

SPIN: Hierarchical Segmentation with Subpart Granularity in Natural Images

Josh Myers-Dean¹, Jarek Reynolds¹, Brian Price², Yifei Fan², and
Danna Gurari^{1,3}

¹ University of Colorado Boulder, ² Adobe, ³ University of Texas at Austin

Abstract. Hierarchical segmentation entails creating segmentations at varying levels of granularity. We introduce the first hierarchical semantic segmentation dataset with subpart annotations for natural images, which we call SPIN (**S**ub**P**art**I**mage**N**et). We also introduce two novel evaluation metrics to evaluate how well algorithms capture spatial and semantic relationships across hierarchical levels. We benchmark modern models across three different tasks and analyze their strengths and weaknesses across objects, parts, and subparts. To facilitate community-wide progress, we publicly release our dataset at <https://joshmyersdean.github.io/spin/index.html>.

1 Introduction

When people discuss *hierarchical image analysis* tasks they are typically talking about one of two approaches: **is-a** relationships [27,28], which treat categories at different abstraction levels (*e.g.*, a Subaru is a car), or **is-part-of** relationships, which concentrate on dividing objects into their constituent parts [13, 14, 19] (*e.g.*, a door is part of a car). While the former has been widely explored, the latter—focusing on object decomposition—has received limited attention in computer vision research. Within segmentation research, the focus of this paper, **is-part-of** relationships have primarily been explored only for part-whole hierarchies, ignoring finer-grained details, such as subparts (*i.e.*, parts of parts).

Research to subpart-level segmentation granularity is hindered by a scarcity of data. While synthetic 3D datasets [14, 19, 35, 63] could be used to infer hierarchical segmentations to subpart granularity, it is well-known that models developed with synthetic data typically generalize poorly to natural images (*i.e.*, images taken by a camera) [8]. This lack of annotated natural data has meant that the few models designed to generate subpart granularity segmentations [11, 52, 60] could only be evaluated qualitatively on a small number of examples (rather than quantitatively at scale).

Our work facilitates the development of algorithms for subpart granularity in three key ways. First, we collect over 102,000 subpart segmentations across 203 diverse subpart categories to expand upon PartImageNet [17] and create the **SubPartImageNet** (SPIN) dataset—the first to offer subpart annotations for natural images. We release this dataset publicly to foster community progress.




Fig. 1: Overview of the diversity of SPIN. Panels (a) and (b) depict subparts unique to specific object class members, such as a roll cage in a car and a shell in a turtle. Panels (c) and (d) illustrate the variability in the number of subparts per object of the same class, with examples of 13 and 6 subparts. Panels (e) and (f) highlight the disparity in image area coverage by different subparts, such as a bottle label (large) versus quadruped claws (tiny).

Examples of annotated subparts are shown in Fig. 1. Second, we introduce two novel metrics to address a shortcoming of the prevailing evaluation methods for hierarchical segmentation, which only assess each granularity level in isolation. Our newly proposed metrics are intended to work alongside these traditional metrics (*e.g.*, Intersection over Union) and assess how well models capture spatial and semantic relationships across hierarchical levels. Our third key contribution is that we benchmark modern models on SPIN across objects, parts, and subparts, thereby presenting the first comprehensive quantitative evaluation of subpart performance and key areas needed for future improvement.

Advancing hierarchical segmentation to subpart granularity could significantly benefit both research and societal applications. One potential use is in generating more detailed image descriptions, such as for augmented reality experiences, answering visual questions, captioning images, or visual storytelling. Similarly, this work could be the foundation for individuals to (1) interactively learn how to speak about finer-grained entities for the first time (for children) or in a new language and (2) recall forgotten words, such as when facing temporary or chronic memory injuries or disabilities. More deeply nested, hierarchical segmentation representations can also enrich tactile visual discovery [25, 42], including for use in screen readers employed by visually impaired individuals (*e.g.*, Apple’s object-based exploration feature [1]). Similarly, we believe finer-grained information could facilitate improvements for related problems, including image/video retrieval, image/video editing, automatic magnification, and robotics.

2 Related Works

Datasets with Hierarchically Segmented Objects. Several datasets provide segmentations showing how objects are hierarchically decomposed into their recursively nested parts. A few do so *without semantic labels* specifying the categories of segmented content, including the pioneering Berkeley dataset published in 2001 [36] and the large-scale SA-1B dataset released in 2023 [22]. Our work, in contrast, focuses on hierarchically segmenting objects *with category*

labels. Already, numerous datasets semantically decompose objects and their parts in natural images [7, 14, 15, 17, 21, 31, 45, 50, 56, 61, 67–69] and synthetic images [5, 19, 20, 35, 39]. One dataset, (i.e., ADE20K [69]) even provides subpart annotations for 10% of its objects, however its lack of exhaustive annotations for objects and their associated parts [52] impedes its benefit for evaluating prediction models. Extending prior work, we introduce the first exhaustively labelled hierarchical semantic segmentation dataset with subpart annotations for natural images by extending one of the largest part-based datasets [17].

Hierarchical Semantic Segmentation Evaluation. Most evaluation protocols for hierarchical segmentation algorithms assess each hierarchy level *independently*, such as by reporting for object and part categories separately their Intersection over Union scores [9, 23, 31, 32, 37, 40, 52, 55, 57, 58, 61, 65], Average Precision scores [45, 51], and Panoptic Quality scores [14, 29, 30, 52, 60]. The one exception is Hierarchical Panoptic Quality [52] which recursively measures the extent to which a predicted hierarchy is complete (e.g., a wheel is a part of a car). Complementing prior work, we introduce evaluation metrics for assessing how well a model captures spatial and semantic hierarchical relationships, such as if a part is perfectly contained in its whole.

Hierarchical Semantic Segmentation Algorithms. Thus far, the focus of methods that predict all entities in a hierarchical decomposition largely has centered on predicting just objects and parts [16, 51, 58], aligning with existing datasets which only recently increasingly have added part segmentations to object segmentations. Only a few works, HIPIE [60], VDT [11], Semantic-SAM [26], and ViRReq [52]¹, examined predicting more layers of a hierarchy to also include subparts. However, these works relied on qualitative assessments for subpart predictions due to a lack of annotated datasets supporting quantitative assessments. Complementing prior work, we quantitatively evaluate HIPIE and over 20 model variants to show their performance on our new dataset challenge across all three layers of the hierarchy (*i.e.*, object, parts, and subparts).

3 SPIN Dataset

We now introduce and characterize our new subpart semantic segmentation dataset that we call SPIN, which is short for **SubPart ImageNet**.

3.1 Dataset Creation

Data Source. Our dataset extends an existing dataset that provides both object and associated part annotations: PartImageNet [17]. It includes 24,080 natural images across 158 ImageNet categories (112 non-rigid and 38 rigid, *e.g.* animals and vehicles respectively). The authors of PartImageNet used the

¹ No code is publicly available for VDT, and ViRReq does not offer complete code. At the time of writing, Semantic-SAM has not released their semantic prediction code.

WordNet [38] taxonomy to establish these categories and their 11 parent super-categories. The authors then identified 40 part categories pertinent to the super-categories and annotated these in all images. Every image in PartImageNet contain one semantically segmented object, and we used only those with at least one segmented part. Due to our limited annotation budget, we capped each super-category at 1,200 images. This resulted in the following amounts per super-category: 311 - Aeroplane, 483 - Bottle, 559 - Boat, 634 - Bicycle, and 1,200 each for Biped, Bird, Car, Fish, Quadruped, Reptile, and Snake.

Subpart Category Selection. We identified 206 subpart categories to segment for 34 PartImageNet part categories that we identified as being decomposable into subparts. Of these, 168 subpart categories generalize across many part categories and 38 categories pertain to only a few. For example, in our dataset, *mouths* generalize while *shells* don't, with most reptiles having heads with *mouths* but only one type of reptile having a body with a *shell* (*i.e.*, turtles). We show examples of specific subparts in Fig. 1(a-b). A histogram showing how many subparts are assigned to each of the 34 part categories is shown in Fig. 2, and a full list is provided in the supplementary materials.

We created our subpart taxonomy through a multi-step process. We identified the general subpart categories by prompting GPT-4 [2] to list expected subparts for every object-part category in PartImageNet,² and then three authors edited the list to exclude non-visible subparts (*e.g.*, a skull). We identified specific subpart categories through manual inspection of at least 15 images per object category, excluding any categories that had ambiguous subpart decompositions; for example, it's not obvious to a lay person what boundaries to use when decomposing a tail into a tip, shaft, and base.

Subpart Annotation Task Design. We created a task interface for collecting all subpart annotations. It presents each image-object pair alongside one of the object's nested parts, outlined with an overlaid polygon. For each part, the task

² We prompted GPT-4 with "Please list the canonical subparts of a <object>-<part>. Only include subparts that are clearly visible and recognizable to a layperson."




Fig. 2: Histogram the number of unique subpart category labels for each of the 34 part categories. (Aero=Aeroplane; Quad=Quadruped)

#Images	#Object Cat	#Part Cat	#Sub-Part Cat
10,387	11 (10,387)	40 (29,818)	203 (106,324)

Table 1: SPIN Composition: Each column lists category counts alongside the total number of annotated entities within those categories. (Cat=Categories)

interface has multiple steps. First, users are asked, “Can you locate the <object>-<part> enclosed in a polygon?” If the target content is present, annotators are then asked for each possible subpart, “Can you locate any <subpart> on the <object>-<part>?” If the response is yes, annotators segment the subpart. The interface supports segmentation by collecting a series of clicked points to create a connected polygon. It also supports annotating multiple polygons in three scenarios: (1) multiple instances of a subpart (*e.g.*, two eyes), (2) subparts with holes (*e.g.*, coiled snakes), and (3) occlusions breaking a subpart into multiple, disconnected pieces.

Annotation Collection. We hired 18 highly trusted annotators from Amazon Mechanical Turk (AMT) who we had vetted through repeated employment for previous segmentation tasks. During data collection, we ensured high annotation quality through five methods: an onboarding qualification test, detailed instructions, live “office hours” during annotation periods, phased task roll-out for worker feedback, and continuous inspection of submitted results.

Dataset Splits. Our dataset adheres to the PartImageNet training, validation, and testing splits of 85%, 5%, and 10% of the data, respectively. This results in 8,828 training, 519 validation, and 1,040 test images in our SPIN dataset.

3.2 Dataset Analysis

We now characterize SPIN’s overall composition as well as its subparts.

Overall Dataset Composition. We characterize SPIN with respect to the following seven factors: number of images, number of object categories, total number of annotated objects, number of part categories, total number of annotated parts, number of subpart categories, and total number of annotated subparts. Results are shown in Table 1. As shown, our contribution of 203 entity categories is over a five-order-of-magnitude increase compared to the 40 categories added for PartImageNet [17], while providing nearly an order of magnitude more annotations (*i.e.* 106,324 semantic annotations versus 11,960 part segmentations in PartImageNet).

Subpart Statistics. We next compute the mean number of subparts per part and characterize the typical appearance of subparts with respect to five metrics:

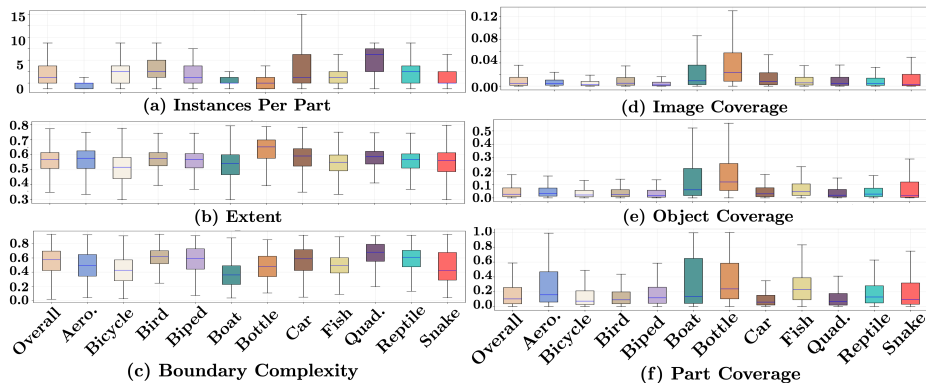


Fig. 3: Boxplots showing the distribution of subpart image occupation (left) and boundary complexities per part, per object (right). The blue lines represent medians, bottoms and tops of each box represent the 25th and 75th percentile values respectively, and whiskers represent the most extreme data points not considered outliers. Overall, SPIN’s subparts take up a relatively small number of pixels per image, while featuring a range of geometric complexity. (Aero=Aeroplane; Quad=Quadruped)

- **Boundary complexity:** ratio of a subpart’s area to the length of its perimeter (*i.e.*, isoperimetric quotient). Values range from 0 (highly jagged boundary) to 1 (circular).
- **Extent:** the ratio of the area of a contour to the area of its bounding box. Values are in $(0,1]$, where values approaching 0 mean that a contour occupies little area in it’s bounding box (*e.g.*, a thin diagonal line) and 1 means that a contour is perfectly contained (*e.g.*, a square).
- **Image coverage:** fraction of image pixels occupied by the subpart.
- **Object coverage:** fraction of parent object pixels occupied by the subpart.
- **Part coverage:** fraction of parent part pixels occupied by the subpart.

For all six factors, we report both the overall mean as well as the mean with respect to each object category. Results are shown in Fig. 3.

Objects tend to have 2 to 5 subparts per part (*i.e.* Fig. 3a; 25th to 75th percentile values), with a total of 4 to 15 subparts. There is diversity across object categories with aeroplanes typically having small variability with only 1 to 2 subparts per part (*i.e.* 25th to 75th percentile range) while car typically exhibits great variability with 2 to 8 subparts per part (*i.e.* 25th to 75th percentile range), as exemplified in Fig. 1(c-d). We attribute the large range of subpart per part counts to two factors: the intrinsic diversity in cars and the viewing angle. For example, a go-cart viewed from the front will have no windshield, no windows, and two tires, while a bus viewed from a front-right angle will have a windshield, tires, and windows. Such object-specific patterns around the expected prevalence of subparts could be predictive cues for models decomposing detected objects into their recursively nested subparts.

In terms of shape, most subparts’ contours occupy the majority of their bounding box (*i.e.*, Fig. 3b, median extent value ~ 0.7) with moderate boundary complexity (*i.e.*, Fig. 3c, median boundary complexity values around 0.5). Still, we observe high variability for both metrics that we again attribute to the intrinsic diversity in each object category as well as the viewing angle. For example, a bottle cap or a bicycle tire viewed from the side will be an ellipse versus a nearly perfect circle when viewed from the top (*i.e.*, occupying more of it’s bounding box). Altogether, these findings highlight that our dataset encourages the design of models that can account for a large range of subpart shapes.

Regarding the relative area occupied by each subpart, we find that all subparts occupy small portions of the parent parts, parent objects, and entire image. For example, across every object category, subparts occupy less than 6% of the image for at least 75% of subpart instances (Fig. 3d). Additionally, the subparts tend to occupy less than 20% of their parent parts (Fig. 3e; Overall) and less than 5% of their parent objects (Fig. 3f; Overall). When adopting size thresholds introduced for the MSCOCO dataset [33], where 32^2 and 96^2 are thresholds determining whether an object is small, medium, or large, we find 54.10% (57,525) of SPIN’s subpart annotations qualify as small, 38.08% (40,488) as medium, and only 7.82% (8,311) as large. Most subparts are already considered small according to mainstream research despite that we annotated subparts for a dataset known to typically contain a single, large prominent object in the image. Consequently, embedded in our problem is the well-known challenge of small entity detection [53], which will only grow as subparts are explored for more complex scenes where less prominent objects must also be hierarchically decomposed.

When comparing subparts belonging to the rigid object categories (*i.e.*, human-made) versus non-rigid object categories (*i.e.*, animals), we found no clear distinction. For instance, the mean number of subparts per part (Fig. 3a) is the smallest amount with least variability in aeroplanes and the greatest amount with greatest variability in cars. While prior work has emphasized distinguishing between these two extremes [17], our analysis suggests that rigid and non-rigid objects appear quite similar at the subpart level.

4 Evaluating Hierarchical Consistency

We facilitate assessing how well relationships across varying levels of granularity are captured for hierarchical segmentation by introducing two metrics that account for spatial and semantic relationships respectively.

We first introduce the **Spatial Consistency Score (SpCS)**, which focuses on structural hierarchical alignment rather than semantic alignment. It indicates the proportion of a child’s segmentation (*e.g.*, subparts) that is contained in its super-region’s segmentation (*e.g.*, part, object), with values ranging from 0 (no containment) to 1 (perfect containment). It leverages binarized segmentations distinguishing the foreground entity from the background. Formally, it is com-

puted as the average containment ratio across all pairs as follows:

$$SpCS = \frac{1}{|\mathcal{R}|} \sum_{(\text{child}, \text{parent}) \in \mathcal{R}} \underbrace{\frac{|\text{child} \cap \text{parent}|}{|\text{child}|}}_{\text{containment ratio}} \quad (1)$$

where \mathcal{R} denotes the set of prediction pairs with a ground truth parent-child relationship within a hierarchical structure.³ $|\mathcal{R}|$ denotes the total number of parent-child pairs, $|\text{child} \cap \text{parent}|$ is the number of pixels in the intersection of the child segmentation with the parent, $|\text{child}|$ is the number of pixels in the child segmentation.

We next introduce the **Semantic Consistency Score (SeCS)**, which focuses on semantic alignment by measuring how well predictions at different levels (subpart, part, object) logically correspond. For example, an eye (subpart) should correspond to a head (part) and, by extension, to a quadruped (object), rather than illogical associations like a windshield (subpart) with a head (part) or a bottle (object). Let FG_S , FG_P , and FG_O be the sets of non-background predictions in subpart, part, and object predictions, respectively. Define $X = FG_S \cap FG_P \cap FG_O$ as the intersection of these predictions, representing pixels with consistent foreground labels across levels. For each pixel $x \in X$, with predicted categories $C(S_x)$, $C(P_x)$, and $C(O_x)$ for subparts (S_x), parts (P_x), and objects (O_x), respectively, category entailment is evaluated using $\mathbb{1}[C(S_x), C(P_x)]$ and $\mathbb{1}[C(P_x), C(O_x)]$, where $\mathbb{1}$ is an indicator function returning 1 if the hierarchical entailment is correct, using ground truth relations, otherwise 0. We compute the entailment measure $M(x)$ as:

$$M(x) = \begin{cases} 1 & \text{if } \mathbb{1}[C(S_x), C(P_x)] \wedge \mathbb{1}[C(P_x), C(O_x)], \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

where \wedge is the logical **and** operator. SeCS is the mean of $M(x)$ across all $x \in X$. Values range from 0 (no semantic coherence) to 1 (perfect semantic coherence).

5 Algorithm Benchmarking

We now assess modern models’ ability to account for hierarchical decompositions to the subpart granularity.

5.1 Open-Vocabulary Localization in a Zero-Shot Setting

We first evaluate models for open-vocabulary localization in zero-shot mode across three hierarchy levels: object, part, and subpart. This approach avoids the drawbacks of the alternative options of training models from scratch or fine-tuning, including the computational expense and risks of overfitting to the target task and so diminishing generalizable knowledge.

³ For a quadruped, for instance, pairs such as {(eyes, head), (chest, torso), (torso, quadruped)} could be present in \mathcal{R} .

Models. We benchmark 11 model variants.

Most related to our work is HIPIE [60], which is designed for open-vocabulary hierarchical segmentation to the subpart level. It takes as input a list of all candidate categories that can appear and then predicts which categories are present where in the given image. We test two variants with different backbones: ResNet-50 [18] and ViT-H [12].

We also evaluate 10 variants of foundation models that, by design, are intended for use on a variety of downstream tasks. Five produce pixel-wise segmentation masks: LISA-7B [24], LISA-13B [24], GLaMM [46], and PixelLLM-7B [48], and PixelLLM-13B [48]. The other five can only produce bounding boxes: CoGVLM [59], Ferret-7B [64], Ferret-13B [64], Shikra [6], and Kosmos2 [43]. For these models, we analyze their upper bound of their potential performance by notifying the model what object category needs to be located in a given image. We examine two types of requests: "Can you please locate the <object> in the image?" for objects, and "Can you please locate the <(sub)part> of the <object> in the image?"⁴ for each part and subpart. We evaluate with these prompts a **Specific** category by using the original WordNet category from the 158 ImageNet categories as well as a **General** category by referring to the 11 super-categories. Consequently, we test each model with four prompts.

Evaluation Metrics. We use two evaluation metrics for all models: (1) Intersection over Union (IoU) to assess segmentation performance for objects, parts, and subparts independently and (2) our Spatial Consistency Score (SpCS) to assess structural coherence across object, part, and subpart levels. We also use our new Semantic Consistency Score (SeCS) to assess semantic coherence between predicted hierarchical levels for the only relevant model, HIPIE (it predicts from a list of candidate categories rather than for the target, single category).

Localization Results (mIoU). Results are shown in Table 2.

We observe a consistent trend across all models: localization performance is best for object segmentation, followed by increasingly worse performance for more granular categories of parts and then subparts. The best score for localizing objects is 86 followed by 40 for parts and 14 for subparts. The worst scores come from the only model directly designed for hierarchical segmentation, HIPIE. It does worse across all granularity levels than the foundation models; *e.g.*, for subpart localization, 1 vs 3 from the worst-performing foundation model, Kosmos2. We attribute the poor performance to HIPIE’s limited training data (*e.g.*, limited part vocabulary), which only included part-level categories from Pascal Parts [7]. Moreover, HIPIE’s heuristic grouping assumes that parts are the sum of their subparts, which is not the case in SPIN (*e.g.*, unlabelled "space" exists between a cheek and an eye in a head).

We explored to what extent entity size is correlated to the resulting IoU scores by using linear regression for the better-performing foundation models.

⁴ For subparts, we refer to the parent object rather than the parent part, as broader context aids in processing finer details [47].

Method	# Params.	mIoU _S		mIoU _P		mIoU _O		SpCS - Avg		SpCS - S2P		SpCS - P2O	
		Specific	General	Specific	General	Specific	General	Specific	General	Specific	General	Specific	General
HIPIE [60] (R-50)	200M	0.80	0.10	8.05	7.22	51.36	13.87	97.82	98.77	100	100	95.64	97.54
HIPIE [60] (ViT-H)	800M	0.90	0.92	7.21	8.23	66.77	21.69	98.04	99.25	100	98.50	96.08	96.06
PixelLLM [48]	7B	9.08	9.24	32.37	31.37	83.93	79.87	73.79	82.46	69.57	79.20	89.49	90.73
PixelLLM [48]	13B	9.53	10.13	32.04	32.96	82.90	79.96	84.86	83.06	82.13	80.87	92.60	88.90
LISA [24]	7B	9.61	9.14	27.72	27.28	86.23	83.36	79.30	77.44	74.61	72.93	92.10	89.78
LISA [24]	13B	11.52	11.55	32.29	31.36	87.78	85.45	86.28	79.98	82.87	75.70	96.02	92.41
GLaMM [46]	7B	11.03	11.00	39.41	40.00	86.29	86.31	87.12	87.75	84.21	84.96	95.72	95.38
GLaMM - FT	7B	24.25	24.56	59.37	60.76	86.42	91.08	87.13	87.95	75.93	85.16	90.23	96.04
Ferret [64]	7B	7.22	7.30	25.69	26.54	47.84	47.98	78.70	73.02	46.50	65.68	93.65	94.60
Ferret [64]	13B	6.37	6.25	23.67	23.65	47.99	47.50	81.04	77.30	48.70	70.98	94.17	95.40
CoGVLM [59]	17B	13.94	14.29	38.13	38.56	46.47	43.65	76.33	77.20	72.37	73.32	86.19	87.59
Shikra [6]	7B	8.41	8.72	26.50	27.42	45.50	29.14	68.57	67.01	63.82	66.76	81.85	67.89
Kosmos2 [43]	2.6B	3.45	3.59	19.00	19.48	48.63	48.89	82.74	81.60	78.34	77.03	96.01	95.58
SAM [22]	630M	49.61	49.61	69.23	69.23	90.06	90.06	86.42	86.42	83.85	83.85	92.30	92.30
SAM + CoGVLM	630M	19.88	19.98	50.93	50.90	77.94	75.09	66.71	67.33	80.50	80.73	60.65	61.44

Table 2: Performance of modern models with respect to localization at three granularity levels—subparts (mIoU_S), parts (mIoU_P), objects (mIoU_O)—and hierarchical consistency at two levels—subparts to parts (SpCS-S2P), parts to objects (SpCS-P2O), and their mean (SpCS - Avg). Results are shown for hierarchical semantic segmentation models (described in Section 5.1), hierarchical object detection models (Section 5.1), and interactive segmentation models that perform category-agnostic localization (Section 5.2). The method families are separated by dashed lines in that order. Best performing scores are shown in bold.

Due to space constraints, we provide further experimental details and results in the supplementary materials. In summary though, our findings suggest that size has little correlation to performance for objects but greater correlations for more granular levels of parts and subparts. We suspect that one contributing factor is the sensitivity of the IoU metric for entities occupying a small number of pixels, as small errors are especially prone to yield dramatic changes in the IoU scores for such cases.

We next explored the impact of a prompt’s category specificity on model performance by comparing performance for when general versus specific terms are employed. Most models perform worse when trying to locate *objects* categorized under general terms (*e.g.*, ‘quadruped’) rather than specific ones (*e.g.*, ‘dog’). For example, CoGVLM performs 2.82 percentage points (pp) worse with general categories compared to specific categories. Similarly, the 13B variant of LISA has a 2.33pp performance drop and its 7B counterpart has a 2.87pp performance drop. The exceptions to this are GLaMM, Ferret, and Kosmos2, whose performance are relatively stable and unaffected by the level of category specificity. A potential reason is that these models utilize training data that explicitly use both general and specific terms. For example, GCG [46], a source of training data for GLaMM, refers to entities with both abstract and specific terms (*e.g.*, person and toddler). Overall, this variance highlights the challenge in choosing an appropriate level of category specificity for *objects*. However, the story changes when trying to locate *parts* and *subparts*, with the performance gap between general and specific categories falling to within 1pp. An interesting direction for future work is better understanding how and why an entity’s granularity level influences models’ abilities to overcome abstraction challenges.

We also examine how the training data of the Large Language Models (LLM) Llama [54], since it serves as the basis for the majority of foundation models evaluated in our study. Specifically, we analyze the empirical frequency (*i.e.*, total occurrence) of uni-grams (*e.g.*, eye) and n-grams (*e.g.*, eye of the quadruped) within the dataset using the ∞ -gram API [34], which uses byte array pattern matching to find all occurrences of a given BPE tokenized [49] n -gram. Due to space constraints, results are provided in the supplementary materials. In summary though, we observe a decline in average uni-gram occurrence frequency when increasing granularity (*e.g.*, part to subpart). An interesting direction for future work is to remedy this imbalance, such as by designing parameter-efficient methods to restructure the text embedding space of LLMs to recognize hierarchical relationships, akin to MERU [10] restructuring CLIP [44]-space for image-level is-a relationships.

Performance When Training on Our SPIN Dataset. We next examine the potential benefit of our SPIN training data for modern models. To do so, we fine-tuned the top-performing GLaMM algorithm [46] for all hierarchy levels in SPIN, and we refer to this variant as GLaMM-FT. Results are shown in Table 2.

Overall, we observe a considerable performance boost. For example, when prompting with general categories, we observe a 13.56 percentage point (pp) increase (123% relative increase) for subparts, a 20.76pp increase (51.90% relative increase) for parts, and a 4.77pp increase for objects (5.51% relative increase). While such performance gains are promising, we still observe very low absolute scores for subparts, indicating their challenge for modern models. We suspect a promising direction for future work is providing cost-effective ways to further increase the amount of available training data, such as through synthetic data creation and 2D projections from 3D models.

Hierarchical Consistency Results. SpCS results are shown in Table 2. Similar trends are observed for part-object and subpart-part relationships. Part-object pairs consistently exhibit high SpCS-P2O scores, indicating models tend to correctly predict parts within their corresponding whole objects. This provides evidence that models’ have some understanding of spatial containment among part-object relationships. A potential reason for this is that models are indirectly trained on part-whole relationships, because they are already taught that an object is a part of a scene. The aforementioned trend is slightly worse when examining consistency between subparts and parts (*i.e.*, SpCS-S2P), with scores varying from 46.50 (*i.e.*, Ferret 7B) to 84.21 (*i.e.*, GLaMM) when employing general terms for objects and from 65.68 (*i.e.*, Ferret 7B) to 84.96 (*i.e.*, GLaMM) when employing more specific categories for objects. This variability, coupled with low mIoU scores for subparts, reveals that models often predict subparts outside their intended part boundaries, underscoring poor performance in properly contextualizing subparts within an object.

5.2 Interactive Segmentation for Vocabulary-Free Localization

We next assess to what extent models can *localize* entities at all hierarchy levels in our dataset if they are notified where to look in an image. This sets an upper bound of what is possible. To do so, we benchmark the state-of-the-art interactive segmentation model, SAM [22], which takes as input a coarse marking on an image specifying where to create the segmentation. We use bounding boxes as our coarse image markings, because they are shown to outperform alternatives [41] when using SAM. We explore the upper bound performance for this model by using the ground truth object, part, and subpart segmentations to generate the bounding boxes (*i.e.*, by using $[x_{min}, y_{min}, x_{max}, y_{max}]$ for each entity’s pixels). Results are shown in Table 2.

Despite the idealized guidance from the bounding box prompts, SAM’s performance is still imperfect. As granularity increases, localization performance consistently decreases by about 20pp when going down a granularity level (*e.g.*, objects to parts), revealing that SAM struggles more as segmentation requests become more fine-grained. Still, our findings demonstrate SAM’s potential for yielding considerably higher quality hierarchical segmentations than observed from existing localization models benchmarked in the previous section. Compared to the top-performing zero-shot segmentation-based model (*i.e.*, GLaMM), SAM achieves a boost of 38.61pp for subparts, 20.23pp for parts, and 3.75pp for objects, alongside similar SpCS scores.

We also evaluate a fully-automated system with SAM, by incorporating bounding box predictions from the top-performing bounding box producing model, CoGVLM. Results are shown in Table 2. We similarly observe that it outperforms the leading zero-shot segmentation method, GLaMM, in segmenting *subparts* and *parts* (9.24 percentage point (pp) boost for subparts and 10.87pp for parts). We attribute these performance gains to SAM’s ability to leverage a location prior (*e.g.*, bounding box), whereas models like GLaMM rely solely on language priors. In contrast, this approach performs worse than the leading zero-shot segmentation method, GLaMM, for *objects*, with a 11.22pp decrease. We attribute this performance drop to CoGVLM’s inferior performance in consistently and accurately detecting bounding boxes.

5.3 Recognizing Hierarchical Semantics in a Zero-Shot Setting

We finally evaluate the ability of models to *recognize* hierarchical semantic labels for objects in images (*i.e.*, without localization). We assess this in an idealized scenario in which we notify models what to look for as well as where to look in the image. This sets an upper bound of what may be possible with such models as they are explicitly told *where* to look.

Models. We evaluate eight off-the-shelf model variants: CoGVLM [59], Ferret-7B [64], Ferret-13B [64], Shikra [6], Osprey [66], Kosmos2 [43], ViP-Llava-7B [4], and ViP-Llava-13B [4]. All prompts to the models resemble the following basic request: "Is there a <object> in the <region>?" for objects and "Is there

Method	# Params.	Prompt	mACC _S	mACC _{SS}	mACC _P	mACC _{PS}	mACC _O	mACC _{OS}
Shikra [6]	7B	Bbox	66.12	61.28	79.67	74.56	77.04	82.05
Ferret [64]	7B	Bbox	59.01	57.19	59.95	66.17	86.78	89.21
Ferret [64]	13B	Bbox	50.31	69.42	61.72	67.84	88.93	89.02
CoGVLM [59]	17B	CoT	42.16	45.54	81.75	83.31	87.15	83.67
Kosmos2 [43]	2.6B	SG	79.12	64.32	74.23	68.53	84.55	78.78
Osprey [66]	7B	Mask	77.43	55.95	86.30	81.20	96.43	84.85
ViP-Llava [4]	7B	Mask	53.57	58.52	64.13	64.96	71.62	65.25
ViP-Llava [4]	13B	Mask	94.78	97.43	99.90	99.80	98.33	95.06

Table 3: Accuracy metrics for object hierarchical decomposition: general (mACC_S, mACC_P, mACC_O) and specific (mACC_{SS}, mACC_{PS}, mACC_{OS}) categories across subparts, parts, and objects. (CoT = Chain of Thought; SG = Self-Grounding)

a <(sub)part> of the <object> in the <region>?" for parts and subparts. Regions are specified either with segmentations (ViP-Llava, Ferret, Osprey) or bounding boxes (CoGVLM, Shikra, Kosmos2). We generate these regions directly from the ground truth annotations (*e.g.*, bounding boxes as described in Section 5.2) for objects, parts, and subparts. We prompt models for each category known to be in the region, one at a time. An answer is ‘yes’ if ‘yes’ is present in the model’s output, and otherwise ‘no’. As done in Section 5.1, we test both general and specific categories in the prompts.

Evaluation Metrics. We use *accuracy* to evaluate at each granularity level.

Results. Results are reported in Table 3.

While most models struggle overall to recognize categories in images (despite being told where to look in the images), one model has nearly perfect accuracy: ViP-Llava-13B. One hypothesis for this atypically strong performance is the model lacks understanding and instead always answers yes when prompted about the presence of categories. We explored this with further experiments using adversarial examples that are instead expected to always answer “no”.⁵ However, our findings from these experiments are none-conclusive. We observe the model can answer “no”, but it also suffers in this case from a significant decrease in accuracy for the adversarial prompts. This performance drop could be due to a separate issue that the model struggles with negation (*i.e.*, the reversal curse [3]). Altogether, this evaluation represents the best-case scenario for performance; it is likely to fall further in real-world applications due to imperfect human inputs or model-generated segmentations, both of which we have shown to exhibit poorer performance at finer granularity levels, as discussed in Sections 5.1 and 5.2.

We observed mixed outcomes regarding which granularity levels are most difficult to recognize. For Ferret-13B and Osprey, an increase in granularity correlates with improved accuracy; *e.g.*, transitioning from subparts to objects using general category names results in a 28.62pp increase for Ferret-13B and a 19.00pp

⁵ Results are shown in the supplementary materials for two prompts asking “Is the category not present” and “Is the [different category] present”.

increase for Osprey. Conversely, for most other models, object identification consistently outperforms subpart identification. The differential performance across granularity levels may partly stem from the disparate prevalence of tokens in the training data across different granularity levels, particularly within the Llama dataset (as previously discussed). For parts and subparts, the mixed results observed—where parts sometimes outperform subparts (*e.g.*, CoGVLM) and vice versa (*e.g.*, Ferret-7B)—could be attributed to their relative frequency in the training data. Future research could focus on examining how foundation models differentiate responses based on prompt granularity and on enriching datasets like GRIT [64] with hierarchical relationships to improve model understanding of intra-entity relationships. Additionally, future work could explore enhancing hierarchical-relationship understanding so success at different hierarchical levels can facilitate finding entities at other granularity levels.

We also observed mixed outcomes when comparing the efficacy of specific versus general categories. Half the models performed better using general terms for subparts, with a similar pattern observed in parts and objects. A potential reason for this discrepancy, compared to the relatively consistent results in Sec 5.1, is that localization models often rely on feature similarities, such as CLIP [44] features. Even if a model does not capture every nuance of a category query, it still attempts find a match based on overall similarity to known categories (*e.g.*, mini-van vs. van) and is less likely to reject user queries [62]. Conversely, interactive understanding models, designed to maximize accuracy, may strictly search for precise category matches within their trained representations. Given our task’s structure (*i.e.*, answers are restricted to ‘yes’ or ‘no’), these models are predisposed to outright reject the query (*i.e.*, answer ‘no’) if an exact match is absent, rather than proposing close or related category matches.

6 Conclusion

We introduce SPIN, the first dataset challenge for hierarchical segmentation at the subpart granularity in natural images. Our analysis reveals SPIN’s characteristics, our two new evaluation metrics enable gauging algorithmic proficiency in capturing the spatial and semantic relationships across hierarchy levels, and our benchmarking across three tasks demonstrates that models struggle to recognize and segment subparts in SPIN, even under idealized guidance. We publicly release SPIN to encourage further advancements in hierarchical segmentation, sharing the dataset at <https://joshmyersdean.github.io/spin/index.html>.

Acknowledgments. This work was supported by Adobe Research Gift Funds and utilized the Blanca condo computing resource at the University of Colorado Boulder. Josh Myers-Dean is supported by a NSF GRFP fellowship (#1917573). We thank the crowdworkers for contributing their time for the construction of SPIN and the authors of our benchmarked models for open-sourcing their work.

References

1. Explore images. <https://support.apple.com/guide/iphone/use-voiceover-for-images-and-videos-iph37e6b3844/ios>
2. Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)
3. Berglund, L., Tong, M., Kaufmann, M., Balesni, M., Stickland, A.C., Korbak, T., Evans, O.: The reversal curse: Llms trained on " a is b" fail to learn " b is a". arXiv preprint arXiv:2309.12288 (2023)
4. Cai, M., Liu, H., Mustikovela, S.K., Meyer, G.P., Chai, Y., Park, D., Lee, Y.J.: Making large multimodal models understand arbitrary visual prompts. In: IEEE Conference on Computer Vision and Pattern Recognition (2024)
5. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
6. Chen, K., Zhang, Z., Zeng, W., Zhang, R., Zhu, F., Zhao, R.: Shikra: Unleashing multimodal llm’s referential dialogue magic. arXiv preprint arXiv:2306.15195 (2023)
7. Chen, X., Mottaghi, R., Liu, X., Fidler, S., Urtasun, R., Yuille, A.: Detect what you can: Detecting and representing objects using holistic models and body parts. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1971–1978 (2014)
8. Deitke, M., Han, W., Herrasti, A., Kembhavi, A., Kolve, E., Mottaghi, R., Salvador, J., Schwenk, D., VanderBilt, E., Wallingford, M., Weihs, L., Yatskar, M., Farhadi, A.: Robothon: An open simulation-to-real embodied ai platform. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
9. Deng, B., Genova, K., Yazdani, S., Bouaziz, S., Hinton, G., Tagliasacchi, A.: Cvxnet: Learnable convex decomposition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 31–44 (2020)
10. Desai, K., Nickel, M., Rajpurohit, T., Johnson, J., Vedantam, S.R.: Hyperbolic image-text representations. In: International Conference on Machine Learning. pp. 7694–7731. PMLR (2023)
11. Ding, M., Shen, Y., Fan, L., Chen, Z., Chen, Z., Luo, P., Tenenbaum, J.B., Gan, C.: Visual dependency transformers: Dependency tree emerges from reversed attention. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14528–14539 (2023)
12. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
13. Geng, H., Xu, H., Zhao, C., Xu, C., Yi, L., Huang, S., Wang, H.: Gapartnet: Cross-category domain-generalizable object perception and manipulation via generalizable and actionable parts. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7081–7091 (2023)
14. de Geus, D., Meletis, P., Lu, C., Wen, X., Dubbelman, G.: Part-aware panoptic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5485–5494 (2021)

15. Gong, K., Liang, X., Li, Y., Chen, Y., Yang, M., Lin, L.: Instance-level human parsing via part grouping network. In: Proceedings of the European conference on computer vision (ECCV). pp. 770–785 (2018)
16. He, J., Chen, J., Lin, M.X., Yu, Q., Yuille, A.L.: Compositor: Bottom-up clustering and compositing for robust part and object segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11259–11268 (June 2023)
17. He, J., Yang, S., Yang, S., Kortylewski, A., Yuan, X., Chen, J.N., Liu, S., Yang, C., Yu, Q., Yuille, A.: Partimagenet: A large, high-quality dataset of parts. In: European Conference on Computer Vision. pp. 128–145. Springer (2022)
18. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
19. Hong, Y., Li, Q., Zhu, S.C., Huang, S.: Vlgrammar: Grounded grammar induction of vision and language. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1665–1674 (2021)
20. Hong, Y., Yi, L., Tenenbaum, J., Torralba, A., Gan, C.: Ptr: A benchmark for part-based conceptual, relational, and physical reasoning. *Advances in Neural Information Processing Systems* **34**, 17427–17440 (2021)
21. Jia, M., Shi, M., Sirotenko, M., Cui, Y., Cardie, C., Hariharan, B., Adam, H., Belongie, S.: Fashionpedia: Ontology, segmentation, and an attribute localization dataset. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16. pp. 316–332. Springer (2020)
22. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollar, P., Girshick, R.: Segment anything. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 4015–4026 (October 2023)
23. Koo, J., Huang, I., Achlioptas, P., Guibas, L.J., Sung, M.: Partglot: Learning shape part segmentation from language reference games. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16505–16514 (2022)
24. Lai, X., Tian, Z., Chen, Y., Li, Y., Yuan, Y., Liu, S., Jia, J.: Lisa: Reasoning segmentation via large language model. arXiv preprint arXiv:2308.00692 (2023)
25. Lee, J., Peng, Y.H., Herskovitz, J., Guo, A.: Image explorer: Multi-layered touch exploration to make images accessible. In: Proceedings of the 23rd International ACM SIGACCESS Conference on Computers and Accessibility. ASSETS '21, Association for Computing Machinery, New York, NY, USA (2021). <https://doi.org/10.1145/3441852.3476548>, <https://doi.org/10.1145/3441852.3476548>
26. Li, F., Zhang, H., Sun, P., Zou, X., Liu, S., Yang, J., Li, C., Zhang, L., Gao, J.: Semantic-sam: Segment and recognize anything at any granularity. arXiv preprint arXiv:2307.04767 (2023)
27. Li, L., Zhou, T., Wang, W., Li, J., Yang, Y.: Deep hierarchical semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1246–1257 (2022)
28. Li, T., Gupta, V., Mehta, M., Srikumar, V.: A logic-driven framework for consistency of neural models. arXiv preprint arXiv:1909.00126 (2019)
29. Li, X., Xu, S., Yang, Y., Cheng, G., Tong, Y., Tao, D.: Panoptic-partformer: Learning a unified model for panoptic part segmentation. In: European Conference on Computer Vision. pp. 729–747. Springer (2022)

30. Li, X., Xu, S., Yang, Y., Yuan, H., Cheng, G., Tong, Y., Lin, Z., Yang, M.H., Tao, D.: Panopticpartformer++: A unified and decoupled view for panoptic part segmentation. arXiv preprint arXiv:2301.00954 (2023)
31. Liang, X., Gong, K., Shen, X., Lin, L.: Look into person: Joint body parsing & pose estimation network and a new benchmark. *IEEE transactions on pattern analysis and machine intelligence* **41**(4), 871–885 (2018)
32. Liang, X., Shen, X., Feng, J., Lin, L., Yan, S.: Semantic object parsing with graph lstm. In: *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*. pp. 125–143. Springer (2016)
33. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*. pp. 740–755. Springer (2014)
34. Liu, J., Min, S., Zettlemoyer, L., Choi, Y., Hajishirzi, H.: Infini-gram: Scaling unbounded n-gram language models to a trillion tokens. arXiv preprint arXiv:2401.17377 (2024)
35. Liu, Q., Kortylewski, A., Zhang, Z., Li, Z., Guo, M., Liu, Q., Yuan, X., Mu, J., Qiu, W., Yuille, A.: Cgpart: A part segmentation dataset based on 3d computer graphics models. arXiv preprint arXiv:2103.14098 (2021)
36. Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In: *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001. vol. 2*, pp. 416–423. IEEE (2001)
37. Michieli, U., Zanuttigh, P.: Edge-aware graph matching network for part-based semantic segmentation. *International Journal of Computer Vision* **130**(11), 2797–2821 (2022)
38. Miller, G.A.: Wordnet: a lexical database for english. *Commun. ACM* **38**(11), 39–41 (nov 1995). <https://doi.org/10.1145/219717.219748>, <https://doi.org/10.1145/219717.219748>
39. Mo, K., Zhu, S., Chang, A.X., Yi, L., Tripathi, S., Guibas, L.J., Su, H.: Partnet: A large-scale benchmark for fine-grained and hierarchical part-level 3d object understanding. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 909–918 (2019)
40. Mo, K., Zhu, S., Chang, A.X., Yi, L., Tripathi, S., Guibas, L.J., Su, H.: PartNet: A large-scale benchmark for fine-grained and hierarchical part-level 3D object understanding. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019)
41. Myers-Dean, J., Fan, Y., Price, B., Chan, W., Gurari, D.: Interactive segmentation for diverse gesture types without context. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 7198–7208 (January 2024)
42. Nair, V., Zhu, H.H., Smith, B.A.: Imageassist: Tools for enhancing touchscreen-based image exploration systems for blind and low vision users. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. CHI '23, Association for Computing Machinery, New York, NY, USA* (2023). <https://doi.org/10.1145/3544548.3581302>, <https://doi.org/10.1145/3544548.3581302>
43. Peng, Z., Wang, W., Dong, L., Hao, Y., Huang, S., Ma, S., Wei, F.: Kosmos-2: Grounding multimodal large language models to the world. arXiv preprint arXiv:2306.14824 (2023)

44. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
45. Ramanathan, V., Kalia, A., Petrovic, V., Wen, Y., Zheng, B., Guo, B., Wang, R., Marquez, A., Kovvuri, R., Kadian, A., et al.: Paco: Parts and attributes of common objects. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7141–7151 (2023)
46. Rasheed, H., Maaz, M., Shaji, S., Shaker, A., Khan, S., Cholakkal, H., Anwer, R.M., Xing, E., Yang, M.H., Khan, F.S.: Glamm: Pixel grounding large multimodal model. arXiv preprint arXiv:2311.03356 (2023)
47. RAYMOND W. GIBBS, J., Matlock, T.: Psycholinguistics and mental representations (2000)
48. Ren, Z., Huang, Z., Wei, Y., Zhao, Y., Fu, D., Feng, J., Jin, X.: Pixellm: Pixel reasoning with large multimodal model (2023)
49. Sennrich, R., Haddow, B., Birch, A.: Neural machine translation of rare words with subword units. arXiv preprint arXiv:1508.07909 (2015)
50. Song, X., Wang, P., Zhou, D., Zhu, R., Guan, C., Dai, Y., Su, H., Li, H., Yang, R.: Apollocar3d: A large 3d car instance understanding benchmark for autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5452–5462 (2019)
51. Sun, P., Chen, S., Zhu, C., Xiao, F., Luo, P., Xie, S., Yan, Z.: Going denser with open-vocabulary part segmentation. 2023 IEEE/CVF International Conference on Computer Vision (ICCV) pp. 15407–15419 (2023), <https://api.semanticscholar.org/CorpusID:258762519>
52. Tang, C., Xie, L., Zhang, X., Hu, X., Tian, Q.: Visual recognition by request. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15265–15274 (2023)
53. Tong, K., Wu, Y.: Deep learning-based detection from the perspective of small or tiny objects: A survey. *Image and Vision Computing* **123**, 104471 (2022)
54. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., Lample, G.: Llama: Open and efficient foundation language models (2023)
55. Tsogkas, S., Kokkinos, I., Papandreou, G., Vedaldi, A.: Deep learning for semantic part segmentation with high-level guidance. arXiv preprint arXiv:1505.02438 (2015)
56. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset (2011)
57. Wang, J., Yuille, A.L.: Semantic part segmentation using compositional model combining shape and appearance. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1788–1797 (2015)
58. Wang, P., Shen, X., Lin, Z., Cohen, S., Price, B., Yuille, A.L.: Joint object and part segmentation using deep learned potentials. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1573–1581 (2015)
59. Wang, W., Lv, Q., Yu, W., Hong, W., Qi, J., Wang, Y., Ji, J., Yang, Z., Zhao, L., Song, X., et al.: Cogvlm: Visual expert for pretrained language models. arXiv preprint arXiv:2311.03079 (2023)
60. Wang, X., Li, S., Kallidromitis, K., Kato, Y., Kozuka, K., Darrell, T.: Hierarchical open-vocabulary universal image segmentation. *Advances in Neural Information Processing Systems* **36** (2024)

61. Wei, M., Yue, X., Zhang, W., Kong, S., Liu, X., Pang, J.: Ov-parts: Towards open-vocabulary part segmentation. *Advances in Neural Information Processing Systems* **36** (2024)
62. Wu, T.H., Biamby, G., Chan, D., Dunlap, L., Gupta, R., Wang, X., Gonzalez, J.E., Darrell, T.: See, say, and segment: Teaching lmms to overcome false premises. *arXiv preprint arXiv:2312.08366* (2023)
63. Xiang, F., Qin, Y., Mo, K., Xia, Y., Zhu, H., Liu, F., Liu, M., Jiang, H., Yuan, Y., Wang, H., Yi, L., Chang, A.X., Guibas, L.J., Su, H.: SAPIEN: A simulated part-based interactive environment. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2020)
64. You, H., Zhang, H., Gan, Z., Du, X., Zhang, B., Wang, Z., Cao, L., Chang, S.F., Yang, Y.: Ferret: Refer and ground anything anywhere at any granularity. *arXiv preprint arXiv:2310.07704* (2023)
65. Yu, F., Liu, K., Zhang, Y., Zhu, C., Xu, K.: Partnet: A recursive part decomposition network for fine-grained and hierarchical shape segmentation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9491–9500 (2019)
66. Yuan, Y., Li, W., Liu, J., Tang, D., Luo, X., Qin, C., Zhang, L., Zhu, J.: Osprey: Pixel understanding with visual instruction tuning (2023)
67. Zhao, J., Li, J., Cheng, Y., Sim, T., Yan, S., Feng, J.: Understanding humans in crowded scenes: Deep nested adversarial learning and a new benchmark for multi-human parsing. In: *Proceedings of the 26th ACM international conference on Multimedia*. pp. 792–800 (2018)
68. Zheng, S., Yang, F., Kiapour, M.H., Piramuthu, R.: Modanet: A large-scale street fashion dataset with polygon annotations. In: *Proceedings of the 26th ACM international conference on Multimedia*. pp. 1670–1678 (2018)
69. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 633–641 (2017)