emerson, and Nigel Collier. 2023. [Visual spatial reasoning](#). *Transactions of the Association for Computational Linguistics*.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. [ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks](#).

- In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Huashao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. 2022. [CLIP4Clip: An empirical study of CLIP for end to end video clip retrieval and captioning](#). *Neurocomputing*, 508:293–304.
- Zixian Ma, Jerry Hong, Mustafa Omer Gul, Mona Gandhi, Irena Gao, and Ranjay Krishna. 2023. [CREPE: Can vision-language foundation models reason compositionally?](#) In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10910–10921.
- George A. Miller. 1995. [WordNet: A lexical database for English](#). *Commun. ACM*, 38(11):39–41.
- Ron Mokady, Amir Hertz, and Amit H Bermano. 2021. [ClipCap: CLIP prefix for image captioning](#). *arXiv preprint arXiv:2111.09734*.
- J. Zvi Namenwirth and Robert Philip Weber. 1987. *Dynamics of culture*. Allen & Unwin – Boston, Mass., USA.
- Letitia Parcalabescu, Michele Cafagna, Lilitta Muradjan, Anette Frank, Iacer Calixto, and Albert Gatt. 2022. [VALSE: A task-independent benchmark for vision and language models centered on linguistic phenomena](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8253–8280, Dublin, Ireland. Association for Computational Linguistics.
- Jae Sung Park, Sheng Shen, Ali Farhadi, Trevor Darrell, Yejin Choi, and Anna Rohrbach. 2022. [Exposing the limits of video-text models through contrast sets](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3574–3586, Seattle, United States. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [PyTorch: An Imperative Style, High-Performance Deep Learning Library](#). In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. [Scikit-learn: Machine learning in Python](#). *Journal of Machine Learning Research*, 12:2825–2830.
- James W. Pennebaker, Roger John Booth, and Martha E. Francis. 2007. [Linguistic inquiry and word count \(LIWC2007\)](#).
- James W. Pennebaker, Ryan L. Boyd, Kayla Jordan, and Kate G. Blackburn. 2015. [The development and psychometric properties of LIWC2015](#).
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Nazneen Rajani, Weixin Liang, Lingjiao Chen, Margaret Mitchell, and James Zou. 2022. [SEAL: Interactive tool for systematic error analysis and labeling](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 359–370, Abu Dhabi, UAE. Association for Computational Linguistics.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. [Hierarchical text-conditional image generation with CLIP latents](#). *arXiv preprint arXiv:2204.06125*.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Emmanuelle Salin, Badreddine Farah, S. Ayache, and Benoit Favre. 2022. [Are vision-language transformers learning multimodal representations? a probing perspective](#). In *AAAI Conference on Artificial Intelligence*, pages 11248–11257.
- Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. 2021. [LAION-400M: Open dataset of CLIP-filtered 400 million image-text pairs](#). In *Proceedings of the NeurIPS Data Centric AI Workshop*.
- Skipper Seabold and Josef Perktold. 2010. [statsmodels: Econometric and statistical modeling with python](#). In *9th Python in Science Conference*.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. [Conceptual Captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, Melbourne, Australia. Association for Computational Linguistics.

- Ravi Shekhar, Sandro Pezzelle, Yauhen Klimovich, Aurélie Herbelot, Moin Nabi, Enver Sangineto, and Raffaella Bernardi. 2017. [FOIL it! find one mismatch between image and language caption](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 255–265, Vancouver, Canada. Association for Computational Linguistics.
- Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. 2022. [FLAVA: A foundational language and vision alignment model](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15638–15650.
- Philip J. Stone, Dexter C. Dunphy, and Marshall S. Smith. 1967. [The general inquirer: A computer approach to content analysis](#). *American Educational Research Journal*, 4:397.
- Student. 1908. [The probable error of a mean](#). *Biometrika*, pages 1–25.
- The pandas development team. 2023. [pandas-dev/pandas: Pandas](#).
- Tristan Thrush, Ryan Jiang, Max Bartolo, Amanpreet Singh, Adina Williams, Douwe Kiela, and Candace Ross. 2022. [Winoground: Probing vision and language models for visio-linguistic compositionality](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5238–5248.
- Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. 2020. [SciPy 1.0: fundamental algorithms for scientific computing in python](#). *Nature methods*, 17(3):261–272.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. [MiniLM: Deep self-attention distillation for task-agnostic compression of pre-trained transformers](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 5776–5788. Curran Associates, Inc.
- Michael L. Waskom. 2021. [seaborn: statistical data visualization](#). *Journal of Open Source Software*, 6(60):3021.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Mert Yuksekgonul, Federico Bianchi, Pratyusha Kalluri, Dan Jurafsky, and James Zou. 2023. [When and why vision-language models behave like bags-of-words, and what to do about it?](#) In *The Eleventh International Conference on Learning Representations*.
- Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021. [VinVL: Revisiting visual representations in vision-language models](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5579–5588.
- Tiancheng Zhao, Tianqi Zhang, Mingwei Zhu, Haozhan Shen, Kyusong Lee, XiaoPeng Lu, and Jianwei Yin. 2022. [VL-CheckList: Evaluating pre-trained vision-language models with objects, attributes and relations](#). *ArXiv*, abs/2207.00221.