

# What does a platypus look like?

## Generating customized prompts for zero-shot image classification

Sarah Pratt<sup>1\*</sup> Ian Covert<sup>1</sup> Rosanne Liu<sup>2,3</sup> Ali Farhadi<sup>1</sup>

<sup>1</sup>University of Washington <sup>2</sup>Google Research, Brain Team <sup>3</sup>ML Collective

### Abstract

Open-vocabulary models are a promising new paradigm for image classification. Unlike traditional classification models, open-vocabulary models classify among any arbitrary set of categories specified with natural language during inference. This natural language, called “prompts”, typically consists of a set of hand-written templates (e.g., “a photo of a { }”) which are completed with each of the category names. This work introduces a simple method to generate higher accuracy prompts, without relying on any explicit knowledge of the task domain and with far fewer hand-constructed sentences. To achieve this, we combine open-vocabulary models with large language models (LLMs) to create Customized Prompts via Language models (CuPL, pronounced “couple”). In particular, we leverage the knowledge contained in LLMs in order to generate many descriptive sentences that contain important discriminating characteristics of the image categories. This allows the model to place a greater importance on these regions in the image when making predictions. We find that this straightforward and general approach improves accuracy on a range of zero-shot image classification benchmarks, including over one percentage point gain on ImageNet. Finally, this simple baseline requires no additional training and remains completely zero-shot.

## 1. Introduction

Open-vocabulary models [40, 23, 42, 63] achieve high classification accuracy across a large number of datasets without labeled training data for those tasks. To accomplish this, these models leverage the massive amounts of image-text pairs available on the internet by learning to associate the images with their correct caption, leading to greater flexibility during inference. Unlike standard models, these

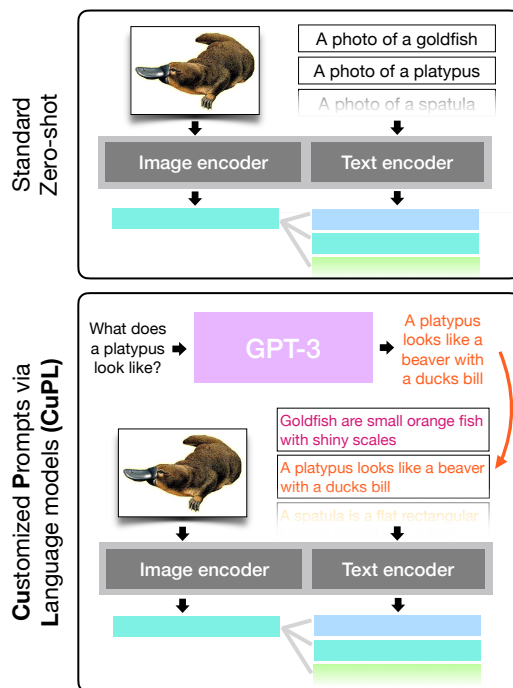


Figure 1. **Schematic of the method.** (Top) The standard method of a zero-shot open-vocabulary image classification model (e.g., CLIP [42]). (Bottom) Our method of CuPL. First, an LLM generates descriptive captions for given class categories. Next, an open-vocabulary model uses these captions as prompts for performing classification.

models classify images by providing a similarity score between an image and a caption. To perform inference, one can generate a caption or “prompt” associated with each of the desired categories, and match each image to the best prompt. This means that categories can be selected ad hoc and adjusted without additional training.

However, this new paradigm poses a challenge:

*How can we best represent an image category through natural language prompts?*

\*Correspondence to spratt3@uw.edu.

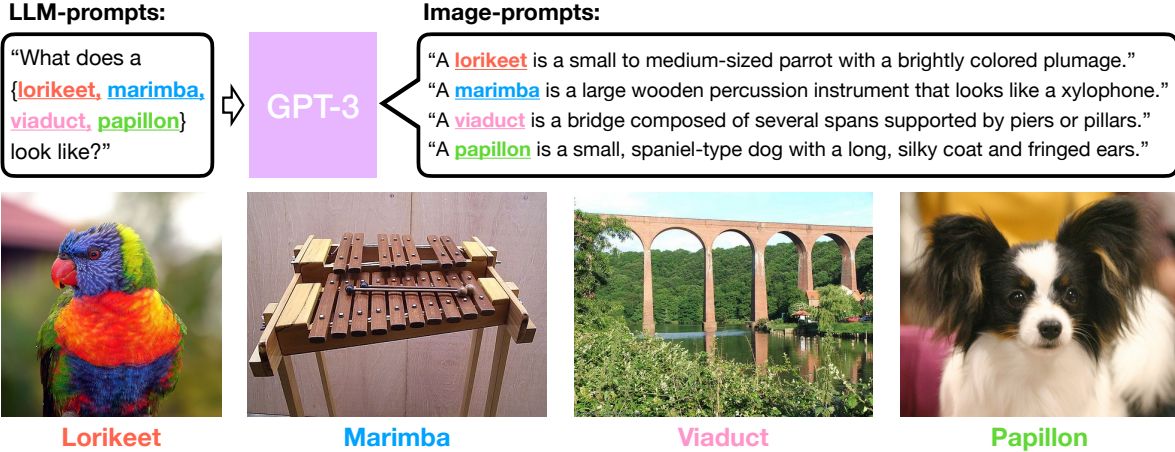


Figure 2. **Example CuPL LLM-prompts and Image-prompts.** LLM-prompts are filled in with a class name and then used as input to GPT-3, which then outputs image-prompts. Example LLM generated image-prompts and associated images from ImageNet are shown. Only image-prompts are used for the downstream image classification.

The standard approach is to hand write a number of prompt templates [42] (e.g., “a photo of a {}”), compile a natural language label for each category in the dataset, and create a set of prompts for each category by filling in each template with the natural language labels. Then, image embeddings are matched to the nearest set of prompt embeddings and labelled with the category associated with that set of prompts (more details in Section 2).

This method has three major drawbacks. Firstly, each prompt template has to be hand-written, so having twice as many prompts for a category requires twice as much human effort. This can become costly as each new dataset typically has a different set of prompt templates [42].

Secondly, the prompt templates must be general enough to apply to all image categories. For example, a prompt for the ImageNet [13] category “platypus” could only be as specific as “a photo of a {platypus}”, and could not be something like “a photo of a {platypus}, a type of aquatic mammal” as that template would no longer be relevant for other image categories. This is limiting, as descriptive details are useful for fine-grained classification. For example, different species of frogs share many of the same characteristics. However, tree frogs can be distinguished with their distinct large eyes. This is a valuable detail for classification but cannot be included in a general template. Therefore, when using these basic templates, the model may not take advantage of this detail in the image, leading to an incorrect categorization as demonstrated in Figure 5.

Lastly, writing high performing prompt templates currently requires prior information about the contents of the dataset. For example, the list of hand-written ImageNet prompts [42] includes “a black and white photo of the {}.”, “a low resolution photo of a {}.”, and “a toy {}.” all of which demonstrate prior knowledge about the type of rep-

resentations present in the dataset. This information is not generalizable to other datasets, as ImageNet contains “black and white” and “toy” representations of its categories, but other datasets do not (e.g., FVGC Aircraft [32]).

To overcome these challenges, we propose Customized Prompts via Language models (CuPL). In this algorithm, we couple a large language model (LLM) with a zero-shot open-vocabulary image classification model. We use the LLM to generate prompts for each of the image categories in a dataset. Using an LLM allows us to generate an arbitrary number of prompts with a fixed number of hand-written sentences. Additionally, these prompts are now customized to each category and contain specified visual descriptions while still remaining zero-shot. This allows prompts to contain details about a class which distinguish it from other similar classes. For example, to describe a tree frog, the LLM generates the sentence “A tree frog looks like a small frog with large eyes.” This not only describes the category, but specifically mentions the eyes, the feature which distinguishes the Tree frog class from the most visually similar classes - other types of frogs. We find that CuPL prompts are rich with these discriminating details and show that the model is able to leverage these details to place more importance on relevant parts of the image when classifying between similar, commonly confused categories (Figure 5).

We find these customized prompts outperform the hand-written templates on 15 zero-shot image classification benchmarks, including a greater than 1 percentage point gain on ImageNet [13] Top-1 accuracy and a greater than 6 percentage point gain on Describable Textures Dataset [11], with fewer hand-written prompts when compared to the standard method used in [42]. Finally, this method requires no additional training or labeled data for either model.

## 2. Methods

The CuPL algorithm consists of two steps: (1) generating customized prompts for each of the categories in a given dataset and (2) using these prompts to perform zero-shot image classification.

### 2.1. Generating Customized Prompts

This step consists of generating prompts using an LLM. For clarity, we distinguish between two different kind of prompts. The first are the prompts which cue the LLM to generate the descriptions of the dataset categories. These prompts do not describe an object, but rather prompt the description of an object (e.g., “What does a platypus look like?”). We will refer to these as “LLM-prompts”.

Secondly, there are the prompts to be matched with images in the zero-shot image classification model. These are the prompts that describe a category (e.g., “A platypus looks like ...”). We call them “image-prompts.” In CuPL, these are the output of the LLM, as exemplified in Figure 2.

In this work, we use GPT-3 [5] as our LLM. To generate our image-prompts, we must first construct a number of LLM-prompt templates. While this does require some engineering by hand, it is significantly less than the amount of hand-engineered sentences used in the standard method of creating image-prompt templates for CLIP. For example, in our ImageNet experiments, we construct 5 LLM-prompt templates compared to the 80 image-prompts used by CLIP for zero-shot ImageNet classification.

After constructing these LLM-prompts, we generate 10 different image-prompts for each of the LLM-prompts. This means for ImageNet we use an LLM to generate a total of 50 customized image-prompts for each image category. For each of these, we generate a maximum of 50 tokens, but halt a generation early if it produces a period. Additionally, we generate with a high temperature of 0.99, which encourages more diversity among the 10 generated image-prompts. We also clean each generated sentence by deleting any blank lines and adding a period at the end.

### 2.2. Utilizing Customized Prompts

After generating image-prompts for each of the categories, we then perform zero-shot image classification. While there are a number of open-vocabulary models [40, 23, 42, 63], we report our results using CLIP [42] as this is the most popular publicly available open-vocabulary model.

CLIP consists of a text encoder and an image encoder (schematic on the top of Figure 1). In the standard setting, there are a number of hand-written templates which can be completed with the relevant category names (e.g. “A photo of a {}”, “A photo of many {}”). To classify the images in a dataset, each of these templates is filled in with a given category name. Then each of these sentences is embedded

via the text encoder, and all sentences completed with the same category name are averaged and normalized. This results in  $n$  embeddings where  $n$  is the number of categories in the dataset. Each of these  $n$  embeddings is the mean of many different sentence embeddings. Then each image in the dataset is embedded using the image encoder. This embedding is compared to each of the  $n$  text embeddings using cosine similarity and is labeled with the most similar one.

CuPL requires only a small adjustment from this standard practice. Instead of filling in the hand-written templates for each category, we simply replace these altogether with the sentences output by GPT-3. This means that for CuPL, hand-written templates are only used as input for the LLM, while the prompts for CLIP are entirely generated text. We present 2 different settings of CuPL (as shown in Table 1), each representing a different trade-off between accuracy and hand-engineering.

**1. CuPL (base).** This setting uses three hand-written sentences across all 15 examined datasets. We do this by constructing general LLM-prompt templates which are filled in with the category names for each dataset. Our three general templates are as follows:

Describe what a/the \_\_\_ looks like:  
Describe a/the \_\_\_ :  
What are the identifying characteristics of a/the \_\_\_ ?

The blank portion of this template is either filled in with the category type plus the category name (e.g. “pet” + {} for the Oxford Pets dataset [38] or “aircraft” + {} for FGVC Aircraft [32]) or just the category name for more general datasets like ImageNet [13]. Type specification is necessary because of words that have multiple meanings. For example “boxer” from the Oxford Pets dataset can also mean a person who boxes, as opposed to a dog breed, so it is necessary to specify “Describe a pet boxer:”. Similarly, “Tornado” from the FGVC Aircraft dataset can be a type of aircraft or a type of weather.

**2. CuPL (full).** In this setting we use different LLM-prompt templates for each dataset, just as [42] uses different image-prompt templates for each dataset. However, we use fewer hand-written templates overall and also contain less specific information about each dataset in the templates. For this work, each dataset has between 2 and 9 LLM-prompts which generate between 20 and 90 image-prompt per category (10 generated sentences per LLM-prompt). For ImageNet, we use the following 5 LLM-prompts: (1) “Describe what a(n) {} looks like”, (2) “How can you identify a(n) {}?”, (3) “What does a(n) {} look like?”, (4) “A caption of an image of a(n) {}”, (5) “Describe an image from the internet of a(n) {}”. Full LLM-prompts for all datasets as well as example image-prompts are given the Appendix.

	ImageNet	DTD	Stanford Cars	SUN397	Food101	FGVC Aircraft	Oxford Pets	Caltech101	Flowers 102	UCF101	Kinetics-700	RESISC45	CIFAR-10	CIFAR-100	Birdsnap	mean	Total	Unique
std	75.54	55.20	77.53	69.31	93.08	32.88	93.33	93.24	78.53	77.45	60.07	71.10	95.59	78.26	50.43	73.43		
# hw	80	8	8	2	1	2	1	34	1	48	28	18	18	18	1		268	175
CuPL (base)	76.19	58.90	76.49	72.74	93.33	36.69	93.37	93.45	78.83	77.74	60.24	68.96	95.81	78.47	51.11	74.15		
$\Delta$ std	+0.65	+3.70	-1.04	+3.43	+0.25	+3.81	+0.04	+0.21	+0.30	+0.29	+0.17	-2.14	+0.22	+0.21	+0.63			
# hw	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3		45	3
CuPL (full)	76.69	61.70	77.63	73.31	93.36	36.11	93.81	93.45	79.67	78.36	60.63	71.69	95.84	78.57	51.11	74.80		
$\Delta$ std	+1.15	+6.50	+0.10	+4.00	+0.28	+3.23	+0.48	+0.21	+1.14	+0.91	+0.56	+0.59	+0.25	+0.31	+0.63			
# hw	5	6	9	3	3	2	2	3	2	5	4	5	3	4	3		59	45

Table 1. **Performance of CuPL prompts compared to the standard, hand-written prompts in CLIP [42] on 15 zero-shot image classification benchmarks.** “ $\Delta$ std” stands for the difference; green shows improvement. In addition to accuracy, we show number of prompt templates that are hand-written (“# hw”) for each dataset using each method, as well as the total and unique number of hand-written templates for each method (unique number only counts templates once even if used for multiple datasets). Note that CuPL (base) uses just three hand-constructed sentence across all datasets compared to 175 in the standard method.

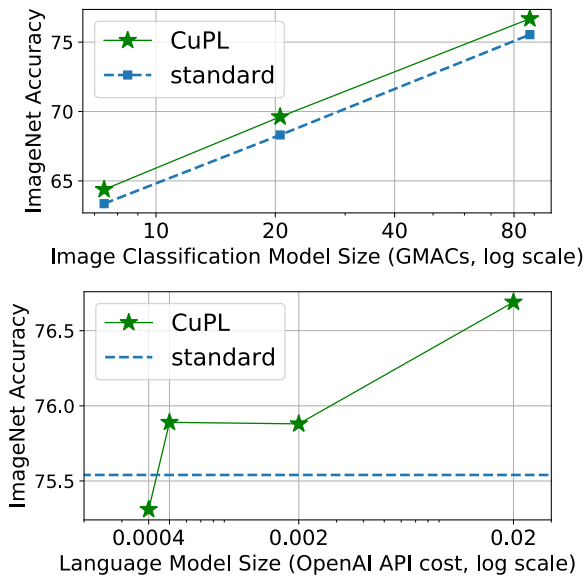


Figure 3. **Performance of CuPL as models scale.** (Top) ImageNet Top-1 accuracy for various scales of CLIP. CuPL prompts remain consistently better than standard prompts even we adjust CLIP model size (ViT-B/32, ViT-B/16, ViT-L/14). GPT-3 model set as DaVinci-002. (Bottom) ImageNet Top-1 accuracy for various scales of GPT-3 (ada, babbage, curie, davinci-002). Larger models produce higher accuracy. CLIP model set as ViT-L/14.

### 3. Experiments and Results

We first discuss the details of our experimental setup. We next show improvements on a wide range of image classification benchmarks. We then examine the scaling behavior with respect to the model size and report obser-

vations regarding hyperparameters such as the number of hand-written prompts. Finally, we provide evidence that the model is able to use CuPL prompts to place more importance on the most relevant parts of the image.

#### 3.1. Setup

Unless specified otherwise, we use CLIP with a backbone of ViT-L/14 [14] and the GPT-3 DaVinci-002 model. Additionally, in order to perform open-vocabulary image classification, each image category needs a natural language label. This is sometimes provided by the dataset, but not always (e.g. ImageNet categories are described by an id number which can map to multiple synonyms). For this work, we use the same natural language labels specified in [42].

We report our findings on 15 zero-shot image recognition benchmarks: ImageNet [13], Describable Textures Dataset (DTD) [11], Stanford Cars [26], Scene UNderstanding (SUN397) [60], Food101 [4], FGVC Aircraft [32], Oxford Pets [38], Caltech101 [16], Flowers 102 [36], UCF101 [52], Kinetics-700 [8], Remote Sensing Image Scene Classification (RESISC45) [10], CIFAR-10 [27], CIFAR-100 [27], and Birdsnap [2]. For the two video datasets, we extract the middle frame of the video, as is done in Radford *et al.* [42].

#### 3.2. Results

Our results for the base prompts setting and the full prompts setting are in Table 1. We present our method’s performance on 15 different image classification benchmarks, comparing both the classification accuracy and the number of hand-written sentence templates needed for each method. Note that for the standard method [42], the hand-written sentences refer to the image-prompts, while for

CuPL the hand-written sentences refer to the LLM-prompts, with which image-prompts are generated.

**1. CuPL (base).** In this setting, we see performance gains in 13 out of the 15 examined datasets. Note this setting uses *just three hand-constructed sentence across all datasets*. This is in comparison to the nearly 175 unique image-prompt templates that are hand-written across all of these datasets in the standard setting. Additionally, in the standard setting these hand-constructed prompts must be very specific to the dataset (e.g., “a black and white photo of a {}.”, “a plastic {}.”). In comparison, CuPL (base) requires only the category type of the overall dataset and still outperforms the hand-written, domain specified baseline in almost all cases. Thus, we present this base prompt setting as a simple standard that matches or exceeds prompt engineering open-vocabulary models.

**2. CuPL (full prompts).** Here we see improvements on all examined datasets. This includes large (over 1 percentage point) gains on ImageNet Top-1, DTD (texture classification), SUN397 (scene classification), FGVC Aircraft (fine-grained aircraft classification), and Flowers 102 (flower classification). While this setting requires more hand-written prompts than setting (1), it still requires significantly fewer than the baseline method (5 sentences versus 80 sentence for ImageNet), and does not include knowledge about the image domain. The full list of hand-constructed sentences for CuPL (full prompts) and the baseline method [42] can be found in the Appendix.

### 3.3. Analysis and Ablations

**Other prompting techniques.** Concurrent work by Menon *et al.* [33] also explores LLM generated descriptions for image classification. This work differs from CuPL as it generates a structured list of identifying attributes in a single generation, which are reformatted into multiple sentences. In contrast, CuPL outputs a single sentence for multiple generations, with no enforced format. The benefit of the structured output used in Menon *et al.* [33] is that the authors can examine the similarity of a given image with each individual attribute to understand which ones most contribute to a prediction. However, unlike CuPL, this method performs worse than standard human-written prompts, as shown in Table 2. This is potentially because this work focuses on explainability, and therefore enforces a strict format on the generated prompts, likely reducing overall accuracy.

**Model Size.** In Figure 3, we show CuPL (full prompts) at different model scales. As there are two different zero-shot models in the CuPL algorithm, we show the effects of varying each model individually. On the top, we vary the CLIP model used while holding the LLM constant. We see consistent gains across all model sizes. On the bottom, we vary the size of the LLM. We plot the accuracy of the

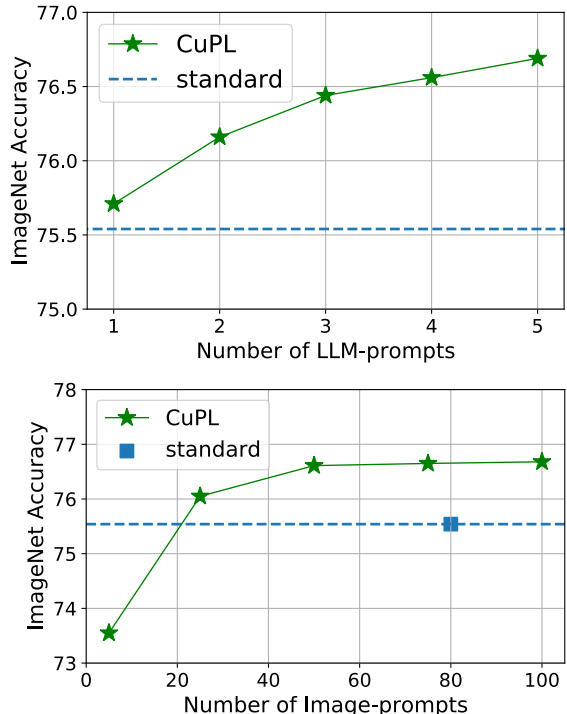


Figure 4. **Ablation on number of LLM-prompts (top) and image-prompts (bottom).** (Top) As number of hand-written LLM-prompts increases, so does accuracy. 10 image-prompts are generated for each LLM-prompt. Note that CuPL outperforms the baseline even with just one hand-written sentence. We add the prompts in a greedy manner, at each step adding the 10 prompts which lead to the largest performance gain. (Bottom) We adjust the number of image-prompts generated by a fixed number (5) of LLM-prompts. Even at 5 Image-prompts per LLM-prompt (25 prompts total), we outperform the baseline which uses 80 image-prompts.

baseline as well, which does not vary as it does not utilize an LLM. We find larger models lead to higher accuracy, though the 2nd and 3rd largest models perform similarly.

**Number of Prompts.** In Figure 4, we present ablations on the number of LLM-prompts and image-prompts for CuPL (full prompts). On the top, we show ImageNet accuracy as we increase the number of LLM-prompts. This also corresponds to the number of sentences that have to be hand-written. Notably, this method outperforms the baseline even when using prompts generated from a single hand-written sentence. On the bottom, we hold the number of LLM-prompts constant at 5 and adjust how many image-prompts we generate per LLM-prompt. We plot the accuracy given the total number of image-prompts (so 10 generated image-prompt per LLM-prompt corresponds to 50 total image-prompts). We see that CuPL begins to outperform the baseline at just 25 image-prompts, well below the 80 image-prompts used in the baseline.

**Additional Analysis.** In the Appendix, we provide com-








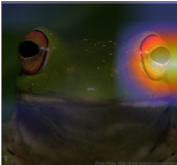


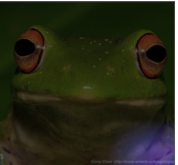

Original Image	Region Importance with CuPL Prompts		Region Importance with Baseline Prompts		Example Image of Distractor Class
(A)	(B)	(C)	(D)	(E)	(F)
<b>GT Label:</b> Schipperke 	<b>Schipperke Prompt:</b> "A Schipperke is a small, black Belgian dog with pointy ears and an <b>upright tail.</b> " 	<b>Groenendael Prompt:</b> "A Groenendael dog can be identified by its black coat and erect ears." 	<b>Schipperke Prompt:</b> "A photo of a Schipperke" 	<b>Groenendael Prompt:</b> "A photo of a Groenendael dog" 	<b>Example:</b> Groenendael dog 
	Prediction: Schipperke ✓		Prediction: Groenendael dog ✗		
<b>GT Label:</b> Tree Frog 	<b>Tree Frog Prompt:</b> "A tree frog looks like a small frog with <b>large eyes.</b> " 	<b>Tailed Frog Prompt:</b> "The tailed frog is a small frog that is found in North America." 	<b>Tree Frog Prompt:</b> "A photo of a tree frog" 	<b>Tailed Frog Prompt:</b> "A photo of a tailed frog" 	<b>Example:</b> Tailed frog 
	Prediction: Tree Frog ✓		Prediction: Tailed Frog ✗		

Figure 5. **CuPL prompts lead the model to focus on semantically important regions of the image.** We use Shapley values (Section 3.4) to visualize the importance of each region in a binary classification problem. We examine which parts of an image lead the model to classify it as the correct class versus a commonly confused class. We present the original image (column A), as well as four heatmaps showing which regions raise the probability of the *correct* class for the *CuPL* model (column B), the *incorrect* class for the *CuPL* model (column C), the *correct* class for the *baseline* model (column D), and the *incorrect* class for the *baseline* model (column E). Additionally, we show that the regions that are more important to CuPL than to the baseline correspond to regions mentioned in the CuPL prompts (i.e. “tail” for Schipperke dog and “eyes” for Tree frog). We also show an example image from the distractor class to demonstrate the level of similarity between these fine-grained classes (column F). Finally, we see that CuPL scores the correct class higher, whereas the baseline scores the incorrect class higher. This series of observations lead us to believe that CuPL is able to correct errors because the descriptive prompts cause the model to weigh semantically important regions more heavily.

Standard	CuPL	Menon <i>et al.</i> [33]
75.54	76.69	75.00

Table 2. **Comparison with Menon *et al.* [33]** on Top-1 Imagenet accuracy with ViT L/14.

parisons between CuPL prompts and descriptive prompts generated with definitions of ImageNet classes as well as with Wikipedia descriptions of ImageNet classes. We find that CuPL prompts outperform both of these baselines. Additionally, we provide results of ensembling CuPL prompts and the baseline hand-written prompts used in [42]. We find that this ensemble outperforms just baseline prompts for all datasets, and outperforms just CuPL prompts for some datasets.

### 3.4. Shapley Value Analysis

We show that CuPL descriptions allow CLIP to place more importance on image regions that are most relevant for the correct classification. In order to measure the importance of regions in the image, we invoke Shapley values [49], a tool from game theory that has become popular for understanding which input information contributes to a model’s final prediction [31, 9]. Shapley values can be computed for any model, and although there are methods designed specifically for vision transformers [12] (the underlying architecture of CLIP), we use a simple model-agnostic calculation [35]. We employ Shapley values to understand the importance of different image regions with CuPL prompts versus baseline prompts, and we find that CuPL places more value on regions that are emphasized in object descriptions, and thus are likely important for obtaining correct classifications. We demonstrate this correlation

in two ways: (1) visualizing heatmaps of importance over images, and (2) measuring the importance of segmented image parts annotated by the PartImageNet Dataset [18].

**Importance Heatmaps** To understand how CuPL captions lead to a change in importance of different image regions, we calculate the Shapley value of small image patches when using CuPL prompts versus when using baseline prompts. We calculate the Shapley values with respect to a binary classification probability between the correct class and a similar distractor class in order to understand how CuPL corrects these errors. As shown in Figure 5, we examine the important regions of an image of a dog when classifying between two very similar dog categories: a “Schipperke dog” versus a “Groenendael dog”. Both of these classes are Belgian dogs that are black with pointy ears. However, they have a few subtle differences including the typical appearance of their tails. Additionally, we show the important regions of an image when classifying between a “Tree frog” and a “Tailed frog”, which also look very similar.

For each binary classification, we show four heatmaps: (1) the regions that contribute to a higher probability of the correct class when using CuPL prompts, (2) the regions that contribute to a higher probability of the incorrect class when using CuPL prompts, (3) the regions that contribute to a higher probability of the correct class when using baseline prompts, (4) the regions that contribute to a higher probability of the incorrect class when using baseline prompts. Interestingly, we find that not only does CuPL place importance on different regions of the image, but these regions correspond to descriptions in the CuPL prompts. For example, the tail of the dog is very important to the “Schipperke” probability when using CuPL prompts, but not when using baseline prompts, and the tail of the Schipperke dog is described 10 times in the CuPL descriptions of this class. Similarly, we find that the eyes in the image of the frog are much more important when classifying with CuPL than with the baseline, and that the eyes are mentioned 10 times in the CuPL description of a tree frog. We provide more examples of this phenomenon in the Appendix.

**Importance of Segmented Parts** In order to understand the correlation between the importance of an image region and its frequency in CuPL prompts on a larger scale, we utilize the PartImageNet Dataset [18]. This dataset contains segmentation maps of the different parts of a class for a subset of ImageNet classes. For example, the dog classes have the parts: ‘head’, ‘body’, ‘leg’ and ‘tail’. We use these segmentation maps to obtain the Shapley value for each part of the animal with respect to the final probability of the ground truth class. To understand the effect of changing to CuPL prompts, we calculate the difference between the Shapley values with CuPL prompts and with baseline prompts, and we average across all images in a class. So for each part in

each examined class we calculate the following (where SV denotes the Shapley value):

$$\frac{1}{|\text{class}|} \sum_{\text{image} \in \text{class}} SV_{\text{CuPL}}(\text{image}, \text{part}) - SV_{\text{base}}(\text{image}, \text{part})$$

This gives us a score for how much more important a part of an animal is to CuPL compared to the baseline for classification. Additionally, we quantify how prevalent each body part is in the CuPL descriptions. We do this using the WordNet [34] database to tag each words as part of the ‘leg’, ‘head’, etc. More details of this tagging system are given in the Appendix. We present our findings in Figure 6. We find that the parts that are more important to CuPL are highly correlated with the parts that are present in the descriptions of the animals (and thus likely important to the identification of the animal). For example, head-related attributes of the Japanese Spaniel class are frequently mentioned in the descriptions. Additionally, the ‘head’ in the image is much more important to the final prediction for CuPL than for baseline. Thus, CuPL is able to extract important information for identifying the animal from the text and incorporate it into classification predictions.

## 4. Related Work

### 4.1. Natural Language Descriptions for Image Classification

Several prior works use text-based knowledge of image categories to improve classification accuracy. [15] extract visual information from unstructured text descriptions collected from the internet to recognize parts of object and classify them in a zero-shot way. [45] and [19] use natural language descriptions of bird types to train a multimodal classification model. [21] use hand-collected attribute tags to attend over relevant features in images. [39] extract visual information from Wikipedia descriptions to enable zero-shot bird classification. Additional works [50, 6] show improvements on large datasets (e.g., ImageNet) using external information from external databases such as Imagenet-wiki and Wordnet. While these works show the effectiveness of augmenting zero-shot models with descriptive text, all of these prior works rely on external natural language databases for descriptions. This often limits the possible categories that can be classified and can require extensive preprocessing to extract visual descriptions from noisy natural language.

### 4.2. Generated Text for Downstream Tasks

Recent work has utilized text generated from LLMs in a number of ways. [47] use an LLM to paraphrase existing image captions to use as data augmentation for CLIP. [30] use GPT-3 to generate knowledge on a topic when given a number of demonstrations, which is then used to improve

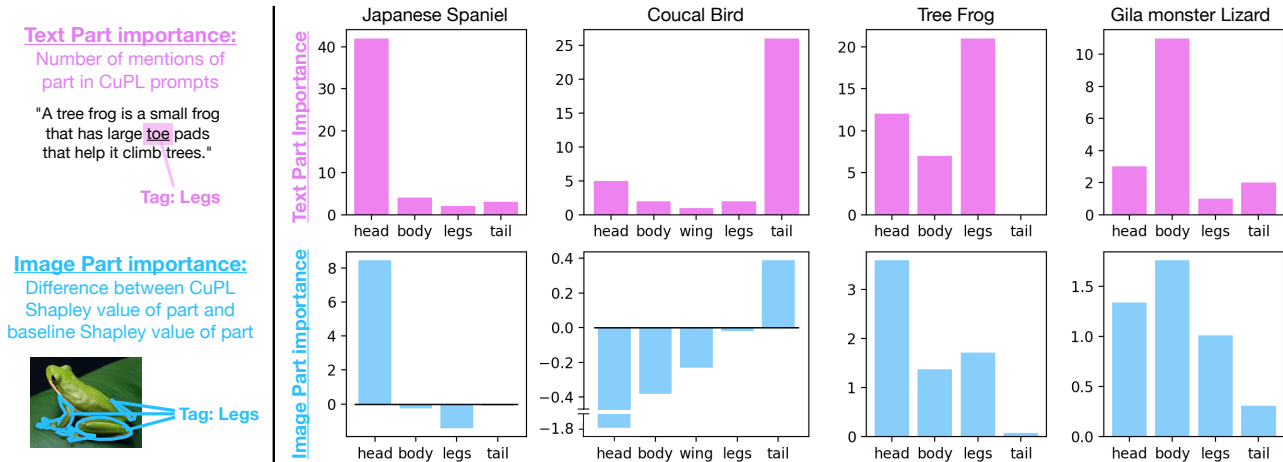


Figure 6. When specific parts of an animal/object are frequently mentioned in CuPL prompts, the CuPL model places more importance on these parts in the image compared to the baseline model. The PartImageNet dataset [18] provides segmentation maps of ImageNet images broken down into parts. For example, Tree Frog is broken down into the parts: ‘head’, ‘leg’, ‘body’ and ‘tail’. We use the WordNet database [34] to tag words in CuPL prompts as belonging to one of these parts. We refer to the number of mentions of the part as the Text Part Importance. We then use the PartImageNet segmentations to compare the Shapley value of each part when using CuPL prompts and baseline prompts, which we call the Image Part Importance. We find a strong correlation between the Text Part Importance and the Image Part Importance, leading to the conclusion that CuPL is able to take advantage of the knowledge contained in the descriptions when making its predictions.

accuracy on common sense reasoning questions. [20] use a LLM to add labels to text to improve text classification accuracy. In [64], the outputs of a GPT-2 model are used to train an encoder on top of a vision model to generate multimodal image representations for a variety of tasks. [53] utilize a language model to perform image captioning by iteratively generating candidate image captions with a LLM and then using feedback from an open-vocabulary model to align it to a given image. Similarly, [62] use GPT-3 along with text descriptions of images for the Visual Question Answering (VQA) task. However, unlike CuPL these prior works are either purely language tasks (common sense reasoning, text classification) or multimodal with some language component (image captioning, VQA). Most similarly, [33] use an LLM to generate a structured list of attributes which are then reformatted into captions for CLIP. However this work differs from ours as it does not improve over human written templates. Additionally, [61] use an LLM to generate a list of natural language attributes for ImageNet classes and then select a subset of these attributes for each class in a few-shot manner. Our work differs from this as we remain in the zero-shot setting.

### 4.3. Prompt Engineering

Previous efforts have explored methods for obtaining successful natural language prompts. For both open-vocabulary image classification models as well as LLMs, the format of prompts is known to highly affect accuracy [48, 42, 5, 17]. This has led to a large effort to find op-

timal prompt formats. Proposed methods include crowdsourcing high performing prompts [1] as well as framing prompts to induce models to give explanations as well as answers [57, 25, 37]. Additional works have proposed learning prompts via gradient based methods [65, 41, 29, 28, 51], retrieval from a database [46], or reformatting/rephrasing existing prompts [24, 46].

Most relevant to this work are a number of methods for designing optimal prompts for zero-shot image classification with open-vocabulary models. These methods learn prompts formats which yield high accuracy for image classification using either supervised [66, 43] or unsupervised [22] methods. However, unlike these prior works this work requires no additional training or labeled data.

## 5. Conclusion

We demonstrate that leveraging knowledge from an LLM can immediately improve zero-shot accuracy on a variety of image classification tasks, with much less hand-engineering efforts to craft natural language prompts. Furthermore, prompts can be customized to the desired categories, rather than a general template that applies to all categories. Finally, using prompts generated by LLMs lowers the barrier of prior knowledge about the dataset, which is often required when crafting prompt templates.

Querying an LLM for prompt construction is simple, straightforward and as our results suggested, immediately beneficial. The hypothesis that a joint force of LLMs and open vocabulary models would improve zero-shot image



classification is thoroughly tested in this work. We hope these findings serve as a useful tool towards understanding and improving zero-shot image classification, and more generally, the consolidation of model capacities and modalities through natural language.

## References

- [1] Stephen H. Bach, Victor Sanh, Zheng Xin Yong, Albert Webson, Colin Raffel, Nihal V. Nayak, Abheesht Sharma, Taewoon Kim, M SAIFUL BARI, Thibault Févry, Zaid Alyafeai, Manan Dey, Andrea Santilli, Zhiqing Sun, Srulik Ben-David, Canwen Xu, Gunjan Chhablani, Han Wang, Jason Alan Fries, Maged S. Al-shaibani, Shanya Sharma, Urmish Thakker, Khalid Almubarak, Xiangru Tang, Mike Tian-Jian Jiang, and Alexander M. Rush. Promptsources: An integrated development environment and repository for natural language prompts. *ArXiv*, abs/2202.01279, 2022. 8
- [2] Thomas Berg, Jiongxin Liu, Seung Woo Lee, Michelle L Alexander, David W Jacobs, and Peter N Belhumeur. Birdsnap: Large-scale fine-grained visual categorization of birds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2011–2018, 2014. 4
- [3] Lucas Beyer, Olivier J Hénaff, Alexander Kolesnikov, Xiaoohua Zhai, and Aäron van den Oord. Are we done with imagenet? *arXiv preprint arXiv:2006.07159*, 2020. 24
- [4] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – mining discriminative components with random forests. In *European Conference on Computer Vision*, 2014. 4
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 3, 8, 20, 21
- [6] Sebastian Bujwid and Josephine Sullivan. Large-scale zero-shot image classification from rich and diverse textual descriptions. *ArXiv*, abs/2103.09669, 2021. 7
- [7] Sebastian Bujwid and Josephine Sullivan. Large-scale zero-shot image classification from rich and diverse textual descriptions. *arXiv preprint arXiv:2103.09669*, 2021. 19
- [8] Joao Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *arXiv preprint arXiv:1907.06987*, 2019. 4
- [9] Hugh Chen, Ian C Covert, Scott M Lundberg, and Su-In Lee. Algorithms to estimate Shapley value feature attributions. *arXiv preprint arXiv:2207.07605*, 2022. 6
- [10] Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017. 4
- [11] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, , and A. Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014. 2, 4
- [12] Ian Covert, Chanwoo Kim, and Su-In Lee. Learning to estimate Shapley values with vision transformers. *arXiv preprint arXiv:2206.05282*, 2022. 6
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 2, 3, 4, 22
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 4
- [15] Mohamed Elhoseiny, Yizhe Zhu, Han Zhang, and A. Elgammal. Link the head to the ”beak”: Zero shot learning from noisy text description at part precision. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6288–6297, 2017. 7
- [16] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *2004 conference on computer vision and pattern recognition workshop*, pages 178–178. IEEE, 2004. 4
- [17] Tianyu Gao, Adam Fisch, and Danqi Chen. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*, 2020. 8
- [18] Ju He, Shuo Yang, Shaokang Yang, Adam Kortylewski, Xiaoding Yuan, Jie-Neng Chen, Shuai Liu, Cheng Yang, Qihang Yu, and Alan Yuille. Partimagenet: A large, high-quality dataset of parts. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VIII*, pages 128–145. Springer, 2022. 7, 8, 22
- [19] Xiangteng He and Yuxin Peng. Fine-grained image classification via combining vision and language. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7332–7340, 2017. 7
- [20] Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Juan-Zi Li, and Maosong Sun. Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer for text classification. In *ACL*, 2022. 8
- [21] Siteng Huang, Min Zhang, Yachen Kang, and Donglin Wang. Attributes-guided and pure-visual attention alignment for few-shot recognition. In *AAAI*, 2021. 7
- [22] Tony Huang, Jack Chu, and Fangyun Wei. Unsupervised prompt learning for vision-language models. *arXiv preprint arXiv:2204.03649*, 2022. 8
- [23] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021. 1, 3
- [24] Zhengbao Jiang, Frank F. Xu, J. Araki, and Graham Neubig. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 8:423–438, 2020. 8

- [25] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *ArXiv*, abs/2205.11916, 2022. 8
- [26] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013. 4
- [27] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 4
- [28] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *ArXiv*, abs/2104.08691, 2021. 8
- [29] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, abs/2101.00190, 2021. 8
- [30] Jiacheng Liu, Alisa Liu, Ximing Lu, Sean Welleck, Peter West, Ronan Le Bras, Yejin Choi, and Hannaneh Hajishirzi. Generated knowledge prompting for commonsense reasoning. In *ACL*, 2022. 7
- [31] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 2017. 6
- [32] S. Maji, J. Kannala, E. Rahtu, M. Blaschko, and A. Vedaldi. Fine-grained visual classification of aircraft. Technical report, 2013. 2, 3, 4
- [33] Sachit Menon and Carl Vondrick. Visual classification via description from large language models. *arXiv preprint arXiv:2210.07183*, 2022. 5, 6, 8
- [34] George A Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41, 1995. 7, 8, 19, 22, 24, 25
- [35] Rory Mitchell, Joshua Cooper, Eibe Frank, and Geoffrey Holmes. Sampling permutations for Shapley value estimation. *Journal of Machine Learning Research*, 23(43):1–46, 2022. 6
- [36] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729. IEEE, 2008. 4
- [37] Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and Augustus Odena. Show your work: Scratchpads for intermediate computation with language models. *ArXiv*, abs/2112.00114, 2021. 8
- [38] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3498–3505. IEEE, 2012. 3, 4
- [39] Tzuf Paz-Argaman, Yuval Atzmon, Gal Chechik, and Reut Tsarfaty. Zest: Zero-shot learning from text descriptions using textual similarity and visual summarization. In *FINDINGS*, 2020. 7
- [40] Hieu Pham, Zihang Dai, Golnaz Ghiasi, Kenji Kawaguchi, Hanxiao Liu, Adams Wei Yu, Jiahui Yu, Yi-Ting Chen, Minh-Thang Luong, Yonghui Wu, et al. Combined scaling for open-vocabulary image classification. *arXiv preprint arXiv:2111.10050*, 2021. 1, 3
- [41] Guanghui Qin and Jas’ Eisner. Learning how to ask: Querying lms with mixtures of soft prompts. *ArXiv*, abs/2104.06599, 2021. 8
- [42] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 1, 2, 3, 4, 5, 6, 8, 12, 20, 22, 23
- [43] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. Denseclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18082–18091, 2022. 8
- [44] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishal Shankar. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning*, pages 5389–5400. PMLR, 2019. 22
- [45] Scott E. Reed, Zeynep Akata, Honglak Lee, and Bernt Schiele. Learning deep representations of fine-grained visual descriptions. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 49–58, 2016. 7
- [46] Ohad Rubin, Jonathan Herzig, and Jonathan Berant. Learning to retrieve prompts for in-context learning. In *NAACL*, 2022. 8
- [47] Shibani Santurkar, Yann Dubois, Rohan Taori, Percy Liang, and Tatsunori Hashimoto. Is a caption worth a thousand images? a controlled study for representation learning. *arXiv preprint arXiv:2207.07635*, 2022. 7
- [48] Timo Schick and Hinrich Schütze. Exploiting cloze-questions for few-shot text classification and natural language inference. In *EACL*, 2021. 8
- [49] Lloyd S Shapley et al. A value for n-person games. 1953. 6
- [50] Sheng Shen, Chunyuan Li, Xiaowei Hu, Yujia Xie, Jianwei Yang, Pengchuan Zhang, Anna Rohrbach, Zhe Gan, Lijuan Wang, Lu Yuan, Ce Liu, Kurt Keutzer, Trevor Darrell, and Jianfeng Gao. K-lite: Learning transferable visual models with external knowledge. *ArXiv*, abs/2204.09222, 2022. 7
- [51] Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. Eliciting knowledge from language models using automatically generated prompts. *ArXiv*, abs/2010.15980, 2020. 8
- [52] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 4
- [53] Yixuan Su, Tian Lan, Yahui Liu, Fangyu Liu, Dani Yogatama, Yan Wang, Lingpeng Kong, and Nigel Collier. Language models can see: Plugging visual controls in text generation. *ArXiv*, abs/2205.02655, 2022. 8

- [54] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. [23](#), [24](#), [25](#), [26](#)
- [55] Ben Wang and Aran Komatsuzaki. Gpt-j-6b: A 6 billion parameter autoregressive language model, 2021. [20](#)
- [56] Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. In *Advances in Neural Information Processing Systems*, pages 10506–10518, 2019. [22](#)
- [57] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. *ArXiv*, abs/2201.11903, 2022. [8](#)
- [58] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*, 2019. [20](#)
- [59] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7959–7971, 2022. [22](#)
- [60] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3485–3492. IEEE, 2010. [4](#)
- [61] Yue Yang, Artemis Panagopoulou, Shenghao Zhou, Daniel Jin, Chris Callison-Burch, and Mark Yatskar. Language in a bottle: Language model guided concept bottlenecks for interpretable image classification. *arXiv preprint arXiv:2211.11158*, 2022. [8](#)
- [62] Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Yumao Lu, Zicheng Liu, and Lijuan Wang. An empirical study of gpt-3 for few-shot knowledge-based vqa. In *AAAI*, 2022. [8](#)
- [63] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022. [1](#), [3](#)
- [64] Youngjae Yu, Jiwan Chung, Heeseung Yun, Jack Hessel, Jae Sung Park, Ximing Lu, Prithviraj Ammanabrolu, Rowan Zellers, Ronan Le Bras, Gunhee Kim, and Yejin Choi. Multimodal knowledge alignment with reinforcement learning. *ArXiv*, abs/2205.12630, 2022. [8](#)
- [65] Ningyu Zhang, Luoqiu Li, Xiang Chen, Shumin Deng, Zhen Bi, Chuanqi Tan, Fei Huang, and Huajun Chen. Differentiable prompt makes pre-trained language models better few-shot learners. *ArXiv*, abs/2108.13161, 2021. [8](#)
- [66] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, pages 1–12, 2022. [8](#)

# APPENDIX: *What does a platypus look like?*

## Generating customized prompts for zero-shot image classification

### Overview

- A. CuPL (Full Prompts) vs Standard Prompts**
- B. CuPL Base Prompts**
- C. Evaluation Metric**
- D. Comparison with Definitions and Wikipedia Descriptions**
- E. Ensembling with Standard Prompts**
- F. Additional Qualitative Results**
- G. Open-Source LLM and ChatGPT**
- H. Single Sentence Baseline**
- I. Robustness**
- J. Tagging with ImageNet Parts details**
- K. CuPL Improvement Analysis**
- L. Error Analysis**
- M. Image-Prompt Distribution**
- N. Temperature Analysis**
- O. Example Generated image-prompts**

### A. CuPL (Full Prompts) vs Standard Prompts

We detail the hand-written prompt templates used for CuPL (full prompts) versus standard CLIP [42] prompt templates. For CuPL, hand-written prompt templates are needed for the LLM-prompts, while for the standard method hand-written prompt templates are needed for the image-prompts.

Note that many of the hand-written templates for the standard method encode information about the datasets. For example, "a toy {}" demonstrates knowledge that objects are sometimes represented as a toy version of an object rather than as the literal object. CuPL prompts remain much more general (e.g. "Describe what a {} looks like").

#### Caltech101

CuPL hand-written	Standard hand-written	
Describe what a(n) {} looks like:	a photo of a {}.	a photo of the {}.
Describe a(n) {}:	a painting of a {}.	a painting of the {}.
What are the identifying characteristics of a(n) {}?	a plastic {}.	the plastic {}.
	a sculpture of a {}.	a sculpture of the {}.
	a sketch of a {}.	a sketch of the {}.
	a tattoo of a {}.	a tattoo of the {}.
	a toy {}.	the toy {}.
	a rendition of a {}.	a rendition of the {}.
	a embroidered {}.	the embroidered {}.
	a cartoon {}.	the cartoon {}.
	a {} in a video game.	the {} in a video game.
	a plushie {}.	the plushie {}.
	a origami {}.	the origami {}.
	art of a {}.	art of the {}.
	graffiti of a {}.	graffiti of the {}.
	a drawing of a {}.	a drawing of the {}.
	a doodle of a {}.	a doodle of the {}.

## Food101

CuPL hand-written	Standard hand-written
Describe what {} looks like Visually describe {} How can you tell that the food in this photo is {}?	a photo of {}, a type of food.

## Stanford Cars

CuPL hand-written	Standard hand-written
How can you identify a(n) {}? Description of a(n) {}, a type of car. A caption of a photo of a(n) {}: What are the primary characteristics of a(n) {}? Description of the exterior of a(n) {} What are the identifying characteristics of a(n) {}, a type of car? Describe an image from the internet of a(n) {} What does a(n) {} look like? Describe what a(n) {}, a type of car, looks like	a photo of a {}. a photo of the {}. a photo of my {}. i love my {}! a photo of my dirty {}. a photo of my clean {}. a photo of my new {}. a photo of my old {}.

## Oxford Pets

CuPL hand-written	Standard hand-written
Describe what a pet {} looks like Visually describe a(n) '{}', a type of pet.	a photo of a {}, a type of pet.

## ImageNet

CuPL hand-written	Standard hand-written
Describe what a(n) {} looks like How can you identify a(n) {}? What does a(n) look like? Describe an image from the internet of a(n) {} A caption of an image of a(n) {}:	<p>a bad photo of a {}.</p> <p>a photo of many {}.</p> <p>a sculpture of a {}.</p> <p>a photo of the hard to see {}.</p> <p>a low resolution photo of the {}.</p> <p>a rendering of a {}.</p> <p>graffiti of a {}.</p> <p>a bad photo of the {}.</p> <p>a cropped photo of the {}.</p> <p>a tattoo of a {}.</p> <p>the embroidered {}.</p> <p>a photo of a hard to see {}.</p> <p>a bright photo of a {}.</p> <p>a photo of a clean {}.</p> <p>a photo of a dirty {}.</p> <p>a dark photo of the {}.</p> <p>a drawing of a {}.</p> <p>a photo of my {}.</p> <p>the plastic {}.</p> <p>a photo of the cool {}.</p> <p>a close-up photo of a {}.</p> <p>a black and white photo of the {}.</p> <p>a painting of the {}.</p> <p>a painting of a {}.</p> <p>a pixelated photo of the {}.</p> <p>a sculpture of the {}.</p> <p>a bright photo of the {}.</p> <p>a cropped photo of a {}.</p> <p>a plastic {}.</p> <p>a photo of the dirty {}.</p> <p>a jpeg corrupted photo of a {}.</p> <p>a blurry photo of the {}.</p> <p>a photo of the {}.</p> <p>a good photo of the {}.</p> <p>a rendering of the {}.</p> <p>a {} in a video game.</p> <p>a photo of one {}.</p> <p>a doodle of a {}.</p> <p>a close-up photo of the {}.</p> <p>a photo of a {}.</p> <p>the origami {}.</p> <p>the {} in a video game.</p> <p>a sketch of a {}.</p> <p>a doodle of the {}.</p> <p>a origami {}.</p> <p>a low resolution photo of a {}.</p> <p>the toy {}.</p> <p>a rendition of the {}.</p> <p>a photo of the clean {}.</p> <p>a photo of a large {}.</p> <p>a rendition of a {}.</p> <p>a photo of a nice {}.</p> <p>a photo of a weird {}.</p> <p>a blurry photo of a {}.</p> <p>a cartoon {}.</p> <p>art of a {}.</p> <p>a sketch of the {}.</p> <p>a embroidered {}.</p> <p>a pixelated photo of a {}.</p> <p>itap of the {}.</p> <p>a jpeg corrupted photo of the {}.</p> <p>a good photo of a {}.</p> <p>a plushie {}.</p> <p>a photo of the nice {}.</p> <p>a photo of the small {}.</p> <p>a photo of the weird {}.</p> <p>the cartoon {}.</p> <p>art of the {}.</p> <p>a drawing of the {}.</p> <p>a photo of the large {}.</p> <p>a black and white photo of a {}.</p> <p>the plushie {}.</p> <p>a dark photo of a {}.</p> <p>itap of a {}.</p> <p>graffiti of the {}.</p> <p>a toy {}.</p> <p>itap of my {}.</p> <p>a photo of a cool {}.</p> <p>a photo of a small {}.</p> <p>a tattoo of the {}.</p>

## FGVC Aircraft

CuPL hand-written	Standard hand-written
Describe a(n) {} aircraft Describe the {} aircraft	<p>a photo of a {}, a type of aircraft.</p> <p>a photo of the {}, a type of aircraft.</p>

## DTD

CuPL hand-written	Standard hand-written
What does “{}” material look like? What does a “{}” surface look like? What does a “{}” texture look like? What does a “{}” object look like? What does a “{}” thing look like? What does a “{}” pattern look like?	<p>a photo of a {} texture.</p> <p>a photo of a {} pattern.</p> <p>a photo of a {} thing.</p> <p>a photo of a {} object.</p> <p>a photo of the {} texture.</p> <p>a photo of the {} pattern.</p> <p>a photo of the {} thing.</p> <p>a photo of the {} object.</p>

## SUN397

CuPL hand-written	Standard hand-written
Describe what a(n) {} looks like How can you identify a(n) {}? Describe a photo of a(n) {}	<p>a photo of a {}.</p> <p>a photo of the {}.</p>

## Kinetics-700

### CuPL hand-written

### Standard hand-written

Describe the action "{}"

What does a person {} look like?

What does the act of {} look like?

Describe "{}"

a photo of a person {}.

a photo of a person using {}.

a photo of a person doing {}.

a photo of a person during {}.

a photo of a person performing {}.

a photo of a person practicing {}.

a video of {}.

a video of a person {}.

a video of a person using {}.

a video of a person doing {}.

a video of a person during {}.

a video of a person performing {}.

a video of a person practicing {}.

a example of {}.

a example of a person {}.

a example of a person using {}.

a example of a person doing {}.

a example of a person during {}.

a example of a person performing {}.

a example of a person practicing {}.

a demonstration of {}.

a demonstration of a person {}.

a demonstration of a person using {}.

a demonstration of a person doing {}.

a demonstration of a person during {}.

a demonstration of a person performing {}.

a demonstration of a person practicing {}.

## UCF 101

### CuPL hand-written

Describe the action of {}  
What does the act of {} look like?  
What does a person doing {} look like?  
Describe “{}”  
Describe the action “{}”

### Standard hand-written

a photo of a person {}.  
a video of a person {}.  
a example of a person {}.  
a demonstration of a person {}.  
a photo of the person {}.  
a video of the person {}.  
a example of the person {}.  
a demonstration of the person {}.  
a photo of a person using {}.  
a video of a person using {}.  
a example of a person using {}.  
a demonstration of a person using {}.  
a photo of the person using {}.  
a video of the person using {}.  
a example of the person using {}.  
a demonstration of the person using {}.  
a photo of a person doing {}.  
a video of a person doing {}.  
a example of a person doing {}.  
a demonstration of a person doing {}.  
a photo of the person doing {}.  
a video of the person doing {}.  
a example of the person doing {}.  
a demonstration of the person doing {}.  
a photo of a person during {}.  
a video of a person during {}.  
a example of a person during {}.  
a demonstration of a person during {}.  
a photo of the person during {}.  
a video of the person during {}.  
a example of the person during {}.  
a demonstration of the person during {}.  
a photo of a person performing {}.  
a video of a person performing {}.  
a example of a person performing {}.  
a demonstration of a person performing {}.  
a photo of the person performing {}.  
a video of the person performing {}.  
a example of the person performing {}.  
a demonstration of the person performing {}.  
a photo of a person practicing {}.  
a video of a person practicing {}.  
a example of a person practicing {}.  
a demonstration of a person practicing {}.  
a photo of the person practicing {}.  
a video of the person practicing {}.  
a example of the person practicing {}.  
a demonstration of the person practicing {}.



## RESISC 45

---

### CuPL hand-written

Describe a satellite photo of a(n) {}  
Describe a(n) {} as it would appear in an aerial image  
How can you identify a(n) {} in an aerial photo?  
Describe the satellite photo of a(n) {}  
Describe an aerial photo of a(n) {}

---

### Standard hand-written

satellite imagery of {}.  
aerial imagery of {}.  
satellite photo of {}.  
aerial photo of {}.  
satellite view of {}.  
aerial view of {}.  
satellite imagery of a {}.  
aerial imagery of a {}.  
satellite photo of a {}.  
aerial photo of a {}.  
satellite view of a {}.  
aerial view of a {}.  
satellite imagery of the {}.  
aerial imagery of the {}.  
satellite photo of the {}.  
aerial photo of the {}.  
satellite view of the {}.  
aerial view of the {}.

## Birdsnap

---

### CuPL hand-written

Describe what the bird {} looks like:  
Describe the bird {}:  
What are the identifying characteristics of the bird {}?

---

### Standard hand-written

a photo of a {}, a type of bird.

## Flowers 102

---

### CuPL hand-written

Describe how to identify a(n) {}, a type of flower  
What does a(n) {} flower look like?

---

### Standard hand-written

a photo of a {}, a type of flower.'

## CIFAR-10

CuPL hand-written	Standard hand-written
Describe what a(n) {} looks like	a photo of a {}.
Describe a(n) {}:	a blurry photo of a {}.
What are the identifying characteristics of a(n) {}?	a black and white photo of a {}.
	a low contrast photo of a {}.
	a high contrast photo of a {}.
	a bad photo of a {}.
	a good photo of a {}.
	a photo of a small {}.
	a photo of a big {}.
	a photo of the {}.
	a blurry photo of the {}.
	a black and white photo of the {}.
	a low contrast photo of the {}.
	a high contrast photo of the {}.
	a bad photo of the {}.
	a good photo of the {}.
	a photo of the small {}.
	a photo of the big {}.

## CIFAR-100

CuPL hand-written	Standard hand-written
Describe a photo of a(n) {}:	a photo of a {}.
What are the identifying characteristics of a(n) {}?	a blurry photo of a {}.
Describe what a(n) {} looks like:	a black and white photo of a {}.
Describe a(n) {}:	a low contrast photo of a {}.
	a high contrast photo of a {}.
	a bad photo of a {}.
	a good photo of a {}.
	a photo of a small {}.
	a photo of a big {}.
	a photo of the {}.
	a blurry photo of the {}.
	a black and white photo of the {}.
	a low contrast photo of the {}.
	a high contrast photo of the {}.
	a bad photo of the {}.
	a good photo of the {}.
	a photo of the small {}.
	a photo of the big {}.

## B. CuPL Base Prompts

The three general sentences used in the base prompt setting are:

Describe what a/the \_\_\_ looks like:  
 Describe a/the \_\_\_ :  
 What are the identifying characteristics of a/the \_\_\_ ?

Here we specify the type filled in for each of the examined datasets, as well as the article used for that dataset (‘a(n)’ or ‘the’):

Dataset	Base LLM-prompt type specification
ImageNet	a(n) {}
DTD	the texture {}
StanfordCars	the car {}
SUN397	a(n) {}
Food 101	the food {}
FGVC Aircraft	the aircraft {}
Oxford Pets	a pet {}
Caltech101	a(n) {}
CIFAR-10	a(n) {}
CIFAR-100	a(n) {}
Flowers 102	the flower {}
Kinetics-700	the action of {}
UCF101	the action of {}
RESISC45	a satellite photo of {}
Birdsnap	the bird {}

## C. Evaluation Metric

ImageNet	DTD	Stanford Cars	SUN397	Food101	FGVC Aircraft	Oxford Pets	Caltech101	Flowers 102	UCF101	Kinetics-700	RESISC45	CIFAR-10	CIFAR-100	Birdsnap
Acc.	Acc.	Acc.	Acc.	Acc.	Mean per class	Mean per class	Mean per class	Mean per class	Acc.	Mean (top1, top5)	Acc.	Acc.	Acc.	Acc.

## D. Comparison with Definitions and Wikipedia Descriptions

**WordNet Definitions and Wikipedia Descriptions.** We also consider two additional methods of obtaining descriptive sentences for each ImageNet category, other than using an LLM. Firstly, we compare CuPL (full) image-prompts with image-prompts generated using definitions of each ImageNet category. Because each ImageNet category is derived from the WordNet database [34], we can use the WordNet definition of each word.

We preprocess these definitions so they are of the form “A(n) {} is a ...” as not all WordNet definitions contain the name of the word itself. We also add a period to the end of each definition, as we find this increases performance. As shown in Table 3, ImageNet Top-1 accuracy with WordNet definition prompts is below that of CuPL or standard prompts. In addition to lower accuracy, this method uses significantly more hand-constructed sentences as it requires 1000 unique hand-written definitions compared to 175 unique hand-written image-prompt templates for the standard method and 45 unique hand-written LLM-prompt templates for CuPL (full).

Secondly, we compare against prompts generated from Wikipedia articles corresponding to each ImageNet category, as collected in [7]. Note that the Wikipedia article does not always exactly match the natural language name of the class used

Standard	CuPL	WordNet	Wiki
75.54	76.69	73.44	68.20

Table 3. **ImageNet Top-1 accuracy for different methods of generating image-prompts.**

by [42]. Additionally, 80 categories map to more than one Wikipedia article (e.g. the category associated with the natural language word “patio” is mapped to the articles for “patio” and “terrace”). In this case, we select the first associated article. We preprocess these by removing the first line (the name of the article), and then extracting the first sentence, including the final period. We find both of these preprocessing steps lead to increase in accuracy. We also truncate this sentence to the maximum allowed input length of CLIP. For the 24 ImageNet categories that do not have an associated Wikipedia page, we use the name of the category as the image-prompts. As shown in Table 3, we find Wikipedia to be less effective than standard prompts, CuPL prompts, or WordNet definitions.

### E. Ensembling with Standard Prompts.

We also consider using LLM generated prompts from CuPL (full) in addition to hand-written prompts. We do this by averaging together all the text embeddings of the CuPL prompts and hand-written prompts. As shown in Table 4, we find that for some datasets, ensembling both types of prompts outperforms CuPL prompts on their own, while for others CuPL prompts perform better. For all datasets, this ensemble performs better than standard prompts alone. However, this ensembling method requires all the hand-written effort and domain knowledge of the standard approach.

	ImageNet	DTD	Stanford Cars	SUN397	Food101	FGVC Aircraft	Oxford Pets	Caltech101	Flowers 102	UCF101	Kinetics-700	RESISC45	CIFAR-10	CIFAR-100	Birdsnap	mean
Ensemble	76.51	61.60	77.66	73.51	93.42	36.47	93.71	93.87	79.73	78.16	61.50	73.03	95.88	79.33	51.09	75.03
$\Delta$ std	+0.97	+6.40	+0.13	+4.20	+0.34	+3.59	+0.38	+0.63	+1.20	+0.71	+1.43	+1.93	+0.29	+1.07	+0.66	
$\Delta$ CuPL	-0.18	-0.1	+0.03	+0.20	+0.06	+0.36	-0.10	+0.42	+0.06	+0.20	+0.87	+1.34	+0.04	+0.76	-0.02	

Table 4. **Performance of the ensemble of CuPL (full) and the standard, hand-written prompts in CLIP [42].** This ensemble outperforms the standard hand-written prompts for all examined datasets (difference shown with  $\Delta$  std), and outperforms CuPL (full) for 11 datasets (difference shown with  $\Delta$  CuPL)

### F. Additional Qualitative Results

In addition to Figure 5, we present additional qualitative results (Figure 7), demonstrating the difference in importance between regions when using CuPL prompts versus baseline prompts. As in Figure 5, we find that CuPL prompts allow the model to place more importance on semantically relevant regions of the image when distinguishing between two visually similar categories.

### G. Open-Source LLMs and ChatGPT

While GPT-3 [5] demonstrates higher performance on a number of tasks compared to smaller open-source models, open-source models are sometime more accessible. We therefore show improvement using GPT-J-6B [55], a small open-source model available on HuggingFace [58]. We find that we are able to surpass human written prompts with prompts generated by this model as shown in Table 5, though we still fall short of those generated by GPT-3. Additionally, we employ a number of strategies to increase the accuracy of the lower quality GPT-J-6B generations. First, we generate at a lower

	ImageNet Acc
standard	75.54
CuPL (GPT-J-6B)	75.62
CuPL (ChatGPT)	76.09
CuPL (GPT-3)	76.69

Table 5. **CuPL with an open-source model and ChatGPT.** CuPL is able to improve over hand-written baselines even for smaller models.









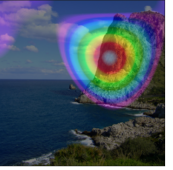





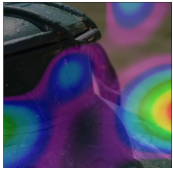



Original Image	Region Importance with CuPL Prompts		Region Importance with Baseline Prompts		Example Image of Distractor Class
	(A)	(B)	(C)	(D)	
<b>GT Label:</b> Eastern Diamondback Rattlesnake	<b>Prompt eastern diamondback:</b> "A eastern diamondback rattlesnake is a snake with a diamond-shaped pattern on its <b>back</b> ."	<b>Prompt sidewinder rattlesnake:</b> "A sidewinder rattlesnake is a type of rattlesnake that has a unique method of locomotion."	<b>Prompt eastern diamondback:</b> A photo of a eastern diamondback rattlesnake.	<b>Prompt sidewinder rattlesnake:</b> A photo of a sidewinder rattlesnake.	<b>Example:</b> Sidewinder Rattlesnake
					
	Prediction: Eastern Diamondback ✓		Prediction: Sidewinder ✗		
<b>GT Label:</b> Promontory	<b>Prompt promontory:</b> "A promontory is a landform that protrudes into a <b>body of water</b> ."	<b>Prompt cliff:</b> "A cliff is a high, steep rock face or slope."	<b>Prompt promontory:</b> A photo of a promontory.	<b>Prompt cliff:</b> A photo of a cliff.	<b>Example:</b> Cliff
					
	Prediction: Promontory ✓		Prediction: Cliff ✗		
<b>GT Label:</b> Milk Can	<b>Prompt milk can:</b> "A milk can looks like a large metal can with a <b>lid</b> ."	<b>Prompt water jug:</b> "A water jug is typically a plastic or glass container with a spout and a handle."	<b>Prompt milk can:</b> A photo of a milk can.	<b>Prompt water jug:</b> A photo of a water jug.	<b>Example:</b> Water Jug
					
	Prediction: Milk Can ✓		Prediction: Water Jug ✗		

Figure 7. More examples as in Figure 5

temperature (0.3) to prevent irrelevant or nonsensical generations, which we find occur more frequently in smaller models. Additionally, we generate 5 times more Image-prompts per LLM-prompt than we do when using GPT-3 [5]. We also add punctuation to the end of all LLM-prompts to encourage the LLM to begin new sentences. Finally, we filter out any Image-prompts which do not contain the name of the ImageNet category they are meant to describe, as well as remove a number of unicode characters from the generations (e.g. ‘u2019’). We present these findings as a way to make CuPL a more accessible options until large high performance models become available to the public.

Additionally, we present CuPL with prompts generated with ChatGPT, a popular chat model with a similar design to GPT-3 [5]. This model is less expensive to run, 10 times cheaper than Davinci-002. As shown in Table 5, it outperforms hand-written prompts by 0.55%. We do modify the original LLM prompts (i.e. “What does a {} look like?”), by formatting it as follows: “Based on what you know as an AI Language model, answer the following question in one sentence:” + LLM prompt + “. Do not mention that you are a language model.” This is because we found that ChatGPT often mentions that it is an AI language model before answering the question or sometimes refused to answer the question. Even with this prompt addition, there were a number of generations mentioning that it is an AI model or refusing to answer, so we filtered out any

responses that contained the words “AI” or “sorry”. We use all the same hyperparameters as the original CuPL, except with max context length increased to 100. We present this as an inexpensive alternative to CuPL full as we were able to generate all ImageNet prompts with ChatGPT for less than \$5 USD.

## H. Single Sentence Baseline

	ImageNet Acc
a photo of a {}	73.46
CuPL (1 hand-written)	75.71
CuPL (1 generated)	74.24

Table 6. **Single sentence baselines.** Comparison of a single hand-written Image-prompt template with a single hand-written LLM-prompt as well as a single CuPL generated Image-prompt.

like’. For this experiment, we generate at a temperature of 0.3 as we find that higher temperatures are only helpful when we are able to ensemble many diverse prompts, not when we are limited to one. We find that CuPL outperforms a single hand-written template under both of these settings, as shown in Table 6.

One of the primary benefits of CuPL is that it decreases the amount of necessary hand-engineering. However, this could also be done by decreasing the number of total hand-written templates used, which comes at a loss in performance. We present this as a ‘low effort’ baseline, where we use only the hand constructed template of ‘a photo of a {}’. We compare this with two CuPL baselines. The first is the baseline in which we also only construct one hand written template: ‘Describe what a {} looks like’. We then use this to generate 10 Image-prompts. The second baseline is a single CuPL generated sentence, generated with the prompt ‘Describe what a {} looks

## I. Robustness

In addition to the previously mentioned benefits of open-vocabulary models, one of the important advances made by CLIP [42] is an increased robustness on out-of-distribution data. Fine-tuning has been shown to degrade performance on out-of-distribution tasks [59], however zero-shot CLIP is robust to these distribution shifts. We show improvement on two common distribution shifts in Table 7, demonstrating that CuPL maintains the robustness of CLIP.

	ImageNet [13]	ImageNet-V2 [44]	ImageNet-Sketch [56]
Standard	75.54	69.86	59.60
CuPL	76.69	70.85	60.05

Table 7. **Robustness of CuPL.** CuPL accuracy on two common ImageNet variants, using CuPL ImageNet Image-prompts. CuPL improves performance on both of these variants, demonstrating that CuPL improves accuracy on in-distribution tasks, while maintaining robustness to distribution shifts.

## J. Tagging with ImageNet Parts details

As demonstrated in Figure 6, we find a strong correlation between the number of mentions of parts of ImageNet classes (‘head’, ‘body’, ‘legs’, etc) in the CuPL prompts and the importance of those parts visually. In this section we detail how we tag these parts in the text of CuPL prompts. The reason that this is difficult is because we cannot simply count the instances of the word ‘head’ to understand the importance of the head in identifying a class. This is because the head contains many other parts, and these may be mentioned instead. For example the face may be important for identification. This is within the ‘head’ segmentation of the PartImageNet [18] category, so should be counted as a mention of the head. To account for this, we take advantage of WordNet [34] which annotates the parts of various objects. Therefore we can use WordNet to obtain a list of parts for each category in PartImageNet. It is also important to note that while ‘face’ is a part of ‘head’, there are also subparts of face. We consider 2 layers of depth to gather parts for a given category.

## K. CuPL Improvement Analysis

### K.1. Per Class Accuracy CuPL versus Baseline

In addition to total accuracy gains, we present the per class accuracy shift between the image-prompts used in [42] and CuPL, shown in Figure 8. As demonstrated, the accuracy gains seen in CuPL are not distributed uniformly through the ImageNet classes, with some classes seeing ~40 percentage point accuracy gains, and others seeing ~40 percentage point

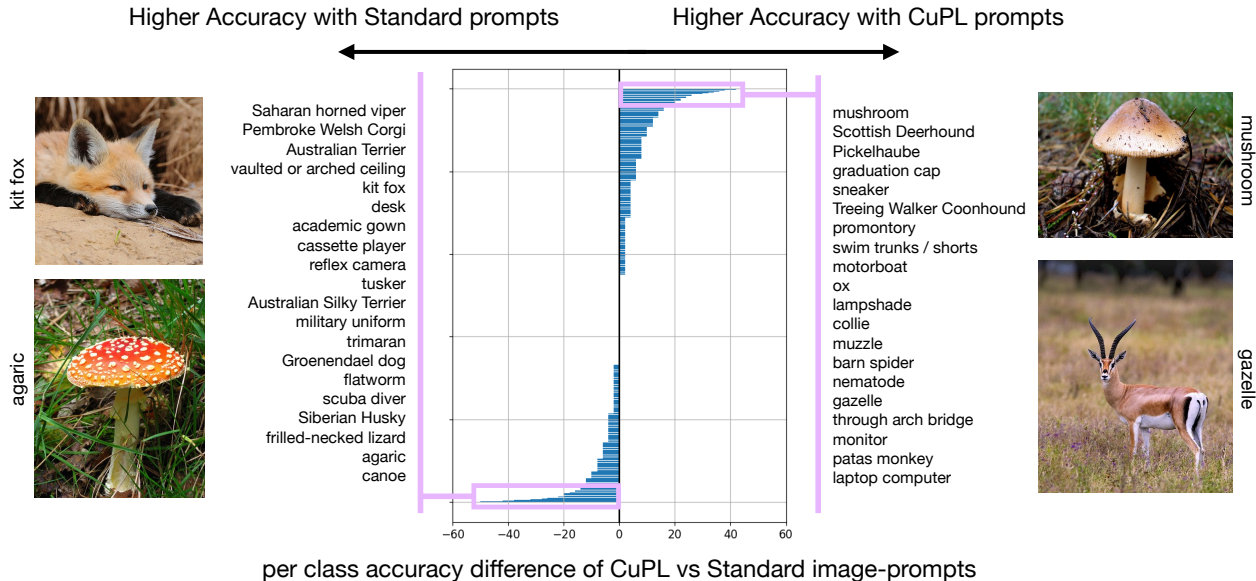


Figure 8. **Distribution of ImageNet per class accuracy difference of CuPL image-prompts versus standard image-prompts.** As shown, the accuracy gains of CuPL are not uniform across all classes. Rather, we see large gains for some classes, and losses for others. In addition, we list the classes which see the largest accuracy gains when switching to CuPL prompts (with “mushroom” having the largest gain), and the 20 classes with the largest accuracy losses (with “canoe” having the largest loss.)

accuracy losses when compared against class accuracy with standard prompts. In other words, while CuPL sees a higher accuracy overall when compared to the standard method, the images which are correctly predicted by the standard prompts are not a subset of the images which are correctly predicted by CuPL. In fact CuPL sees just over a 1 percentage point gain when compared to standard prompts, but differs in its predictions from the standard method for 11.50% of predictions (with CuPL correct for 4.48% of these, standard correct for 3.32%, and neither correct for 3.70%).

Figure 8 also shows the classes with the 20 greatest accuracy gains and losses when comparing class accuracy with standard image-prompts [42] versus with CuPL image-prompts. Interestingly, for many of the classes which see a large accuracy gain, we see a corresponding class with a large accuracy loss that is either similar to the initial class or likely to co-occur with it (e.g. agaric/mushroom, academic gown/graduation cap, military uniform/Pickelhaube, desk/monitor).

## K.2. visual similarity analysis

In Figure 8, we provide initial analysis on the categories where the CuPL algorithm improves the most over standard hand-written prompts. We find that the improvement is not uniformly distributed, but rather some classes see a large improvement, while others see a decrease in per class accuracy. Interestingly, there are often two similar categories where one sees a large increase in accuracy and the other sees a decrease. For example, the ‘mushroom’ class has an approximately 40pp increase, while ‘agaric’ (a subclass of mushroom) is one of the classes with the largest drop in accuracy.

In order to better understand this phenomenon, we examine the change in accuracy between CuPL and the standard method of prompting in the image embedding space. Thus we are able to visualize the close relationship between categories like ‘agaric’ and ‘mushroom’. In order to be able to visualize the high dimensional CLIP image embedding in two dimensions, we utilize the t-distributed stochastic neighbor embedding algorithm [54]. Figure 9 visualizes image features (reduced into two dimensions) in relation to CuPL improvement.

As was suggested by Figure 8, we see in Figure 9 that when there is a class that has a large increase in accuracy with CuPL prompts (‘mushroom’, ‘graduation cap’, ‘monitor’) there is often a decrease in class accuracy for a visually related class. This means that when choosing between two similar or co-occurring classes, CuPL has a different distribution of classification than the standard method (e.g. the standard method prefers ‘canoe’ over ‘paddle’ much more strongly than CuPL). This suggests that the overall accuracy improvement of CuPL over the standard method may come (at least in part) from better distinguishing between two visually similar classes. While it may be over-correcting from the mistakes of the standard method (as demonstrated by the drop in accuracy in one of the two similar classes), the CuPL predictions appear to be overall more accurate, as demonstrated by the overall higher accuracy.

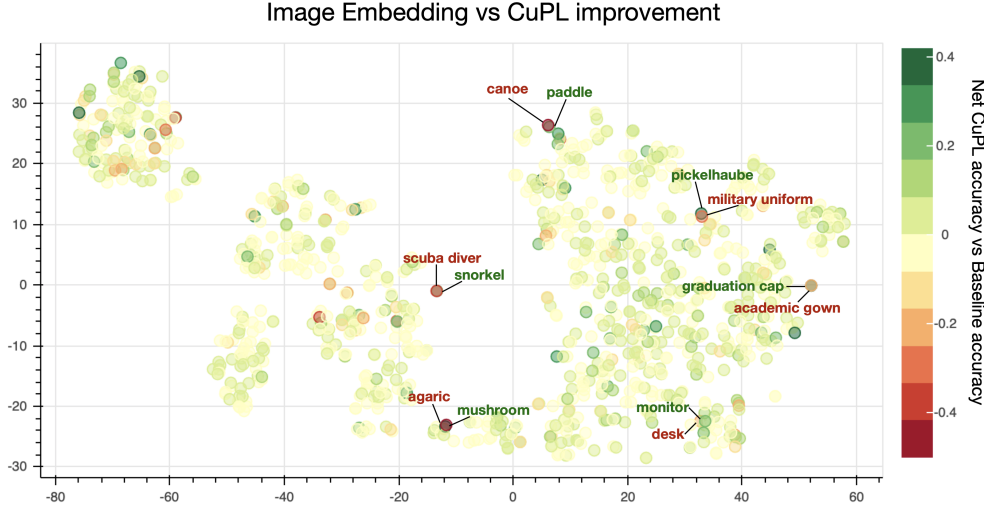


Figure 9. **Visualizing of image embedding of ImageNet classes compared to CuPL improvement on that class.** Each point on this figure represents the average image embedding of an ImageNet class, which has been reduced to two dimensions using t-sne [54]. We see that when there is a class with a large improvement compared to the baseline, it is often visually similar to a class which has a decrease in accuracy. This suggests that CuPL’s improved accuracy may be due in part to an increased ability to distinguish similar classes compared to the baseline.

### K.3. Co-occurrences between objects

Many of the frequently confused pairs in Figure 9 are objects that are likely to occur in an image (i.e. ‘canoe’-‘paddle’ or ‘graduation cap’-‘academic gown’ or ‘monitor’-‘desk’). One potential benefit of CuPL captions is that they are able to capture co-occurrences as well. For example, one CuPL prompt for the ‘canoe’ class is *A canoe is typically a narrow boat with pointed ends that is propelled with a paddle.* This caption contains the word ‘paddle’ which is frequently confused with ‘canoe’ and likely to be present in images, even where the correct label is ‘canoe’. We therefore investigate the effectiveness of CuPL captions on images which contain more than one ImageNet object.

We attain this by using the ImageNet-ReaL dataset [3] which relabels ImageNet images with all applicable labels, so an image with both a ‘canoe’ and a ‘paddle’ would have both labels. We then tag images as having multiple ImageNet objects or only one ImageNet object based on the ReaL dataset. Finally, we compute ImageNet accuracy across each of these two sets (using standard ImageNet labels). Results are given in Table K.3.

	One class present (85.1% of ims)	Multiple classes present (14.9% of ims)
Standard	79.89	51.59
CuPL	80.79	53.58

Table 8. **Standard versus CuPL accuracy based on number of ImageNet classes present in image.** We use the ImageNet-ReaL dataset [3] to find images which have more than one applicable ImageNet label. We then present the accuracy for standard prompts and CuPL prompts using standard ImageNet labels, split by images which contain only one possible ImageNet class and images which may contain multiple classes.

### L. Error Analysis

In Figure 10, we present an error analysis of our model using two different metrics. The first is an analysis between the model prediction and the correct class using the visual similarity of these labels. To capture this, we first attain an average visual embedding of each class by taking the mean of each image in that class and then normalizing that mean. Then for each class we rank how similar each of the other 999 classes are by the distance between these embeddings. If the models makes an incorrect prediction, but it predicts the class with the closest embedding to the correct label, then we refer to this as an *image offset* of 1.

Additionally, we examine the prediction errors in terms of the linguistic similarity of the labels. We do this with the WordNet [34] similarity of two labels. For example, if the label of the prediction and the ground-truth label share the same



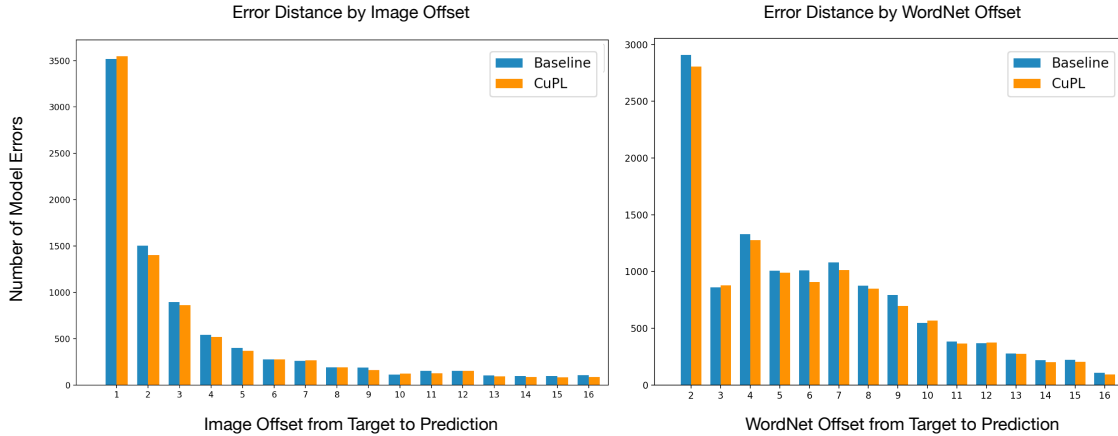


Figure 10. **Error analysis of CuPL and baseline comparing models errors visually and linguistically to ground truth labels.** We compare the errors made by CuPL to the errors made by the baseline methods. We find that the errors made by CuPL are more likely to be the most visually similar class to the ground truth label when compared to the errors made by the baseline. However, the errors made by the baseline are more likely to be the most linguistically similar class according to the WordNet [34] heirarchy. This suggests that visual information is being extracted from the CuPL descriptions to make visually consistent predictions.

parent in the WordNet tree, that is a *WordNet offset* of 2.

While slight, there is a difference in the errors made by CuPL compared to the errors made by the baseline as shown in Figure 10. CuPL is more likely to have an error that has an image offset of 1 than the baseline. However, the baseline is more likely to have an error that has a WordNet offset of 1 than CuPL. This implies that CuPL may be taking advantage of the visually descriptive language of the captions, as even when the model makes errors, they tend to favor categories that are visually similar to the ground truth. However, the baseline method does not have visual descriptions in its Image-prompts which may lead to its errors aligning more linguistically with the ground-truth.

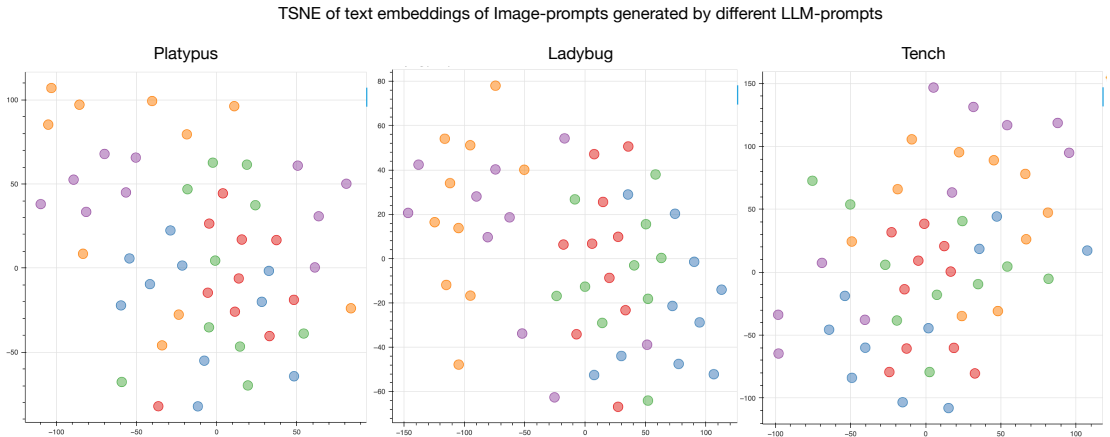


Figure 11. **Visualization of embeddings of Image-prompts generated with various LLM-prompts.** Each point represents the image embedding of one Image-prompts for the stated category with has been reduced to two dimensions using t-sne [54]. Image-prompts with the same color are generated by the same LLM-prompts.

## M. Image-prompt Distribution

When generating Image-prompts, we use an ensemble of prompts generated by different LLM-prompts (e.g. ‘What does a {} look like?’). As we find that the diversity of Image-prompts is correlated with accuracy, it is valuable to understand how different LLM-prompts affect the diversity of Image-prompts. To accomplish this, we visualize the text embedding of Image-prompts for a selection of classes using t-sne [54] dimension reduction. We then color Image-prompts that were

generated by the same LLM-prompt. As shown in Figure 11, we find that while there is a slight clustering of Image-prompts by LLM-prompts, there is also a large amount of overlap.

## N. Temperature Analysis

**Diversity of Prompts.** We also examine the impact of the diversity of image-prompts on ImageNet accuracy. We adjust this parameter by changing the temperature of the GPT-3 model. This value changes the likelihood of selecting lower probability tokens and makes sentences more diverse from each other. As demonstrated in Figure 12, more diverse prompts lead to higher ImageNet accuracy. Note these comparisons are done with a single LLM-prompt to save computational cost.

We provide several further analyses of the effect of temperature in prompt generation. First, we visualize the distribution of Image-prompts that have been generated with a variety of different temperatures. We do this by selecting prompts for 3 different ImageNet classes at 3 different temperatures. We then perform dimensionality reduction on the text embeddings of these prompts in order to visualize their distribution. As shown in Figure 13, Image-prompts generated with a temperature of 0.1 are clustered in a few different locations. Image-prompts generated with a temperature of 0.5 are more widely, and Image-prompts generated with a temperature of 0.9 have a similar, but even wider distribution.

Additionally in Section O, we give all generated prompts for the ImageNet class ‘Tench’ at three different temperatures. At the lowest temperature, the generated Image-prompts are nearly identical when generated with the same LLM-prompts. As the temperature increases, so does the difference in the generated prompts.

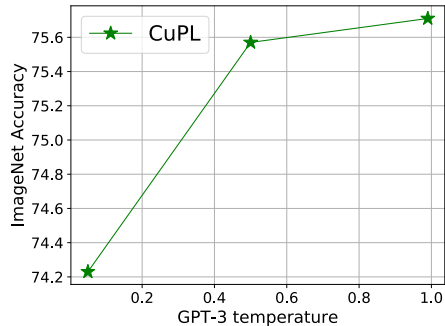


Figure 12. **Effect of LLM temperature.** More prompt diversity leads to higher performance.

As shown in Figure 13, Image-prompts generated with a temperature of 0.1 are clustered in a few different locations. Image-prompts generated with a temperature of 0.5 are more widely, and Image-prompts generated with a temperature of 0.9 have a similar, but even wider distribution.

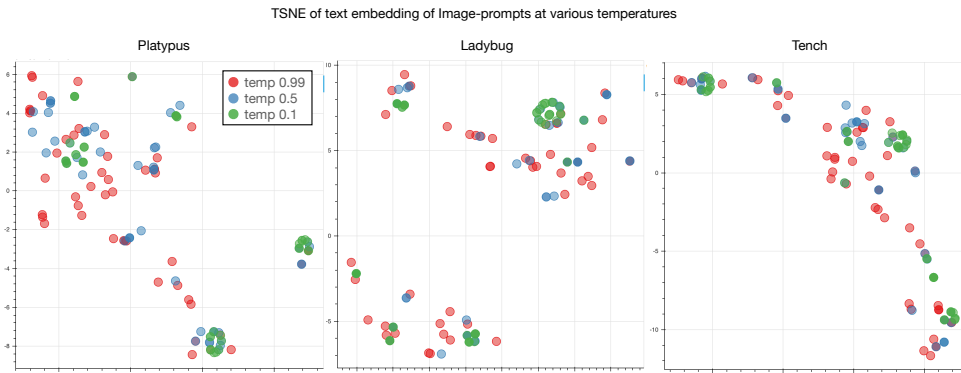


Figure 13. **Visualization of embeddings of Image-prompts generated with various temperatures.** Each point represents the image embedding of one Image-prompts for the stated category with has been reduced to two dimensions using t-sne [54]. Prompts generated with a higher temperature cover a wider distribution.

## O. Example Generated image-prompts

A selection of LLM-generated image-prompts for a subset of ImageNet categories. We give all 50 image-prompts for the first ImageNet category of “Tench” and then 10 randomly selected prompts for a number of randomly selected ImageNet categories.

### O.1. All generated image-prompts for “Tench” category

**Temperature = 0.99**

- "A tench is a freshwater fish of the carp family.",
- "A tench is a freshwater fish that is typically brown or olive in color.",
- "A tench is a fresh water fish that can grow up to 2 feet in length.",
- "A tench is a freshwater fish of the family Cyprinidae.",

"A tench is a freshwater fish of the carp family.",  
 "A tench is a freshwater fish with a dark green back and light-colored sides.",  
 "Tench are a freshwater fish found in Europe.",  
 "A tench is a small freshwater fish in the carp family.",  
 "A tench is a heavyset freshwater fish with a mottled brown body and a small, flat head.",  
 "A tench is a freshwater fish that looks similar to a carp.",  
 "A tench is a freshwater fish in the carp family.",  
 "A tench is a freshwater fish of the Cyprinidae family.",  
 "The tench is a freshwater fish of the Cyprinidae family.",  
 "The tench is a freshwater fish in the family Cyprinidae.",  
 "The easiest way to identify a tench is by its herringbone-patterned scales.",  
 "A tench is a freshwater fish of the carp family.",  
 "Tench are a freshwater fish found in Europe.",  
 "Tench have a large, slimy body with scales that have a green hue.",  
 "The tench is a freshwater fish belonging to the carp family.",  
 "A tench is a freshwater fish of the Cynoglossidae family.",  
 "A tench is a freshwater fish in the carp family.",  
 "Tensch are freshwater fish with Olive Green backs, shading to Yellowish on the sides.",  
 "A tench looks like a green freshwater fish with a brownish hue.",  
 "A tench looks like a freshwater fish with a dark olive-green back, fading to yellowish-brown on the sides.",  
 "A tench usually has olive-green skin with dark spots, and a orange-yellow underbelly.",  
 "Tench are a freshwater fish that can grow up to 70cm long! They have olive-brown skin with dark spots, and their meat is white and firm.",  
 "A tench is a freshwater fish with a sturdy body and a greenish-brown coloration.",  
 "A tench is a freshwater fish that can grow up to about two feet long.",  
 "A tench is a freshwater fish in the carp family.",  
 "A tench is a large, freshwater fish with a thick body and large head.",  
 "The image is of a tench fish swimming in water.",  
 "The image is of a tench fish swimming in a pond.",  
 "The tench is a freshwater fish native to Europe.",  
 "This image shows a large, dark green tench swimming in a pond.",  
 "An image of a tench from the internet would likely show a dark green fish with a lighter underside.",  
 "The image is of a tench fish.",  
 "The image is of a tench fish on a white background.",  
 "A tench is a freshwater fish of the Cyprinidae family.",  
 "The image is of a tench swimming in a murky pond.",  
 "In the image, a tench swims in a pond with lily pads.",  
 " A tench in a river.",  
 "A tench (Tinca tinca) is a freshwater fish in the carp family that is found throughout Europe.",  
 " Tench (Tinca tinca), a member of the carp family (Cyprinidae), native to Eurasia.",  
 " A tench, a freshwater fish in the family Cyprinidae.",  
 " The tench (Tinca tinca) is a freshwater fish of the cyprinid family found throughout Eurasia.",  
 " A tench in a Finnish lake.",  
 "A tench (Tinca tinca) is a freshwater fish belonging to the carp family (Cyprinidae).",  
 "A tench in a fishpond.",  
 " The common tench is a freshwater fish of the cyprinid family found throughout Eurasia.",  
 "Tench (Tinca tinca) in a pond."

### **Temperature = 0.5**

"A tench is a freshwater fish that is typically greenish-brown in color with a brassy sheen.",  
 "A tench is a freshwater fish that typically has a dark green back, light brown sides, and a white belly.",  
 "A tench is a freshwater fish that can grow up to two feet long.",

"Tench are a freshwater fish found in Europe.",  
 "A tench is a freshwater fish that is typically olive green in color with dark spots.",  
 "A tench is a freshwater fish that is typically olive green in color with a brownish tint.",  
 "A tench is a freshwater fish of the Cyprinidae family.",  
 "Tench are a freshwater fish found in Europe.",  
 "A tench is a freshwater fish that has a dark green back, light brown sides, and a white belly.",  
 "A tench is a freshwater fish that is typically olive-green in color with a brownish dorsal fin.",  
 "A tench is a freshwater fish of the cyprinid family.",  
 "A tench is a freshwater fish of the carp family.",  
 "A tench is a freshwater fish that is typically olive green in color with a brownish back.",  
 "The tench is a freshwater fish of the carp family Cyprinidae.",  
 "A tench is a freshwater fish that is typically greenish-brown in color.",  
 "A tench is a freshwater fish of the carp family.",  
 "A tench is a freshwater fish of the carp family.",  
 "Tench have olive green backs and flanks, with yellowish bellies.",  
 "A tench is a freshwater fish of the carp family.",  
 "A tench is a freshwater fish of the cyprinid family.",  
 "A tench is a freshwater fish that can grow up to 30 inches long.",  
 "A tench is a freshwater fish that is typically greenish-brown in color.",  
 "A tench is a freshwater fish that can grow up to about two feet long.",  
 "A tench is a freshwater fish with a brownish-green back and sides, and a yellowish-brown belly.",  
 "A tench is a freshwater fish that can grow up to two feet long.",  
 "A tench is a freshwater fish that looks similar to a carp.",  
 "A tench is a freshwater fish that is typically greenish-brown in color.",  
 "A tench is a freshwater fish that is part of the carp family.",  
 "A tench is a freshwater fish of the cyprinid family.",  
 "A tench is a freshwater fish that can grow up to two feet long.",  
 "The image is of a tench fish swimming in water.",  
 "The image is of a tench fish swimming in a pond.",  
 "The image is of a tench fish swimming in a pond.",  
 "The image is of a tench fish swimming in a pond.",  
 "The image is of a tench fish swimming in a pond.",  
 "In the image, a tench is swimming in a pond with lily pads.",  
 "The image is of a tench fish swimming in a pond.",  
 "The image is of a tench fish swimming in a pond.",  
 "The image is of a tench fish swimming in a pond.",  
 "The image is of a tench fish swimming in a pond.",  
 "A tench fish, native to Europe, characterized by its greenish-brown color and spots.",  
 " A tench (*Tinca tinca*) in a pond.",  
 "A tench (*Tinca tinca*) is a freshwater fish belonging to the carp family (Cyprinidae).",  
 "A tench (*Tinca tinca*) is a freshwater fish in the carp family.",  
 "A tench (*Tinca tinca*) is a freshwater fish in the carp family (Cyprinidae).",  
 " A tench in a river.",  
 "A tench (*Tinca tinca*) is a freshwater fish in the carp family (Cyprinidae).",  
 "A tench (*Tinca tinca*) in a pond.",  
 " A tench (*Tinca tinca*) in a garden pond.",  
 " A tench in a river."

### **Temperature = 0.1**

"A tench is a freshwater fish that is typically olive green in color with dark spots.",  
 "A tench is a freshwater fish that can grow up to 30 inches long.",  
 "A tench is a freshwater fish that can grow to a length of over two feet.",  
 "A tench is a freshwater fish that can grow up to two feet long.",  
 "A tench is a freshwater fish that can grow up to two feet long.",  
 "A tench is a freshwater fish that is typically olive green in color with dark spots.",

"A tench is a freshwater fish that can grow up to two feet long.",  
 "A tench is a freshwater fish that can grow up to two feet long.",  
 "A tench is a freshwater fish that can grow up to two feet long.",  
 "A tench is a freshwater fish that is typically olive green in color with dark spots.",  
 "A tench is a freshwater fish of the carp family.",  
 "A tench is a freshwater fish of the carp family.",  
 "A tench is a freshwater fish of the carp family.",  
 "A tench is a freshwater fish of the carp family.",  
 "A tench is a freshwater fish of the carp family.",  
 "Tench have a dark green back, light olive sides, and a yellowish belly.",  
 "A tench is a freshwater fish of the carp family.",  
 "A tench is a freshwater fish of the carp family.",  
 "A tench is a freshwater fish of the carp family.",  
 "A tench is a freshwater fish that can grow up to two feet long.",  
 "A tench is a freshwater fish that can grow up to two feet long.",  
 "A tench is a freshwater fish that can grow up to two feet long.",  
 "A tench is a freshwater fish that can grow up to two feet long.",  
 "A tench is a freshwater fish that is typically olive green in color with a brownish dorsal side.",  
 "A tench is a freshwater fish that can grow up to two feet long.",  
 "A tench is a freshwater fish that can grow up to two feet long.",  
 "A tench is a freshwater fish that is typically olive green in color with a brownish dorsal side.",  
 "A tench is a freshwater fish that can grow up to two feet long.",  
 "A tench is a freshwater fish that can grow up to two feet long.",  
 "The image is of a tench fish swimming in a pond.",  
 "The image is of a tench fish swimming in a pond.",  
 "The image is of a tench fish swimming in a pond.",  
 "The image is of a tench fish swimming in a pond.",  
 "The image is of a tench fish swimming in a pond.",  
 "The image is of a tench fish swimming in a pond.",  
 "The image is of a tench fish swimming in a pond.",  
 "The image is of a tench fish swimming in a pond.",  
 "The image is of a tench fish swimming in a pond.",  
 "A tench (Tinca tinca) is a freshwater fish in the carp family.",  
 "A tench (Tinca tinca) is a freshwater fish of the carp family (Cyprinidae).",  
 "A tench (Tinca tinca) is a freshwater fish in the carp family.",  
 "A tench (Tinca tinca) is a freshwater fish in the carp family.",  
 "A tench (Tinca tinca) is a freshwater fish in the carp family.",  
 "A tench (Tinca tinca) is a freshwater fish in the carp family.",  
 "A tench (Tinca tinca) is a freshwater fish in the carp family.",  
 "A tench (Tinca tinca) is a freshwater fish in the carp family.",  
 "A tench (Tinca tinca) in a garden pond.",  
 "A tench (Tinca tinca) in a garden pond.",  
 "A tench (Tinca tinca) is a freshwater fish of the carp family (Cyprinidae)."

**O.2. Sample generated image-prompts for randomly selected ImageNet categories**

"bubble":  
 "A bubble is a sustained period of inflated asset prices.",  
 "A bubble looks like a sphere of air.",  
 "A bubble is often characterized by rapidly increasing prices in an asset or security, followed by a sharp decrease in prices.",  
 "A bubble looks like a small, round, thin film of soap filled with air.",  
 "A bubble looks like a round sphere of soap film.",  
 "\"A bubble being blown in the park.",  
 "A close-up of a soap bubble with a thin film of water in between two layers of air.",

"A bubble is a spherical shape made up of a thin film of soap water.",  
"A bubble is a circle of air surrounded by water.",  
"A bubble looks like a round, thin layer of soap surrounding a pocket of air."

"kit fox":

"A kit fox is a small species of fox, about the size of a domestic cat.",  
"You can identify a kit fox by its small size, its big ears, and its long, bushy tail.",  
" A kit fox laying on the ground in a desert habitat.",  
"A kit fox is a small fox found in North America.",  
"A kit fox is a small fox with a sleek coat of fur.",  
"A kit fox is a small fox with large ears, a long, black-tipped tail, and pale fur.",  
"A kit fox has a reddish coat, with white patches on its chest and throat.",  
"This kit fox has a reddish coat and large ears.",  
"A kit fox looks like a small fox with a pointed nose, large ears, and a long, bushy tail.",  
"A kit fox is a small species of fox."

"toy terrier":

"A toy terrier looks like a very small version of a terrier.",  
"The image is of a toy terrier that is mostly white with brown spots.",  
"You can identify a toy terrier by looking for a compact, short-legged dog with a short muzzle.",  
" Cute little guy.",  
"Toy terriers are miniature versions of terriers, such as the Jack Russell Terrier.",  
"A toy terrier is a small, lightweight breed of dog.",  
"The image is of a small, brown toy terrier.",  
"A toy terrier typically has a long, narrow head with pointy ears, and a small, compact body.",  
"A toy terrier is a small, short-legged dog with a long body, pointy nose, and large ears.",  
"A small, brown and white toy terrier is sitting on a beige couch, looking at the camera."

"mousetrap":

"In the image, there is a mousetrap made of wood and metal.",  
"An image of a mousetrap from the internet would most likely show a traditional wooden mousetrap with a metal spring.",  
"A mousetrap is a device made to catch and kill mice.",  
"A mousetrap is a small device that is used to catch mice.",  
"The classic mousetrap consists of a wooden base with a metal spring mounted on one end.",  
"The mousetrap is a small wooden box with a metal spring inside.",  
"The mousetrap is a simple device that has been used for centuries to catch mice.",  
"The classic mousetrap - simple, effective, and deadly.",  
"The most common way to identify a mousetrap is by its small size and rectangular shape.",  
"A mousetrap typically has a wire or wooden frame that is baited with food and springs open quickly to snap shut on the mouse when it attempts to steal the bait."

"dog sled":

"A dog sled is a vehicle on runners, typically with a thin frame and a flat bottom, that is used to convey goods or passengers over snow or ice.",  
"A dog sled is a toboggan pulled by dogs, typically over snow.",  
"A dog sled looks like a bed on runners that is pulled by dogs.",  
"A dog sled is typically a heavy frame on runners that is pulled by one or more dogs.",  
"A dog sled is traditionally a sled pulled by dogs, used for transportation, racing, or other purposes.",  
" teamA dog sled team is a group of dogs that are harnessed together to pull a sled.",  
"A dog sled looks like a small, open vehicle that is pulled by one or more dogs.",  
"Dog sledding in winter.",  
" pulled by huskiesThe image is of a dog sled pulled by huskies.",  
"A dog sled looks like a large cart that is pulled by a team of dogs."

"geyser":

"A geyser is a hot spring that periodically erupts, shooting a column of water and steam into the air.",  
"A geyser looks like a hole in the ground that sometimes spurts hot water and steam into the air.",  
"Geysers are hot springs that periodically spout water and steam into the air.",  
"The image is of a geyser erupting.",  
"A geyser is a hot spring that periodically erupts, spraying water into the air.",  
"A geyser typically looks like a cone of rocks with a small hole at the top.",  
"A geyser is a hot spring where water intermittently boils, sending a jet of hot water and steam into the air.",  
"A geyser is a hot spring that periodically shoots a stream of hot water and steam into the air.",  
"The image is of a geyser shooting water high into the air.",  
"A geyser looks like a column of water that shoots into the air and then falls back down."

"Schipperke":

"A Schipperke is a small, Belgian breed of dog.",  
"A Schipperke is a small black Belgian dog with a rat-like tail.",  
"The image is of a black and white dog with pointy ears and a long body.",  
"A Schipperke is a small, black, Belgian breed of dog.",  
"A Schipperke is a small, black, Belgian breed of dog that closely resembles a fox.",  
"A Schipperke is a small Belgian breed of dog that resembles a fox.",  
"Schipperkes have a long, black coat and a pointed muzzle.",  
"A Schipperke is a small dog breed with a fox-like appearance.",  
"It's a photo of a black and tan Schipperke dog standing in front of a brick wall.",  
"Black, small, spitz-type dog with a long, fox-like snout, large erect ears, and a long, high-set tail."

"go-kart":

"A go-kart typically looks like a small car or buggy with a small engine in the back.",  
"A go-kart is a small vehicle with four wheels, a steering wheel, and a gas pedal.",  
"A go-kart is a small vehicle with a steering wheel, pedals, and an engine.",  
"Two young girls in go-karts race down a path in a park.",  
"A go-kart is a small, lightweight vehicle with four wheels and a simple, open frame.",  
"A go-kart is a small, open-wheeled vehicle used for racing.",  
"Two kids racing go-karts on a dirt track.",  
"A go-kart is a small, racing car.",  
"A go-kart typically looks like a small, open-wheeled car.",  
"A go-kart is a small, lightweight vehicle with four wheels that is propelled by a small engine."

"black-and-white colobus":

"The Black-and-White Colobus is a type of Old World monkey, found in Africa.",  
"A black-and-white colobus has long black fur, and a white face with a black triangle around the eyes.",  
"Colobus Monkey in the TreesThis elegant colobus monkey is swinging through the trees in search of food.",  
" monkeyIn this image, a black-and-white colobus monkey is shown perched atop a tree branch.",  
"The black-and-white colobus monkey is one of the most beautiful and distinctive of all the colobus monkeys.",  
"The black-and-white colobus monkey is a species of primate in the Colobidae family.",  
"Colobus monkeys are generally black with white patches on their face, back, and sides.",  
"The black-and-white colobus is a species of Old World monkey.",  
" monkeyThe image is of a black-and-white colobus monkey sitting on a tree branch.",  
" monkeyIn the image, the black-and-white colobus monkey is sitting in a tree."

"sock":

"A sock usually has a cuff at the top, and a heel at the bottom.",  
" A black sock with a white line running down the middle.",  
"A sock is typically a garment worn on the feet and made from a soft material, such as cotton.",  
"A sock normally has a heel, toe and a cuff at the top.",  
"A sock is a small amount of money that is given to someone without them knowing.",  
"A sock is an article of clothing worn on the feet.",  
"A sock is a piece of clothing that is worn on the feet.",  
"A sock is a tubular garment that covers the foot and ankle.",  
"This image is of a blue and white striped sock.",  
"There are many ways that you can identify a sock."

"Cocker Spaniel":

"A Cocker Spaniel is a medium sized dog with long, floppy ears, and a silky coat that is usually either brown or black.",  
"The Cocker Spaniel has a long floppy ears, a silky coat, and a bushy tail.",  
"A Cocker Spaniel has a long, black muzzle and big, brown eyes.",  
"The Cocker Spaniel is a breed of dog.",  
"An image of a Cocker Spaniel from the internet shows a small brown and white dog with long floppy ears.",  
"The image is of a Cocker Spaniel with short, brown fur and long, floppy ears.",  
"Cocker spaniels have long, floppy ears and a long, silky coat.",  
"A Cocker Spaniel is a small to medium sized dog.",  
"The image is of a light brown and white Cocker Spaniel standing on a green grassy field with its head turned to the side.",  
"A Cocker Spaniel has a long, silky coat that is usually either black, brown, or golden."

"southern black widow":

"A southern black widow spider perched atop a web.",  
"Female southern black widows have a black body with a red hourglass shape on their abdomen.",  
"The southern black widow is black with a red hourglass shape on its belly.",  
"Female southern black widow spiders are black with a characteristic red hourglass-shaped mark on their ventral abdomen.",  
" spiderThe image is of a large, black spider with a red hourglass shape on its abdomen.",  
"A southern black widow is a spider that is black with a red hourglass shape on its abdomen.",  
"A southern black widow spider is a small, black spider with a red hourglass-shaped mark on its underside.",  
"There's an image on the internet of a southern black widow that's really cool.",  
"A southern black widow can be identified by its black coloration with a red hourglass shape on its abdomen.",  
"A Southern black widow is a type of spider that is black with a red hourglass shape on its belly."

"catamaran":

"A catamaran is a sailboat that has two hulls, or wide bodies, that are connected by beams.",  
"The easiest way to identify a catamaran is by its two hulls.",  
"A catamaran is a type of boat that has two parallel hulls.",  
"A catamaran is a multi-hulled vessel with two parallel hulls of equal size.",  
"My dream boat! A sleek catamaran that can zip through the waves.",  
"A catamaran is a multi-hulled vessel with two parallel hulls of equal size.",  
"The image is of a white catamaran with blue trim.",  
"A catamaran is a type of boat that has two hulls, or platforms, that are parallel to each other.",  
"The image is of a yellow catamaran with white trim, sitting in calm water."



"A catamaran is a type of sailing vessel that consists of two parallel hulls of equal size."

"beach":

".",

"Blue skies, white sands, and clear turquoise waters make this beach a paradise.",

"A beach usually has sand and water.",

"Beach identification can be accomplished through the identification of physical characteristics.",

"A beach is a naturally occurring feature of the landscape.",

"The sun sets over the ocean, casting a beautiful orange hue in the sky.",

"A beach is a large body of water with sand or small rocks at the shore.",

"In the image, the beach is Brilliant white with crystal blue waters.",

"The beach looks like a long strip of land next to the ocean.",

"A beach typically looks like a large, flat expanse of sand with some rocks or other natural features nearby."

"rotary dial telephone":

"A rotary dial is a device used to dial telephone numbers.",

"Rotary dial telephone from the mid-20th century.",

"A rotary dial telephone is a phone with a circular dial on the front face.",

"When looking at a rotary telephone, you can tell it is a rotary phone by the Place the phone's receiver on your ear and listen for a dial tone.",

"A rotary dial phone is an older model phone that has a circular device with numbers on it that you rotate with your finger to dial a number.",

"A rotary dial telephone is a type of telephone that uses a mechanical dial to select the telephone number that a user wishes to call.",

"History of the Rotary Dial Telephone.",

"A rotary dial telephone is an old-fashioned telephone that has a round dial on the front of it.",

" rotary dial telephone looks like a classic telephone with a rotary dial on the base unit.",

"The image is of a rotary dial telephone on a black background."