

Visual Recognition by Request

Chufeng Tang¹ Lingxi Xie² Xiaopeng Zhang² Xiaolin Hu^{1,3,4} Qi Tian²

¹ Department of Computer Science and Technology, Tsinghua University

²Huawei Inc. ³Chinese Institute for Brain Research (CIBR)

⁴IDG/McGovern Institute for Brain Research, THBI, Tsinghua University

tcf18@mails.tsinghua.edu.cn {198808xc, zxp@history}@gmail.com

xlhu@mail.tsinghua.edu.cn tian.qil@huawei.com

Abstract

Humans have the ability of recognizing visual semantics in an unlimited granularity, but existing visual recognition algorithms cannot achieve this goal. In this paper, we establish a new paradigm named visual recognition by request (ViRReq¹) to bridge the gap. The key lies in decomposing visual recognition into atomic tasks named requests and leveraging a knowledge base, a hierarchical and text-based dictionary, to assist task definition. ViRReq allows for (i) learning complicated whole-part hierarchies from highly incomplete annotations and (ii) inserting new concepts with minimal efforts. We also establish a solid baseline by integrating language-driven recognition into recent semantic and instance segmentation methods, and demonstrate its flexible recognition ability on CPP and ADE20K, two datasets with hierarchical whole-part annotations.

1. Introduction

Visual recognition is one of the fundamental problems in computer vision. In the past decade, visual recognition algorithms have been largely advanced with the availability of large-scale datasets and deep neural networks [8, 13, 18]. Typical examples include the ability of recognizing 10,000s of object classes [6], segmenting objects into parts or even parts of parts [43], using natural language to refer to open-world semantic concepts [31], etc.

Despite the increasing recognition accuracy in standard benchmarks, we raise a critical yet mostly uncovered issue, namely, the **unlimited granularity** of visual recognition. As shown in Figure 1, given an image, humans have the ability of recognizing arbitrarily fine-grained contents from it, but existing visual recognition algorithms cannot achieve the same goal. Superficially, this is caused by limited annotation budgets so that few training data is available for fine-

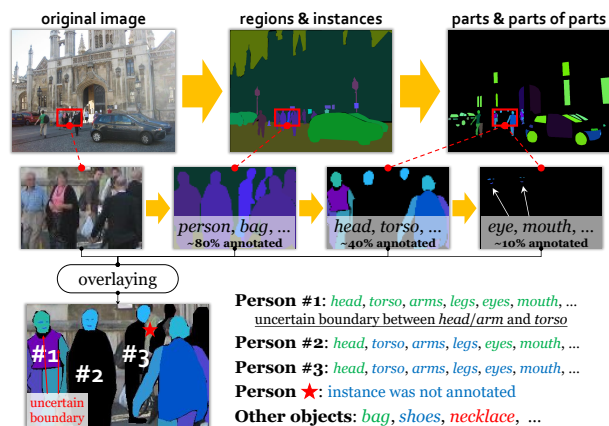


Figure 1. An illustration of unlimited granularity in visual recognition. **Top**: an example image from ADE20K [43] with instance-level and part-level annotations. **Middle**: more and more incomplete annotations of instances, parts, and parts of parts. **Bottom**: as granularity goes finer, higher uncertainty occurs in recognizing the boundary (left) and semantic class (right) of objects and/or parts – here, **green**, **blue** and **red** texts indicate **labeled**, **unlabeled (but defined)**, and **unlabeled (and undefined)** objects, respectively.

grained and/or long-tailed concepts, but we point out that a more essential reason lies in the conflict between granularity and recognition certainty – as shown in Figure 1, *when annotation granularity goes finer, annotation certainty will inevitably be lower*. This motivates us that the granularity of visual recognition shall be variable across instances and scenarios. For this purpose, we propose to interpret semantic annotations into smaller units (*i.e.*, requests) and assume that recognition is performed only when it is asked, so that one can freely adjust the granularity according to object size, importance, clearness, etc.

In this paper, we establish a new paradigm named **visual recognition by request (ViRReq)**. We only consider the segmentation task in this paper because it best fits the need of unlimited granularity in the spatial domain. Compared to

¹We recommend the readers to pronounce ViRReq as /'virik/.

the conventional setting, the core idea is to decompose visual recognition into a set of atomic tasks named **requests**, each of which performs one step of recognition. Specifically, there are two request types, with the first type decomposing an instance into semantic parts and the second type segmenting an instance out from a semantic region. For the first type, a **knowledge base** is available as a hierarchical text-based dictionary to guide the segmentation procedure.

Compared to the existing paradigms of visual recognition, ViRReq has two clear advantages. First, ViRReq naturally has the ability of learning complicated visual concepts (*e.g.*, the whole-part hierarchy in ADE20K [43]) from highly incomplete annotations, while conventional methods can encounter several difficulties. Second, ViRReq allows a new visual concept (*e.g.*, objects, parts, *etc.*) to be added by simply updating the knowledge base and annotating a few training examples. We emphasize that the change of knowledge base does not impact the use of existing training data as each sample is bound to a specific version of the knowledge base. This property, called data versioning, allows for incremental learning with all historical data available.

To deal with ViRReq, we build a query-based recognition algorithm that (i) extracts visual features from the image, (ii) computes text embedding vectors from the request and knowledge base, and (iii) performs interaction between image and text features. The framework is highly modular and the main parts (*e.g.*, feature extractors) can be freely replaced. We evaluate ViRReq with the proposed recognition algorithm on two datasets, namely, the CPP dataset [5] that extends Cityscapes [3] with part-level annotations and the ADE20K dataset [43] that offers a multi-level hierarchy of complicated visual concepts. We parse the annotations of each image into a set of requests for training, and define a new evaluation metric named hierarchical panoptic quality (HPQ) for measuring the segmentation accuracy. Thanks to the ability of learning from incomplete annotations, ViRReq can report part-aware segmentation accuracy on ADE20K, which, to the best of our knowledge, is the first ever work to achieve the goal. In addition, ViRReq shows a promising ability of open-domain recognition, including absorbing new visual concepts (*e.g.*, objects, parts, *etc.*) from a few training samples and understanding anomalous or compositional concepts without training data.

In summary, the contribution of this paper is three fold: (i) pointing out the issue of unlimited granularity, (ii) establishing the paradigm named visual recognition by request, and (iii) setting up a solid baseline for this new direction.

2. Preliminaries and Insights

2.1. Related Work in Visual Recognition

In the deep learning era [19], with deep neural networks [8, 13, 18, 27] being adopted as a generalized tool for

representation learning, the community has been pursuing for an effective method to improve the *granularity* of visual recognition, which we refer to the ability of recognizing rich visual contents. From the perspective of defining a fine-grained visual recognition task, typical examples include collecting image data for more semantic classes (*e.g.*, ImageNet [6] has more than 22K classes) and labeling finer parts (*e.g.*, ADE20K [43] labeled more than 600 classes of parts and parts of parts). Despite the great efforts, these high-quality datasets are still far from the goal of *unlimited granularity* of visual recognition, in particular, recognizing everything that humans can recognize [40]. In what follows, we categorize existing recognition approaches into two types and analyze their drawbacks *assuming the goal being unlimited granularity*.

The first type is the **classification-based** tasks which refer to the targets by assigning a class ID for each semantic unit. The definition of semantic units determines the visual recognition task, *e.g.*, the unit is an image for image classification [6], a bounding box for object detection [9, 25, 33], a masked region for image segmentation [3, 9, 25, 43], a key-point for human pose estimation [1, 25], *etc.* This mechanism has a critical drawback, *i.e.*, the conflict between granularity and certainty. As shown in Figure 1, the certainty (also accuracy) of annotation is not guaranteed when either semantic or spatial granularity goes finer. To alleviate the conflict, we shall allow the granularity to vary across semantic units, *e.g.*, large objects are labeled with parts but small objects are not. *In this paper, we propose a framework that annotations are made upon request.*

The second type is the **language-driven** tasks which leverage texts or other modalities to refer to specific visual contents. Typical examples include visual question answering [2, 4, 11, 38], image captioning [14, 37], visual grounding [15, 26, 28, 45], visual reasoning [42], *etc.* Recently, with the availability of vision-language pre-trained models that learn knowledge from image-text pairs (*e.g.*, CLIP [31], GLIP [22], *etc.*), this mechanism has shown promising abilities for open-world recognition [16, 29, 41], especially in using language as queries for visual recognition [7, 20, 44]. However, the ability of natural language of referring to detailed visual contents (*e.g.*, the parts of a specific **person** in a complex scene with tens of **persons**) is very much limited, hence, it is unlikely that purely relying on language can achieve unlimited granularity. *In this paper, we use language to define a hierarchical dictionary but establish the benchmark mainly based on vision itself.*

2.2. Key Insights towards Unlimited Granularity

Based on the above analysis, we have learned that unlimited granularity is not yet achieved by existing paradigms. We summarize two key insights towards the goal, based on which we present our solution in the next section.

First, we note that unlimited granularity includes the scope of openness (*i.e.*, the open-domain property) and is more challenging. Leveraging data from other modalities (*e.g.*, natural language) is the most promising and convenient way of introducing openness, so we use texts as labels to define semantic concepts. This strategy makes it easier to capture the relationship between concepts (*e.g.*, different kinds of vehicles such as car and bus are closely related) and to learn compositional concepts (*e.g.*, transferring the visual features from the parts of car to the parts of bus).

Second and more importantly, we assume that the recognition granularity is variable across instances and/or scenarios. This calls for a flexible recognition paradigm that *not always* pursues for the finest granularity but performs recognition only with proper requests (*e.g.*, a person can be annotated with parts if the resolution is sufficiently large), yet the instance can be ignored if the resolution is very small. *In our proposal, the granularity is freely controlled by decomposing the recognition task into requests.*

3. Problem: Visual Recognition by Request

As shown in Figure 2, ViRReq offers a novel and unified paradigm for both data annotation and visual recognition. Throughout this paper, we parse existing datasets (CPP and ADE20K) into the ViRReq setting and present a solid baseline. We advocate that the community annotates and organizes visual data using this paradigm, so as to push visual recognition towards unlimited granularity.

3.1. Notations and Task Definition

We consider the image segmentation task in this paper because it mostly fits our goal towards unlimited granularity. Conventional segmentation benchmarks often provide a pre-defined dictionary that contains all visual concepts (*i.e.*, classes). Given an image, $\mathbf{X} = \{x_{w,h}\}_{w=1,h=1}^{W,H}$, where $x_{w,h}$ indicates the pixel at the position (w, h) and W, H are the image width and height, respectively, it is required to predict a pixel-wise segmentation map, $\mathbf{Z} = \{z_{w,h}\}_{w=1,h=1}^{W,H}$, where $z_{w,h}$ represents the semantic and/or instance labels of $x_{w,h}$. As we have analyzed above, such a definition can encounter difficulties when the scene becomes complex and the granularity becomes finer.

The core idea of ViRReq is to decompose the recognition task, $\mathbf{X} \mapsto \mathbf{Z}$, into a series of requests, denoted as $\mathcal{R} = \{\mathbf{r}_k\}_{k=1}^K \doteq \{Q_k, \mathcal{A}_k\}_{k=1}^K$, where Q_k and \mathcal{A}_k denote the k -th request and answer, respectively. These requests are performed sequentially, and some of them may depend on the recognition results of earlier requests, *e.g.*, the system must first segment an instance before further partitioning it into parts. The segmentation results are stored in a tree, denoted as $\mathcal{T} = \{\mathbf{t}_l\}_{l=0}^L \doteq \{\mathbf{M}_l, c_l, \tau_l, \mathcal{U}_l\}_{l=0}^L$, where each node is a semantic region or an instance. For the l -th node, $\mathbf{M}_l \in \{0, 1\}^{W \times H}$ is the binary mask, $c_l \in \{1, 2, \dots, C\}$

is the class index in the knowledge base (see the next paragraph), $\tau_l \in \{0, 1\}$ indicates whether the node corresponds to an instance, and $\mathcal{U}_l \subset \{l+1, \dots, L\}$ is the set of child nodes of \mathbf{t}_l . If \mathbf{t}_1 is the parent node of \mathbf{t}_2 , then \mathbf{t}_2 is recognized by answering a request on \mathbf{t}_1 . In our definition, an instance must be a child node of a semantic region. Note that \mathcal{R} and \mathcal{T} are tightly related: each non-leaf node \mathbf{t}_l in \mathcal{T} corresponds to a request in \mathcal{R} , say \mathbf{r}_k , and the answer \mathcal{A}_k corresponds to the set of all child nodes of \mathbf{t}_l .

Based on the above definition, a typical procedure of visual recognition is described as follows and illustrated in Figure 2. Initially, nothing is annotated and \mathcal{T} has only a root node \mathbf{t}_0 , in which $\mathbf{M}_0 = \mathbf{1}^{W \times H}$, $c_0 = 1$ (indicating a special class named **scene**), $\tau_0 = 1$ (*i.e.*, this is an instance of **scene**), and $\mathcal{U}_0 = \emptyset$. Each request, \mathbf{r}_k , finds the corresponding node, \mathbf{t}_l , performs the recognition task, translates the answer into child nodes of \mathbf{t}_l , and adds them back to \mathcal{T} . There are two types of requests:

Type-I Whole-to-part semantic segmentation. Prerequisite: \mathbf{t}_l must be an instance (*i.e.*, $\tau_l = 1$). The system fetches the class label, c_l , and looks up the dictionary for the set of its sub-classes in texts (*e.g.*, **person** has sub-classes (parts) of **head**, **torso**, **arm**, and **leg**, *etc.*).

Type-II Instance segmentation. Prerequisite: \mathbf{t}_l must be a semantic region (*i.e.*, $\tau_l = 0$). Each time, a probing pixel (named probe) (a_k, b_k) is provided with the request, $M_l(a_k, b_k) = 1$, and the task is to segment the instance that occupies this probe.

At the beginning, the entire image is regarded as a special instance named **scene**, and the top-level semantic classes (*e.g.*, **sky**, **building**, **person**, *etc.*) are considered as its parts. We use such definition to simplify the overall notation system. Briefly, Type-I requests segment instances into semantics, while Type-II requests find instances from semantics. By executing them iteratively and alternately (*i.e.*, the path that yields one unit must be Type-I→Type-II→Type-I→...), one can segment arbitrarily fine-grained units (*i.e.*, towards unlimited granularity) from the input image.

To define the whole-part hierarchy for Type-I requests, we establish a directed graph known as the **knowledge base** – a toy example is shown in Figure 2. Each node of the graph is an object/part concept, and it may have a few child nodes that correspond to its named parts. In our formulation, the class labels appear as texts rather than a fixed integer ID – this is to ease the generalization towards open-domain recognition: given a pre-trained text embedding (*e.g.*, CLIP [31]), some classes are recognizable by language even if they never appear in the training set. The root class is **scene** that corresponds to the entire picture.

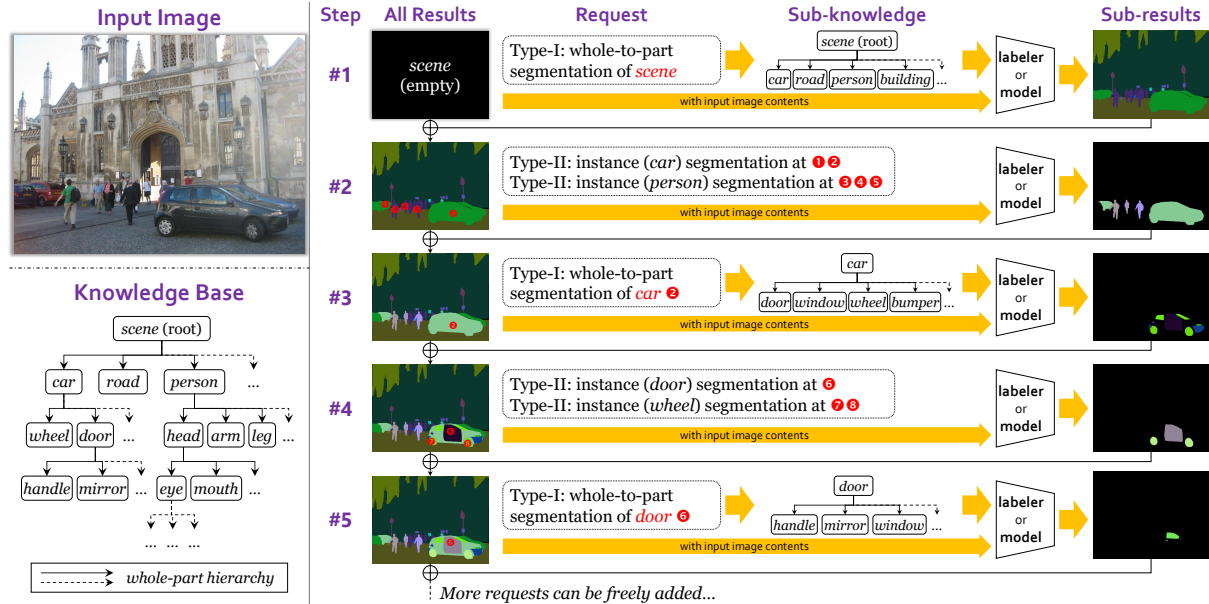


Figure 2. An illustration of the overall pipeline of ViRReq that is built upon a pre-defined knowledge base. Given an input image, ViRReq starts with an empty field (*i.e.*, nothing is annotated or recognized) and iteratively executes Type-I/II requests towards recognition results of finer granularity. The sub-knowledge for Type-I requests is extracted from the knowledge base. The labeler can stop after any number of steps, yet the model can learn from incomplete annotations without difficulties. To save space, we combine a few Type-II requests into one step. We only display five steps here, while a typical annotation/recognition procedure often include much more steps (and requests).

3.2. Advantages over Existing Paradigms

ViRReq is different from existing visual recognition paradigms, in particular, language-driven visual recognition mentioned above. We decompose visual recognition into requests, guided by the knowledge base, to pursue for the ultimate goal, *i.e.*, visual recognition of unlimited granularity. This brings the following two concrete advantages.

First, from a **micro** view, ViRReq allows us to learn complex whole-part hierarchies from highly incomplete annotations. As shown in Figure 2, the parsed requests do not deliver noisy supervision even though the training data (i) ignores many instances for dense objects and/or (ii) annotates fine-grained parts only for a small amount of objects. More importantly, it alleviates the conflict between granularity and certainty by only annotating the contents that labelers are sure about. In other words, ViRReq boosts certainty by sacrificing completeness for a single image, but the entire dataset covers the knowledge base sufficiently well.

Second, from a **macro** view, ViRReq allows us to add new concepts (objects and/or parts) easily. To do this, one only needs to (i) insert a text-based node to the knowledge base and (ii) annotate a few images that contains the concept. Note that although the knowledge base augments with time, this does not prevent us from using old training data, because we can associate each training image to the knowledge base that was used for annotating new concepts (which we call data versioning). That said, the new concept can re-

main unlabeled in old images even it appears, as long as the old images are associated to the old knowledge base that does not contain the new concept. In Section 5.3, we will show the benefits of data versioning in the incremental learning experiments in ADE20K.

4. Query-Based Recognition: A Baseline

ViRReq calls for an algorithm that deals with requests $\{\mathbf{r}_k\}_{k=1}^K$ one by one, similar to the illustration in Figure 2. Specifically, in each step, the input involves the image \mathbf{X} , the current recognition result \mathbf{Z}_k (prior to the k -th request), the request Q_k , and the knowledge base. Processing each request involves extracting visual features, constructing queries, performing recognition and filtering, and updating the current recognition results.

Visual features. We extract visual features from \mathbf{X} using a deep neural network (*e.g.*, a conventional convolutional neural network [13, 18] or vision transformer [8, 27]), obtaining a feature map $\mathbf{F} \in \mathbb{R}^{H' \times W' \times D}$, where $H' \times W'$ is usually a down-sampled scale of the original image.

Language-based queries. Each request is transferred into a set of query embedding vectors \mathbf{E} , for which a pre-trained text encoder (*e.g.*, CLIP [31]) and the knowledge base are required. (1) For Type-I requests (*i.e.*, whole-to-part semantic segmentation), the target class (in text, *e.g.*, **PERSON**) is used to look up the knowledge base, and a total of P_k child nodes are found (in text, *e.g.*, head, arm, leg,

torso). We feed them into the text encoder to obtain P_k embedding vectors, $\mathbf{E} = [\mathbf{e}_{k,1}; \dots; \mathbf{e}_{k,P_k}]$. (2) For Type-II requests (*i.e.*, instance segmentation), given a probing pixel, a triplet $\{a_k, b_k, c_k\}$ is obtained, where (a_k, b_k) is the pixel coordinates and c_k is the target class that is determined by the current recognition result. To construct the query, the target class (in text) is directly fed into the text encoder to obtain the semantic embedding vector \mathbf{e}_k , and the pixel coordinates are fed into a positional encoder to obtain the positional embedding \mathbf{p}_k , where $\mathbf{p}_k = (a_k, b_k)$ for simplicity in our implementation. The query embedding is obtained by combining them, $\mathbf{E} = (\mathbf{e}_k, \mathbf{p}_k)$. All text embedding vectors are D -dimensional, *i.e.*, same as the visual features.

Language-driven recognition. Visual features \mathbf{F} are interacted with the language-based queries \mathbf{E} to perform segmentation. (1) For Type-I requests, we directly compute the inner-product between the visual feature vector $\mathbf{f}_{w,h} \in \mathbb{R}^D$ of each pixel and the embedding vectors \mathbf{E} , obtaining a P_k -dimensional class score vector, $\mathbf{u}_{w,h} = \mathbf{E}^\top \cdot \mathbf{f}_{w,h}$. The entry with the maximal response is the predicted class label. To allow open-set recognition, we augment $\mathbf{u}_{w,h}$ with an additional entry of 0, *i.e.*, $\hat{\mathbf{u}}_{w,h} = [\mathbf{u}_{w,h}; 0]$, where the added entry stands for the *others* class – that said, if the responses of all P_k normal entries are smaller than 0, the corresponding position is considered an unseen (anomalous) class – see such examples in Section 5.3. (2) For Type-II requests, different from semantic segmentation, existing instance segmentation methods usually generate massive proposals (*e.g.*, box proposals in Mask R-CNN [12], feature locations in CondInst [35]) and predict the class label, bounding box, and binary mask for each proposal. To achieve instance segmentation *by request*, we first select the spatially related proposals by the positional embedding \mathbf{p}_k . Specifically, feature locations (in CondInst) near the probing pixel at each feature pyramid level are selected for subsequent prediction. Prediction with the highest categorical score, obtained as similar as processing Type-I requests (inner-product with text embedding), is chosen as the final result.

Top-down filtering. We combine the request r_k and the current recognition result \mathbf{Z}_k to constrain the segmentation masks predicted above. That said, we follow the top-down criterion to deal with the segmentation conflict of different levels requests. For example, part segmentation results should strictly inside the corresponding instance region. Applying advanced mask fusion methods [21, 24, 30, 32] may produce better results. The filtered masks are added back to the current recognition results, which lays the foundation of subsequent requests.

Implementation and training details. We implement the method with two individual models, and use a pre-trained and frozen CLIP [31] model to generate text embeddings. (1) For Type-I requests, we adopt a prevailing Transformer-based model, SegFormer [39] with its customized backbone




Figure 3. An illustration of how Type-I requests are processed for *one pixel*. Each slot indicates the predicted logit (normalized inner-product) of one class. The *selector* is used to look up the knowledge base for child nodes of current class (decided by ground-truth and upper-level prediction during training and testing, respectively). Only labeled classes are used for computing loss and different losses are added together. In the testing stage, dashed lines are activated and the ground-truth is not accessible.

MiT-B0/B5, to extract pixel-wise visual features. For efficient training, requests of all object and part classes are processed *in parallel*. For each visual feature vector, we compute the inner-product between it and all text embeddings. Then, the softmax with cross-entropy loss is computed. Note that only the sub-classes that appear in the knowledge base and are also labeled in the current instance are considered – other classes are ignored because it is unclear if they appear but are not annotated. A pixel may have multiple labels corresponding to different semantic levels, and thus multiple loss terms may be computed, as shown in Figure 3. This strategy is crucial for learning from highly incomplete annotations (see Section 5.3). (2) For Type-II requests, we adopt CondInst [35] where all positive feature locations are viewed as probes and trained *in parallel* for each image. Note that in ViRReq, the probes are allowed to lie anywhere on the instance, so we modify CondInst to sample positive training locations from the entire instance mask, instead of the central region [36] of instance. This training trick, named *mask sampling*, improves instance segmentation accuracy in our setting. For further implementation details, please refer to Appendix A.

5. Experiments

5.1. Datasets and Metric

We investigate ViRReq on two segmentation datasets with whole-part hierarchies. (1) The Cityscapes-Panoptic-Parts (CPP) dataset [5] inherited the images and 19 object classes from Cityscapes [3], and further annotated two groups of part classes for vehicles and persons. As a result, there are 5 object classes having parts, with 9 non-duplicate part classes in total. There is a high ratio of instances annotated with parts, so it is relatively easy for the conventional approaches to make use of this dataset. (2) The ADE20K dataset [43] has a much larger number of

Table 1. Individual segmentation accuracy on CPP. **Left:** Semantic segmentation (Type-I) on Level-1 (*i.e.*, scene-level, *e.g.*, car) and Level-2 (*i.e.*, part-level, *e.g.*, wheel) classes. **Right:** Instance segmentation (Type-II) results on all instances under the non-probing protocol. * indicates that BPR [34] is used.

mIoU (%)	Lv-1	Lv-2	AP (%)	Inst.
SegFormer (B0) [39]	76.54	–	CondInst [35] (R50)	36.6
w/ CLIP	77.35	–	w/ CLIP	36.8
w/ CLIP & parts (ours)	77.39	75.42	w/ CLIP & mask samp.	37.8
SegFormer (B5) [39]	82.25	–	CondInst [35] (R50)*	39.2
w/ CLIP	82.40	–	w/ CLIP	39.6
w/ CLIP & parts (ours)	82.33	78.48	w/ CLIP & mask samp.	40.5

object and part classes. We follow the convention to use the 150 most frequent scene classes for top-level semantic segmentation, and 100 countable objects out of them are used for instance segmentation. Besides, for finer-grained whole-to-part segmentation, we select 82 part classes (non-duplicate) belonging to 40 instance classes, and 29 part-of-parts classes belonging to 17 upper-level part classes. Please refer to Appendix C.1 for the details of data preparation. ADE20K has highly sparse and incomplete annotations of parts (and parts of parts), raising difficulties for training and evaluation – to the best of our knowledge, conventional approaches have not yet quantitatively reported part-based segmentation results on this dataset.

To measure the segmentation quality, we design a metric named hierarchical panoptic quality (HPQ) which can measure the accuracy of a recognition tree of *any depth*. HPQ is a recursive metric. For a leaf node, HPQ equals to mask IoU, otherwise, we first compute the class-wise HPQ:

$$\text{HPQ}_c(\mathbf{t}_l) = \frac{\sum_{\mathbf{t}_{l'} \in \mathcal{TP}_c \cap \mathcal{U}_l} \text{HPQ}(\mathbf{t}_{l'})}{|\mathcal{TP}_c \cap \mathcal{U}_l| + \frac{1}{2} |\mathcal{FP}_c \cap \mathcal{U}_l| + \frac{1}{2} |\mathcal{FN}_c \cap \mathcal{U}_l|}, \quad (1)$$

where \mathcal{TP}_c , \mathcal{FP}_c , and \mathcal{FN}_c denote the sets of true-positives, false-positives, and false-negatives of the c -th class, respectively. The true-positive criterion is $\text{HPQ}_c(\mathbf{t}_{l'}) \geq 0.5$. The HPQ values of all active classes $\text{HPQ}_c(\mathbf{t}_0)$ are averaged into $\text{HPQ}(\mathbf{t}_0)$ at the root node. HPQ is related to prior metrics, *e.g.*, it degenerates to the original PQ [17] if there is no object-part hierarchy, and gets similar to PartPQ [5] if the object-part relationship is one-level (*i.e.*, parts cannot have sub-parts) and parts are only semantically labeled (*i.e.*, no part instances) – this is the case of the CPP dataset (see Appendix B.2). Yet, HPQ is easily generalized to more complex hierarchies.

5.2. Results on CPP

Individual tests. In CPP, a dataset with relatively complete part annotations, we evaluate Type-I and Type-II requests individually. Results are shown in Table 1. For Type-I requests (*i.e.*, semantic segmentation), introducing

Table 2. PartPQ on all, NP (*i.e.*, without parts), and P (*i.e.*, with parts) classes on CPP. * indicates that instance segmentation is improved by BPR [34]. We also show HPQ values of our method but they cannot be compared against prior methods.

PartPQ (%)	All	NP	P
UPSNNet + DeepLabv3+ [5]	55.1	59.7	42.3
HRNet-OCR + PolyTransform + BSAntNet [5]	61.4	67.0	45.8
Panoptic-PartFormer (ResNet50) [23]	57.4	62.2	43.9
Panoptic-PartFormer (Swin-base) [23]	61.9	68.0	45.6
ViRReq: SegFormer (B0) + CondInst (R50)	57.1	61.6	44.2
<i>evaluated using HPQ (%)</i>	<i>56.0</i>	<i>60.5</i>	<i>43.3</i>
ViRReq: SegFormer (B5) + CondInst (R50)*	62.5	67.5	48.6
<i>evaluated using HPQ (%)</i>	<i>61.6</i>	<i>66.6</i>	<i>47.4</i>

text embedding as flexible classes improves the accuracy slightly, and training together with part-level classes also produces reasonable results. That said, language-based segmentation models can handle non-part and part classes simultaneously. For Type-II requests (*i.e.*, instance segmentation), to fairly compete with other methods, we do not use hand-crafted probes in the testing stage. Hence, we densely sample probes on the semantic segmentation regions, each of which generating a candidate instance, and filter the results using NMS (see Appendix A.3 for details). This is called the non-probing-based setting, which will be used throughout this section, and we also study a more flexible probing-based setting (requires user interaction) in Appendix A.4. As shown in Table 1, we achieve higher instance segmentation accuracy. A useful training trick is mask sampling, which makes the model insensitive to the position of probes, as elaborated in Appendix A.4.

Comparison to recent approaches. We combine the best practice into overall segmentation. Results are shown in Table 2. We compute both the PartPQ and HPQ metrics and use PartPQ to compare against previous approaches. Note that computing PartPQ is only possible on two-level datasets such as CPP. As shown, our method reports competitive accuracy among recent works [5, 23]. In addition, when we evaluate NP (without defined parts) and P (with parts) classes separately, we find that ViRReq has clear advantages in the latter set, validating the benefit of step-wise visual recognition, *i.e.*, the model can recognizing objects first and then parts. *Note that this work is actually not optimized for higher results on the fully-annotated CPP dataset, the most important advantages are the abilities to learn complex hierarchies from incomplete annotations and adapt to new concepts easily, as investigated below.*

Learning from incomplete annotations. We have discussed how ViRReq deals with incomplete annotations in Section 3.2 and Section 4. Here, we study this mechanism by performing two experiments on CPP. (1) We preserve

Table 3. Segmentation accuracy of using different ratios of annotations. **Left:** Setting (1) in which all semantic/instance and a portion of part annotations are preserved. **Right:** Setting (2) in which both semantic and part annotations are preserved to the corresponding ratio. See the main texts for details.

Ratio	mIoU (%)		HPQ	Ratio	mIoU (%)		HPQ
	Lv-1	Lv-2	(%)		Lv-1	Lv-2	(%)
100%	82.33	78.48	60.2	100%	82.33	78.48	60.2
50%	81.82	77.59	59.8	75%	81.49	77.81	59.3
30%	81.99	76.52	59.4	50%	79.01	76.23	56.9
15%	81.72	65.51	57.3	30%	73.53	69.72	51.7

all semantic and instance annotations but randomly choose a certain portion of part annotations. (2) Beyond (1), we further remove part of semantic annotations – note that this requires associating each image with an individual knowledge base. Please refer to Appendix B.3 for details. Results are shown in Table 3. We find that ViRReq, without any modification, adapts to both scenarios easily. This property largely benefits our investigation on the ADE20K dataset, where a considerable portion of part annotations is missing and there is an eager call for few-shot incremental learning.

5.3. Results on ADE20K

Adapting conventional methods to ADE20K. We first note that the annotations on ADE20K are highly incomplete, making it difficult for the conventional segmentation methods to learn the complex, hierarchical concepts. We conjecture it to be the main reason that no prior works have ever reported quantitative results for part-aware segmentation on ADE20K. Conventional solutions usually adopt two ways to handle unlabeled pixels, either ignoring all of them or assigning them to an extra background class. We argue that the latter solution is prone to introducing erroneous information because it cannot distinguish undefined parts from unlabeled parts. For comparison, we establish a competitor method (which we refer to as *convention*) in which all three levels (*i.e.*, objects, object parts, parts of parts) are jointly trained in a multi-task manner, following the former conventional solution to ignore all unlabeled pixels. Meanwhile, ViRReq is easily adapted to this scenario. We simply inherit the best practice from the CPP experiments. The segmentation models (SegFormer and CondInst) and training epochs remain the same for fair comparison. That said, the main difference between ViRReq and the competitor lies in the way of organizing training data to ease learning multi-level hierarchies from incomplete annotations.

Quantitative comparison. Results are shown in Table 4. We report the HPQ values for different subsets of concepts. Without complex adaptation, ViRReq surpasses the convention, showing the ability of ViRReq in learning from highly incomplete annotations. We expect the improvement to be

Table 4. Comparison of HPQ on all, NP (*i.e.*, without parts), and P (*i.e.*, with parts defined) classes on ADE20K. Both methods use CondInst (R50) for instance segmentation. P[†] indicates that instances without part annotations are ignored, so that the metric focuses more on the quality of part-level segmentation.

HPQ (%)		All	NP	P	P [†]
Convention	SegFormer-B0	25.6	24.9	27.7	16.8
ViRReq		27.2	25.7	31.1	21.9
Convention	SegFormer-B5	31.8	31.8	32.0	18.9
ViRReq		33.9	32.8	36.9	27.1

larger when more and more parts or parts of parts are introduced – in such scenarios, ViRReq enjoys the ability of decomposing recognition into requests which gets rid of the burden of distinguishing from a large number of part classes. ViRReq offers a benchmark in this direction and offers room of improvement for future research.

Qualitative comparison. We visualize some segmentation results of ViRReq in Figure 4 to show its ability of recognizing complex whole-part hierarchies. Moreover, we compare ViRReq to the competitor in learning from incomplete part annotations. As shown in Figure 5, ViRReq produces much cleaner segmentation results in which most undefined pixels are correctly recognized as *others*.

Few-shot incremental learning. We add 50 top-level objects and 19 parts to the original knowledge base built in ADE20K. For each new concept, 20 instances are labeled, yet the existing data may contain unlabeled instances of these classes. We mix old and new data and fine-tune the trained SegFormer-B5 for 1/4 of base training iterations. We follow CLIP-Adapter [10] to integrate an additional bottleneck layer for feature alignment, which is the only layer that gets updated during the fine-tuning. More details on data preparation and implementation are provided in Appendix C.5. We show qualitative results in Figure 6, where ViRReq easily absorbs new classes into the knowledge base meanwhile the base classes remain mostly unaffected. We emphasize that the incremental learning ability comes from (i) the query-based, language-driven recognition framework and (ii) data versioning (see Section 3.2).

Open-domain recognition. We investigate two interesting scenarios of open-domain recognition. The first one is anomaly segmentation, *i.e.*, finding unknown concepts in an image, which is naturally supported in ViRReq due to the existence of the *others* class in each sub-task. Object-level and part-level examples are shown in Figures 5 and 6, respectively. Note that ViRReq can easily convert unknown classes to known classes via incremental learning. The second one involves compositional segmentation, *i.e.*, transferring part-level knowledge from one class (*e.g.*, *car*) to *others* (*e.g.*, *bus* or *van*) without annotating new data but directly copying the sub-knowledge from the old class to new

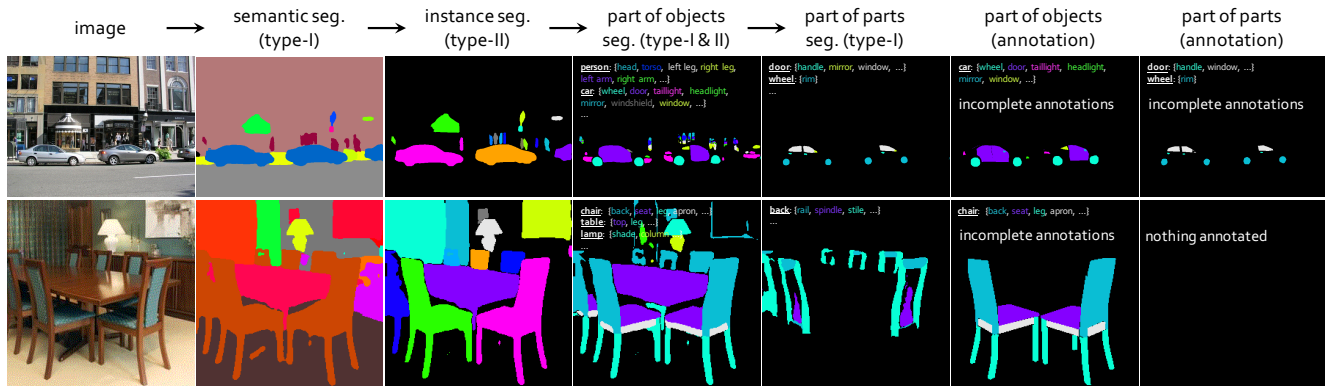


Figure 4. Visualization results of ViRReq on ADE20K. Black areas in prediction indicate the others (*i.e.*, unknown) class. The corresponding sub-knowledge is listed in the blank area for reference. *Best viewed in color and by zooming in for details.*

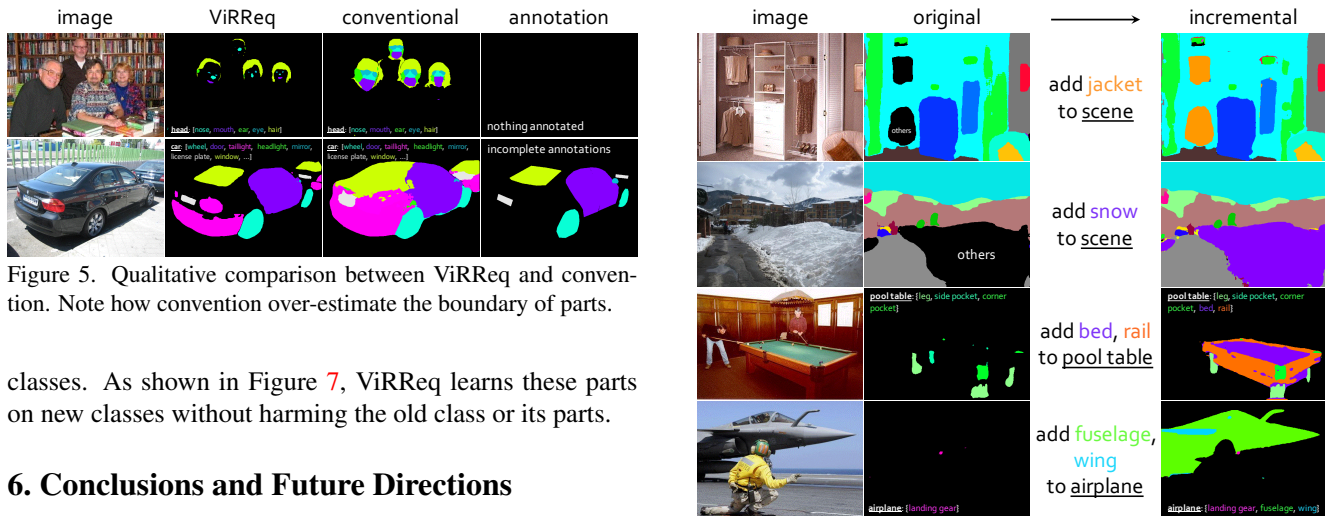


Figure 5. Qualitative comparison between ViRReq and convention. Note how convention over-estimate the boundary of parts.

classes. As shown in Figure 7, ViRReq learns these parts on new classes without harming the old class or its parts.

6. Conclusions and Future Directions

In this paper, we present ViRReq, a novel setting that pushes visual recognition towards unlimited granularity. The core idea is to decompose the end-to-end recognition task into requests, each of which performs a single step of recognition, hence alleviating the conflict between granularity and certainty. We offer a simple baseline that uses language as class index, and achieves satisfying segmentation accuracy on two segmentation datasets with multiple whole-part hierarchies. In particular, thanks to the ability of learning from incomplete annotations, we report part-aware segmentation accuracy (using HPQ, a newly defined metric for ViRReq) on ADE20K for the first time.

From this preliminary work, we learn the lesson that visual recognition has some important yet unsolved issues (*e.g.*, unlimited granularity). Language can assist the definition (*e.g.*, using text-based class indices for open-domain recognition), but solving the essential difficulty requires insights from vision itself. From another perspective, the baseline can be seen as pre-training a language-assisted vision backbone and querying it using vision-friendly prompts (*i.e.*, Type-I/II requests). In the future, we will continue with this paradigm to explore the possibility of unifying various visual recognition tasks, for which two promising directions emerge: (i) designing an automatic

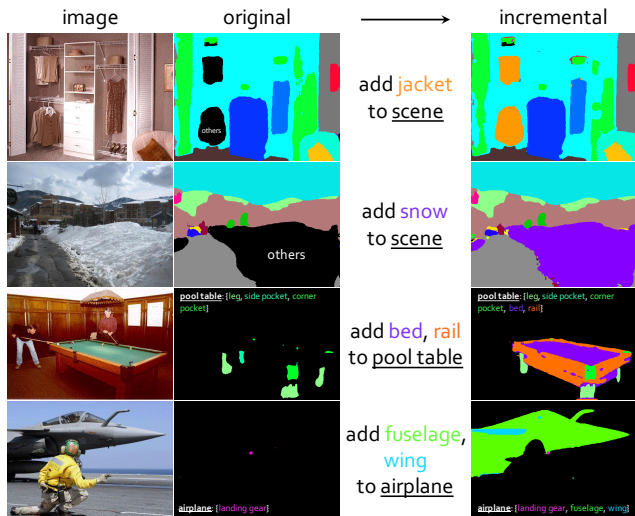


Figure 6. Qualitative results of few-shot incremental learning on objects (top two rows), and object parts (bottom two rows). Prior to incremental learning, these new concepts are recognized as others (the black areas) which is as expected.

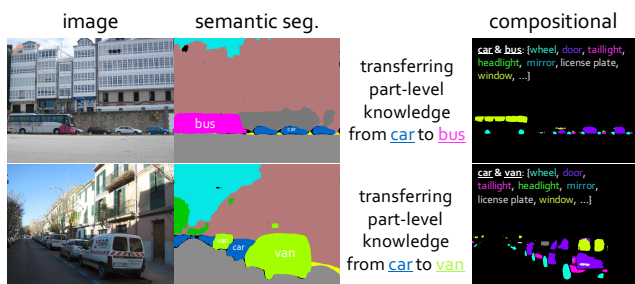


Figure 7. Visualization of compositional segmentation. Note that the base class (car) and its parts are unaffected.

method for learning and updating the knowledge base from training data, and (ii) closing the gap between upstream pre-training and downstream fine-tuning with better prompts.

Acknowledgements. This work was supported by the National Key Research and Development Program of China (Grant 2021ZD0200301), the National Natural Science Foundation of China (Grants 62061136001 and 61836014).

References

- [1] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *Computer Vision and Pattern Recognition*, 2014. 2
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *International Conference on Computer Vision*, 2015. 2
- [3] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Computer Vision and Pattern Recognition*, 2016. 2, 5, 13
- [4] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied question answering. In *Computer Vision and Pattern Recognition*, 2018. 2
- [5] Daan de Geus, Panagiotis Meletis, Chenyang Lu, Xiaoxiao Wen, and Gijs Dubbelman. Part-aware panoptic segmentation. In *Computer Vision and Pattern Recognition*, 2021. 2, 5, 6, 13, 14
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition*, 2009. 1, 2
- [7] Zheng Ding, Jieke Wang, and Zhuowen Tu. Open-vocabulary panoptic segmentation with maskclip. *arXiv preprint arXiv:2208.08984*, 2022. 2
- [8] 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. In *International Conference on Learning Representations*, 2021. 1, 2, 4
- [9] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010. 2
- [10] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544*, 2021. 7, 18
- [11] Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, and Ali Farhadi. Iqa: Visual question answering in interactive environments. In *Computer Vision and Pattern Recognition*, 2018. 2
- [12] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *International Conference on Computer Vision*, 2017. 5, 11
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition*, 2016. 1, 2, 4
- [14] MD Zakir Hossain, Ferdous Sohel, Mohd Fairuz Shiratuddin, and Hamid Laga. A comprehensive survey of deep learning for image captioning. *ACM Computing Surveys*, 51(6):1–36, 2019. 2
- [15] Ronghang Hu, Huazhe Xu, Marcus Rohrbach, Jiashi Feng, Kate Saenko, and Trevor Darrell. Natural language object retrieval. In *Computer Vision and Pattern Recognition*, 2016. 2
- [16] KJ Joseph, Salman Khan, Fahad Shahbaz Khan, and Vineeth N Balasubramanian. Towards open world object detection. In *Computer Vision and Pattern Recognition*, 2021. 2
- [17] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Computer Vision and Pattern Recognition*, 2019. 6
- [18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2012. 1, 2, 4
- [19] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015. 2
- [20] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. In *International Conference on Learning Representations*, 2022. 2, 11
- [21] Jie Li, Allan Raventos, Arjun Bhargava, Takaaki Tagawa, and Adrien Gaidon. Learning to fuse things and stuff. *arXiv preprint arXiv:1812.01192*, 2018. 5, 19
- [22] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. In *Computer Vision and Pattern Recognition*, 2022. 2, 11
- [23] Xiangtai Li, Shilin Xu, Yibo Yang Cheng, Yunhai Tong, Dacheng Tao, et al. Panoptic-partformer: Learning a unified model for panoptic part segmentation. In *European Conference on Computer Vision*, 2022. 6, 14
- [24] Yanwei Li, Xinze Chen, Zheng Zhu, Lingxi Xie, Guan Huang, Dalong Du, and Xingang Wang. Attention-guided unified network for panoptic segmentation. In *Computer Vision and Pattern Recognition*, 2019. 5, 19
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, 2014. 2, 12
- [26] Chenxi Liu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, and Alan Yuille. Recurrent multimodal interaction for referring image segmentation. In *International Conference on Computer Vision*, 2017. 2
- [27] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, et al. Swin transformer v2: Scaling up capacity and resolution. *arXiv preprint arXiv:2111.09883*, 2021. 2, 4
- [28] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *Computer Vision and Pattern Recognition*, 2016. 2
- [29] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple open-vocabulary object detection with

- vision transformers. In *European Conference on Computer Vision*, 2022. [2](#)
- [30] Rohit Mohan and Abhinav Valada. Efficientpts: Efficient panoptic segmentation. *International Journal of Computer Vision*, 129(5):1551–1579, 2021. [5](#), [19](#)
- [31] 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*, 2021. [1](#), [2](#), [3](#), [4](#), [5](#), [11](#)
- [32] Jiawei Ren, Cunjun Yu, Zhongang Cai, Mingyuan Zhang, Chongsong Chen, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Refine: prediction fusion network for panoptic segmentation. In *AAAI Conference on Artificial Intelligence*, 2021. [5](#), [19](#)
- [33] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *International Conference on Computer Vision*, 2019. [2](#)
- [34] Chufeng Tang, Hang Chen, Xiao Li, Jianmin Li, Zhaoxiang Zhang, and Xiaolin Hu. Look closer to segment better: Boundary patch refinement for instance segmentation. In *Computer Vision and Pattern Recognition*, 2021. [6](#), [13](#)
- [35] Zhi Tian, Chunhua Shen, and Hao Chen. Conditional convolutions for instance segmentation. In *European Conference on Computer Vision*, 2020. [5](#), [6](#), [11](#), [12](#), [13](#), [17](#)
- [36] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *International Conference on Computer Vision*, 2019. [5](#), [13](#)
- [37] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *Computer Vision and Pattern Recognition*, 2015. [2](#)
- [38] Qi Wu, Damien Teney, Peng Wang, Chunhua Shen, Anthony Dick, and Anton van den Hengel. Visual question answering: A survey of methods and datasets. *Computer Vision and Image Understanding*, 163:21–40, 2017. [2](#)
- [39] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *Advances in Neural Information Processing Systems*, 2021. [5](#), [6](#), [11](#), [17](#)
- [40] Lingxi Xie, Xiaopeng Zhang, Longhui Wei, Jianlong Chang, and Qi Tian. What is considered complete for visual recognition? *arXiv preprint arXiv:2105.13978*, 2021. [2](#)
- [41] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *Computer Vision and Pattern Recognition*, 2021. [2](#)
- [42] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. In *Computer Vision and Pattern Recognition*, 2019. [2](#)
- [43] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision*, 127(3):302–321, 2019. [1](#), [2](#), [5](#), [14](#), [16](#)
- [44] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. [2](#), [18](#)
- [45] Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. Visual7w: Grounded question answering in images. In *Computer Vision and Pattern Recognition*, 2016. [2](#)