

Panoptic Narrative Grounding

Cristina González¹ Nicolás Ayobi¹ Isabela Hernández¹ José Hernández¹
 Jordi Pont-Tuset² Pablo Arbeláez¹

¹Center for Research and Formation in Artificial Intelligence, Universidad de los Andes, Colombia

²Google Research, Switzerland

Abstract

This paper proposes Panoptic Narrative Grounding, a spatially fine and general formulation of the natural language visual grounding problem. We establish an experimental framework for the study of this new task, including new ground truth and metrics, and we propose a strong baseline method to serve as stepping stone for future work. We exploit the intrinsic semantic richness in an image by including panoptic categories, and we approach visual grounding at a fine-grained level by using segmentations. In terms of ground truth, we propose an algorithm to automatically transfer Localized Narratives annotations to specific regions in the panoptic segmentations of the MS COCO dataset. To guarantee the quality of our annotations, we take advantage of the semantic structure contained in WordNet to exclusively incorporate noun phrases that are grounded to a meaningfully related panoptic segmentation region. The proposed baseline achieves a performance of 55.4 absolute Average Recall points. This result is a suitable foundation to push the envelope further in the development of methods for Panoptic Narrative Grounding.

1. Introduction

Vision and language skills play a key role in humans' understanding of the world and they are rarely used independently. Their interaction is crucial in achieving high-level tasks such as describing objects, narrating a visual scene, or answering questions based on visual cues. Inspired by these capabilities of human intelligence, researchers have formulated tasks at the intersection of computer vision and natural language processing, such as image captioning [5, 18, 25, 40], referring expression comprehension and segmentation [15, 19, 31, 32], visual question answering [2], among many others.

Current experimental frameworks approach vision and language tasks at different granularity levels. Image cap-



Figure 1: **Panoptic Narrative Grounding.** Given an input image (left) with an associated caption (right), our goal is to produce a panoptic segmentation that grounds its visual objects densely (left).

tioning is among the coarsest, aiming at pairing an image with a textual description of its content. With increasing granularity, there are frameworks that intend to assign specific regions in the image to short descriptions, such as referring expression comprehension and segmentation, and region descriptions in the Visual Genome [19] dataset. The finest approaches tackle visual grounding at word level, with bounding boxes in the image linked to noun phrases in the caption, as in the Flickr30k Entities [49, 35] dataset.

Given the general dichotomy in computer vision tasks between *things* (countable objects) and *stuff* (amorphous regions of similar texture), these datasets mainly focus on things categories. However, several works [8, 1, 10, 17] have highlighted the importance of jointly considering both things and stuff classes towards real-world applications. Kirillov *et al.* [17] defined the panoptic segmentation task as the unified formulation of semantic segmentation, which recognizes *stuff*, and instance segmentation, which detects and segments *things*.

Recently, Pont-Tuset *et al.* [36] proposed Localized Narratives, a multimodal dataset for visual grounding based on the natural human task of describing images while simultaneously pointing with the mouse at the regions being described. Their grounding annotations have the densest sampling of current datasets, in that they visually ground every word in the caption with mouse trace segments. Moreover,

this dataset naturally contains panoptic descriptions of the content of each image, including *stuff* regions.

Although some of the experimental frameworks above contain sufficiently dense annotations over language, visual annotations are still very sparse and rough. This paper proposes *Panoptic Narrative Grounding* (Figure 1), a spatially finer and more general task formulation of the natural language visual grounding problem in which (i) we propose a spatially detailed visual grounding that uses segmentations instead of bounding boxes, and (ii) we include all panoptic categories to fully exploit the intrinsic semantic richness in a visual scene. We establish an experimental framework for the study of this task, including new ground-truth annotations and metrics, and we propose a strong baseline to serve as stepping stone for future work.

Considering that collecting pixel-wise annotations has a significant cost, we design a method to transfer the Localized Narrative annotations to regions into the panoptic segmentation annotations provided by the MS COCO [5, 25] dataset. We select a specific region for each noun phrase considering the trace segment associated to it. Since humans use different styles to pointing objects (circling, scribbling, and underlining) we evaluate the quality of the annotation by determining if noun phrases are grounded to a meaningfully related panoptic segmentation region. To do so, we leverage the underlying semantic information in the WordNet ontology [34]. To handle trace segments being not fully synchronized with the object description, we consider all the regions in the image ranked by relative distance to the visual grounding and assign the closest region with a strong meaning relationship. Finally, for plural noun phrases, we select all regions that (i) are within the tightest bounding box around the mouse trace and (ii) are from the same category as the main selected region. In terms of evaluation, in contrast to the traditional metric use in phrase-grounding task, we introduce a stricter metric that calculates recall at different Intersection over Union (IoU) thresholds between the target and the predicted mask. Thus, measuring both recognition accuracy and segmentation quality.

We propose a strong baseline that builds upon a state-of-the-art method, Cross-Modality Relevance [52] (CMR), developed for reasoning tasks on language and vision. Specifically, this method is designed to perform two classification tasks: Visual Questions Answering [2] and Natural Language for Visual Reasoning [43, 44]. We generalize the model to perform Panoptic Narrative Grounding. Thus, our baseline is the first natural language visual grounding method able to align multiple noun phrases with panoptic region segmentations within the image.

Our main contributions can be summarized as follows:

- (1) We propose Panoptic Narrative Grounding, a new formulation of the natural language visual grounding problem which, by using panoptic segmentation regions as

visual grounding, is spatially denser and more general in semantic terms.

- (2) We establish an experimental framework for the study of this problem, with annotations coming from the transfer of Localized Narrative annotations to panoptic segmentations in the MS COCO dataset.
- (3) We introduce the first visual grounding method that matches segmentation regions to specific noun phrases in the caption, which serves as a strong baseline for the task of Panoptic Narrative Grounding.

To ensure the reproducibility of our results and to promote further research on Panoptic Narrative Grounding, we make all the resources of this paper publicly available in our project web page¹: our benchmark dataset annotations for the train and validation splits in MS COCO, an implementation of the evaluation metrics, and the pretrained models and source code for our baseline.

2. Related Work

2.1. Vision and Language Datasets

The first datasets at the intersection of vision and language match images to descriptions without any form of visual grounding [5, 18, 25, 40]. MS COCO Captions [5, 25] comprises 123,287 images for training, validation, and testing; with five human-annotated captions for each one. Several works, however, have suggested that models developed for this task do not base their decisions on the visual information contained in the image, but rather on easier-to-learn language priors [39].

Referring Expressions Comprehension (REC) and Referring Expression Segmentation (RES) are tasks intended to detect and segment, respectively, a target object instance described by a natural language expression. There exist three datasets for these tasks, built on top of the MS COCO dataset: RefCOCO [15], RefCOCO+ [15], and RefCOCOg [32]. RefCOCO+’s descriptions, in contrast to those on RefCOCO, focus on appearance attributes rather than locations of the referent object. RefCOCOg was collected in a non-interactive setup and contains longer descriptions including both appearance and location attributes of the target instances. These datasets, while having pixel-level annotations, do not ground language at the word-level, and their annotations’ semantics boil down mostly to salient objects in images.

Visual Genome [19] is a general-purpose dataset meant to connect vision and language while not being restricted to a specific task. It contains 108,077 images with 50 region descriptions and 35 objects per image on average. Each short description is localized within the image with a bounding box, making it possible to study visual grounding at the phrase level. There is no an explicit relationship between

¹<https://github.com/BCV-Uniandes/PNG>

the objects localized within an image and the visual entities included in the textual descriptions. Having multiple regions descriptions allows to have a more complete understanding of a scene, covering *things* and *stuff* categories.

Flickr30k Entities [35] extends the original Flickr30k dataset [49] with manually annotated bounding boxes grounding each noun phrase in the caption. Despite its granularity over language, this dataset does not provide pixel-wise annotations and focus predominantly on *things*.

Localized Narratives [36] is a large-scale dataset that provides annotations for the whole MS COCO, Flickr30k [50], and ADE20K [53] datasets, and 671K images from Open Images [20]. Its annotations consist of synchronized voice recordings, text transcriptions, and mouse traces made by human annotators. They describe the image’s contents and the mouse traces visually ground each word in the narrative, which makes them the densest annotations over language available to date. Grounding goes beyond nouns and includes visual relationship indicators, verbs, *etc.* Annotators were asked to make descriptions that comprise as much content of the image as possible, which generally entails a high semantic coverage of the scene. The visual grounding using mouse trace segments, however, is spatially very coarse.

Table 1 summarizes the characteristics of existing datasets at the intersection of vision and language and compares them in terms of: (i) language granularity, (ii) visual granularity, and (iii) semantic generality. Our benchmark dataset is designed to fill the gaps between all previous datasets: (i) It maintains the finest granularity over language by grounding specific words in the narrative, as the latest developed natural language grounding datasets. (ii) It provides fine-grained visual grounding annotations using segmentations, as in referring expression segmentation datasets. This allows studying the problem with a stricter experimental framework in which the location on the objects is not approximate as in the object detection task. (iii) Our annotations include *panoptic* categories [17] instead of just focusing on *things* categories, which gives way to a global analysis of the scene which is relevant for reasoning about the world around us.

2.2. Natural Language Visual Grounding Methods

The methods developed for the coarsest natural language visual grounding tasks on language, specifically for the REC and RES tasks, can be grouped into two general approaches: top-down and bottom-up. The former paradigm [6, 27, 51] starts from region proposals extracted by a general method of object detection or segmentation. From this set of objects, the method selects the one that is described to a greater extent by the natural language expression. The strengths of this paradigm are that they can take advantage of what general object detection and seg-

Dataset	Language Granularity	Visual Granularity	Semantic Generality
MS COCO Captions [5, 25]	Caption	×	×
Conceptual Captions [40]	Caption	×	×
Stanford Visual Par. [18]	Caption	×	×
ReferIt [15, 32]	Short phrase	Segmentation	Things
Google Refexp [31]	Short phrase	Segmentation	Things
Visual Genome [19]	Short phrase	Bounding box	Things + Stuff
Flickr30k Entities [49, 35]	Noun phrase	Bounding box	Mainly Things
Localized Narratives [36]	Each word	Traces	Things + Stuff
Panoptic Narrative Grounding (Ours)	Noun phrase	Segmentation	Things + Stuff

Table 1: Panoptic Narrative Grounding compared with major captioning and natural language grounding datasets.

mentation methods have learned from the existing large-scale datasets for these general tasks, rather than directly learning how to detect objects within an image. However, since these methods do not update the region proposals, the upper bound of these methods is significantly impacted by the performance of the base method. In contrast, the latter approaches produce the referred object segmentation by grouping pixels [4, 11, 12, 13, 14, 23, 24, 26, 29, 30, 33, 37, 38, 41, 47, 48]. These methods use single networks, which leverage high and low-level features to refine segmentation masks along levels of their architectures. However, complying with a larger search space than its top-down counterpart, leads to higher rates of false positive segmentations.

Some existing methods approach phrase grounding through a combination of base networks for visual and linguistic information. To this end, Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN) are used [11, 33], and their features leveraged to achieve word-to-image interactions. Other works apply Transformer attention [45] to generate contextualized representations of both modalities. The latter architectures are also used in cross-modality analyses to aggregate and align visual cues and linguistic meanings [21, 22, 28, 42, 52].

3. Panoptic Narrative Grounding Benchmark

This section describes our proposed benchmark. We describe how we generate the ground-truth annotations (Sec. 3.1) and report some statistics (Sec. 3.2). Finally, we explain the proposed metrics to evaluate results on our benchmark (Sec. 3.3).

3.1. Ground-Truth Annotations

We transfer the Localized Narrative [36] annotations to MS COCO panoptic segmentations by synchronizing timed captions and trace points. We consider a set of utterances $U = \{u_1, \dots, u_n\}$ in the caption and a set of trace points $T = \{t_1, \dots, t_m\}$ for each Localized Narrative, where u_i and t_j are the timestamped verbal and spatial units that, respectively, compose a caption and a mouse trace.

Synchronization occurs by selecting the trace segments

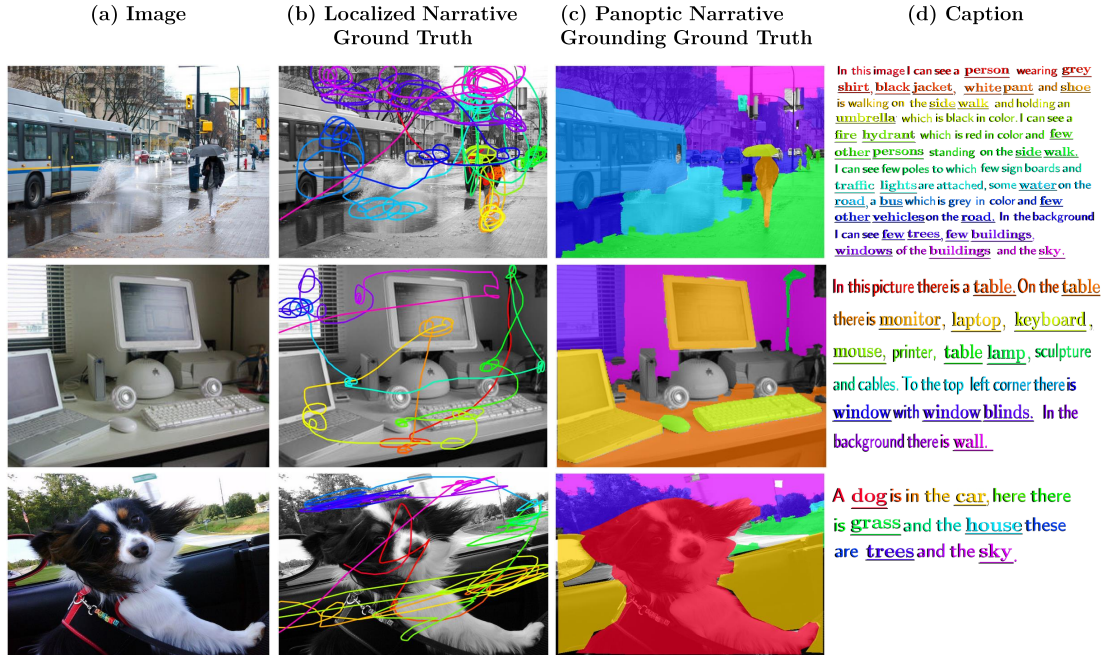


Figure 2: **Ground-Truth Annotations Results.** Examples of different Panoptic Narrative Grounding ground truth resulting from the proposed annotation transfer algorithm (c). We show the input image (a) and Localized Narrative traces (b) and caption with the matched panoptic segmentation regions (d). Color gradient in the trace, panoptic segmentation and caption indicates time over the language. The segmentation regions are visualized with the color of their corresponding noun phrases, according to the last associated spoken word. More qualitative results can be found in the Supplementary Material.

that are drawn over an image between the start and end times of an utterance, or a subset of contiguous utterances, containing a noun phrase (u_{in}). Each u_{in} is hence associated with a subset of trace points $T' = \{t_a, \dots, t_b\}$ (where t_a and t_b are within the timestamp of u_i). T' provides a spatial reference for each utterance through their point coordinates. We average these coordinates to obtain a single point p in the image plane, which we refer to as the Center of Mass (CoM). In the Supplementary Material, we exemplify how this strategy of summarizing the annotator’s grounding is beneficial for irregular and circling patterns in the Localized Narrative annotations.

To identify noun phrases, we perform a chunking parsing using the Natural Language Toolkit (NLTK) for Python [3], and select sequences of consecutive nouns that may or may not have a previous adjective or a cardinal digit. These noun phrases are considered a single grounding unit.

Let $S = \{s_1, \dots, s_k\}$ be the regions from the MS COCO panoptic segmentation of the same image, with corresponding *thing* or *stuff* category labels $C = \{c_1, \dots, c_k\}$. We select the region $s_i \in S$ which contains the location of the CoM p of the mouse trace and establishes it as a candidate for the ground truth of the noun phrase u_{in} . This matching, however, is not guaranteed to be correct, as Localized Narrative mouse traces might not fall exactly on the grounded

object. We, therefore, filter these matches by comparing the noun phrase in the caption (u_{in}) to the MS COCO object category associated with s_i (c_i), to which should *agree*. Intuitively, if the annotator said $u_{in} = \text{“red vehicle”}$, we want to consider the matching correct if *e.g.* $c_i = \text{“vehicle”}$ or $c_i = \text{“car”}$; but we want to discard it if $c_i = \text{“tree”}$, since it might mean that the annotator was pointing to a tree next to a vehicle when they pronounced the word “vehicle”. Furthermore, if the caption refers to the same object in the visual scene multiple times, our benchmark will associate it with all the related noun phrases. Figure 2 shows some examples of our visual grounding annotations.

To evaluate whether the caption utterance(s) u_{in} , or its composing noun(s), correspond to the matched panoptic category c_i , we consider the following criteria: (1) *exact matches* (u_{in} is strictly equal to c_i), (2) *synonyms* (u_{in} is a synonym of c_i), (3) *hierarchical relationship* (u_{in} is either a hypernym or hyponym of c_i), and (4) *meronyms* (u_{in} is a meronym of c_i). We evaluate (1) by simple string comparison, and (2), (3) and (4) accessing WordNet [34] using NLTK [3]. Additionally, (5) we *manually* relate specific words to certain MS COCO categories that are omitted by the WordNet ontology. Examples of these words include clothing pieces, body parts and female figures, as members of the MS COCO “person” category. We consider (1) as



Figure 3: **Examples of grounded plural noun phrases.** We match multiple regions (red) beyond the seed region (purple). Increasing the spatial scope of our annotations agrees with the multi-instance nature of plural noun phrases and better captures the annotators’ intention of referring to several instances. Best viewed in color.

the highest rank in matching, given the similarity between words in this level, followed by (2), (3), (4) and (5). Tag clouds are reported in the Supplementary Material as examples of criteria (2) - (4).

We find two special cases that require additional measures when transferring annotations. First, there is a possibility of time shifts [36]: an annotator moving their mouse slightly before or after describing objects. To address this issue, we consider neighboring segmentation regions of the selected one via p (the CoM of the mouse traces), as potential matches to each noun phrase u_{in} . If no matching occurs with the center region, we proceed to select the closest candidate that matches u_{in} , with any of the semantic relationships defined above, as the grounding annotation for the noun phrase. We consider the minimum distance between segmentation regions as a measure of their closeness (following a *single-linkage* concept). An example of vicinity region analysis can be found in the Supplementary Material.

Second, we consider the implications of plural noun phrases within narratives. These phrases are inherently related to various instances of an object, which are all globally pointed at during the utterance(s). This motivates the mapping of a plural noun phrase to $S' \subseteq S$, which contains several regions with a common category, as opposed to a single segmentation s_i . The common category is defined using p and the vicinity region analysis, from which we select a *seed* region. We then augment this set with all the regions of the same category that are contained in the tightest bounding box around the mouse trace T' . All regions in S' are considered the grounding of the plural noun phrase u_{in} . Figure 3 shows examples of plural groundings.

3.2. Dataset Statistics

Our complete annotation transfer methodology matches an average of 5.1 noun phrases per narrative and 726,445 noun phrases in the whole Localized Narratives annota-

tions for MS COCO dataset. Localized Narratives contain an average of 11.3 noun phrases per narrative. This accounts for 45.1% of noun phrases coverage, which we ground through panoptic segmentation annotations. In the visual domain, this translates to a total of 741,697 segments matched, from which 659,298 are unique. Thus, our benchmark covers 47.5% of all the segments annotated in MS COCO panoptic with at least one noun phrase match. In turn, this coverage accounts for 58.5% of all pixels in 123,287 MS COCO images and 65.4% relative to all annotated pixels in MS COCO panoptic annotations.

The proposed Panoptic Narrative Grounding dataset ground both *things* and *stuff* categories. Specifically, 57.0% of the noun phrases grounded are *things* and 43.0% are *stuff*. 64.7% of matches between grounded noun phrases and MS COCO panoptic categories are exact matches between the category name and a noun in the phrase. Thus, using synonyms, manual matching, and hierarchical and meronymy relationships significantly extend the grounding scope (35.3%) and allow us to better capture the diversity in natural language expressions in Localized Narratives. Lastly, 29.3% of the matched noun phrases correspond to plurals and 44.6% of matches are due to the vicinity region analysis. Consequently, each of the proposed matching steps plays an important role in improving the visual grounding annotations. A detailed report of these statistics can be found in the Supplementary Material. Moreover, we carry out a final manual curation to verify the integrity of our annotations in the 1% of the annotations.

3.3. Metrics

In contrast to the recall metric, traditionally used for phrase grounding, we propose to calculate the Average Recall. This metric evaluates the performance of the method in the Panoptic Narrative Grounding task by considering different Intersection over Union (IoU) thresholds between the panoptic segmentation proposal and the ground-truth for each noun phrase. Hence, the quality of the segmentations affects the method’s performance since the IoU metric determines whether a detection is considered a true positive or not. We obtain a curve that at very low IoU values has recall close to one and at higher IoU values recall drops. The final metric, named Average Recall, is the area under the curve described above. For plural noun phrases, we do not explicitly match the ground-truth instances and the prediction proposals to calculate the Average Recall. Instead, we aggregate all instances of the ground-truth annotation into a single segmentation and compute the IoU with respect to the segmentation composed by all the prediction proposals. With this strategy, we avoid errors or variations in the matching process between annotations and predictions while assessing the overall segmentation quality.

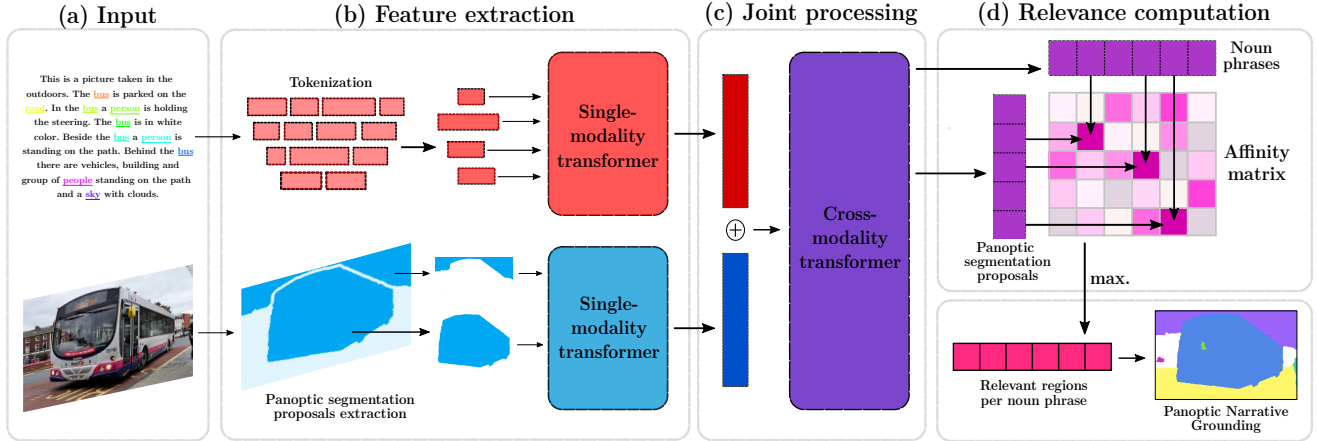


Figure 4: **Panoptic Narrative Grounding Baseline.** Our model takes an image and a corresponding caption as input (a), followed by a top-down approach by extracting panoptic segmentation proposals with a backbone segmenter. We then use single-modality transformers to independently process the segmentation proposals and the tokens from the caption (b) and jointly process the information with a cross-modality transformer (c). Finally, we calculate an affinity matrix between visual and textual information and select the more relevant panoptic segmentation proposal for each noun phrase (d).

4. Baseline for Panoptic Narrative Grounding

We build upon the Cross-Modality Relevance (CMR) [52] model developed for reasoning tasks on language and vision. This method has competitive results in two classification tasks at the intersection of vision and language: Visual Question Answering [2] and Natural Language for Visual Reasoning [43, 44]. CMR model introduces entity relevance representation that explicitly expressed the relevance of textual entities with respect to visual entities. The model uses this affinity matrix as an intermediate representation for the final tasks. Their results suggest using the word alignment with regions in the image results in an increase in performance for the final task. We adapt the model and optimize the architecture for our task.

Figure 4 depicts an overview of our baseline method. Given an image and its caption (a), we extract features from each of the panoptic segmentation region proposals and process each proposal and word with single modality transformers. (b). We then concatenate the outputs from the visual modality transformer and the textual modality transformer and use it as input to the cross-modality transformer that updates the representations of each entity considering not only those of the same modality but also those of the other modality (c). Afterwards, we calculate the affinity matrix from the matrix multiplication between the representations of each word and region proposal. Finally, we make an average pooling over the language dimension of all the words included in each noun phrase. This affinity matrix explicitly indicates which is the most relevant panoptic segmentation region for each noun phrase in the caption (d).

Implementation details: We use a Panoptic Feature Pyra-

mid Network [17] (FPN) with a ResNet-101 [9] backbone pre-trained on MS COCO [5, 25] with $3\times$ schedule using the official implementation [46]. Parameters in Panoptic FPN are fixed. For stuff proposals, features correspond to the ones extracted from the semantic segmentation branch after combining the features extracted from each FPN level and before the final upsampling. For things categories, they are the mask features in the instance segmentation branch. The visual single modality transformer has 3 attention layers. We use the pre-trained BERT “base” model [7] to generate the single-modality representation. The cross-modality transformer has 5 attention layers. During training, we performed a multi-label classification task to include several regions when the noun phrase is plural. Additionally, for plurals in the final prediction we consider all the panoptic segmentation region proposal with a match score greater than 0.1. We train our method on an NVIDIA Quadro RTX 8000 GPU for 25 epochs with an initial learning rate of 0.000,01, 60 images per batch and use an Adam [16] optimizer with standard parameters.

5. Experiments

Experimental setup: We use our Panoptic Narrative Grounding benchmark with the standard training and validation splits of the MS COCO dataset. We only include the images for which our annotation transfer process assigns at least one noun phrase of the Localized Narrative caption to a panoptic segmentation region. The final splits have 133,103 localized narratives for train and 8533 for validation.

Quantitative results: Table 2 and Figure 5a show the quantitative results of our method using the proposed Average

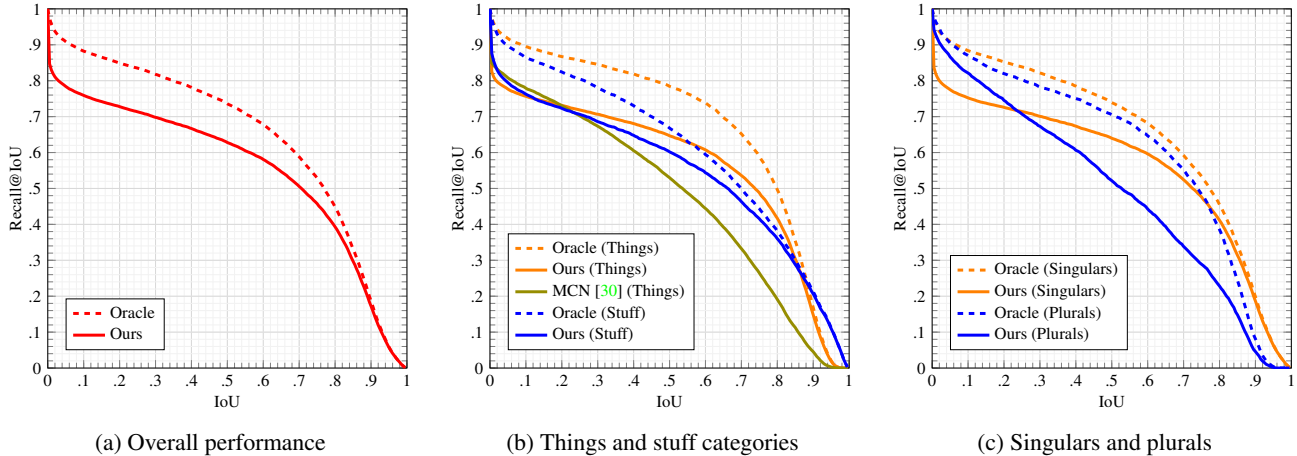


Figure 5: **Average Recall Curve** for our baseline method performance (a) compared to the oracle, and disaggregated into (b) things and stuff categories, and (c) singulars and plurals noun phrases.

Method	Average Recall		
	things + stuff	things	stuff
Oracle	64.4	67.3	60.4
Ours	55.4	56.2	54.3
MCN[30]	-	48.2	-

(a) Things and stuff categories.

Method	Average Recall		
	singulars + plurals	singulars	plurals
Oracle	64.4	64.8	60.7
Ours	55.4	56.2	48.8

(b) Singulars and plurals noun phrases.

Table 2: **Results of our method** for Panoptic Narrative Grounding task compared to the oracle performance, disaggregated into (a) things and stuff categories, and (b) singulars and plurals noun phrases.

Recall evaluation metric (Sec. 3.3). Our method’s relative performance is 86.0% with a drop of only 9.0 absolute Average Recall points with respect to the oracle performance given by the base panoptic segmentation method (55.4 vs. 64.4). We calculate this oracle by choosing as prediction the panoptic segmentation proposal with the highest IoU with the ground-truth annotation. Therefore, the oracle version of our method assumes a perfect match between segments and noun phrases, and only measures segmentation errors. These results suggest that our proposed baseline is very strong and a good starting point to promote research in Panoptic Narrative Grounding. However, our proposed baseline method is limited by the quality of the proposals’ segmentations and the recall of the base panoptic segmentation method.

Furthermore, to compare our baseline for Panoptic Narrative Grounding with respect to a visual grounding state-of-the-art method, Multi-task Collaborative Network [30] (MCN). We generalize our task and make a natural extension of RES task. For this purpose, we split the captions into phrases and include as ground truth all segmentation regions that match a noun phrase within each sentence. Since this method was developed in the scope of RES task, we

just include things categories. This state-of-the-art method for RES follows and bottom-up approach and leverages the complementary relationship between REC and RES tasks by using their properties to benefit the other task. MCN achieves a performance of 48.2 Average Recall. This result is comparable with the performance of our disaggregated method in only things, which shows that our baseline method better address the problem of natural language visual grounding. We hypothesize that approaching this task from a fine-grained formulation allows the model to easily abstract information using the semantics of each word.

By disaggregating the performance of our baseline method into *things* and *stuff* (Tab. 2a, Fig. 5b), it is possible to observe that, even though our method’s performance for *stuff* categories is higher, the drop with respect to the oracle performance for *things* categories is 5.0 absolute Average Recall points greater. We hypothesize that this is due to (i) greater variations in position and appearance in *things* categories. These results suggest that aligning noun phrases and panoptic region proposals gets more difficult as the intra-class visual variability increases. Also, (ii) there can be multiple instances in an image of the same object category in the case of *things*, as opposed to *stuff* categories, so the

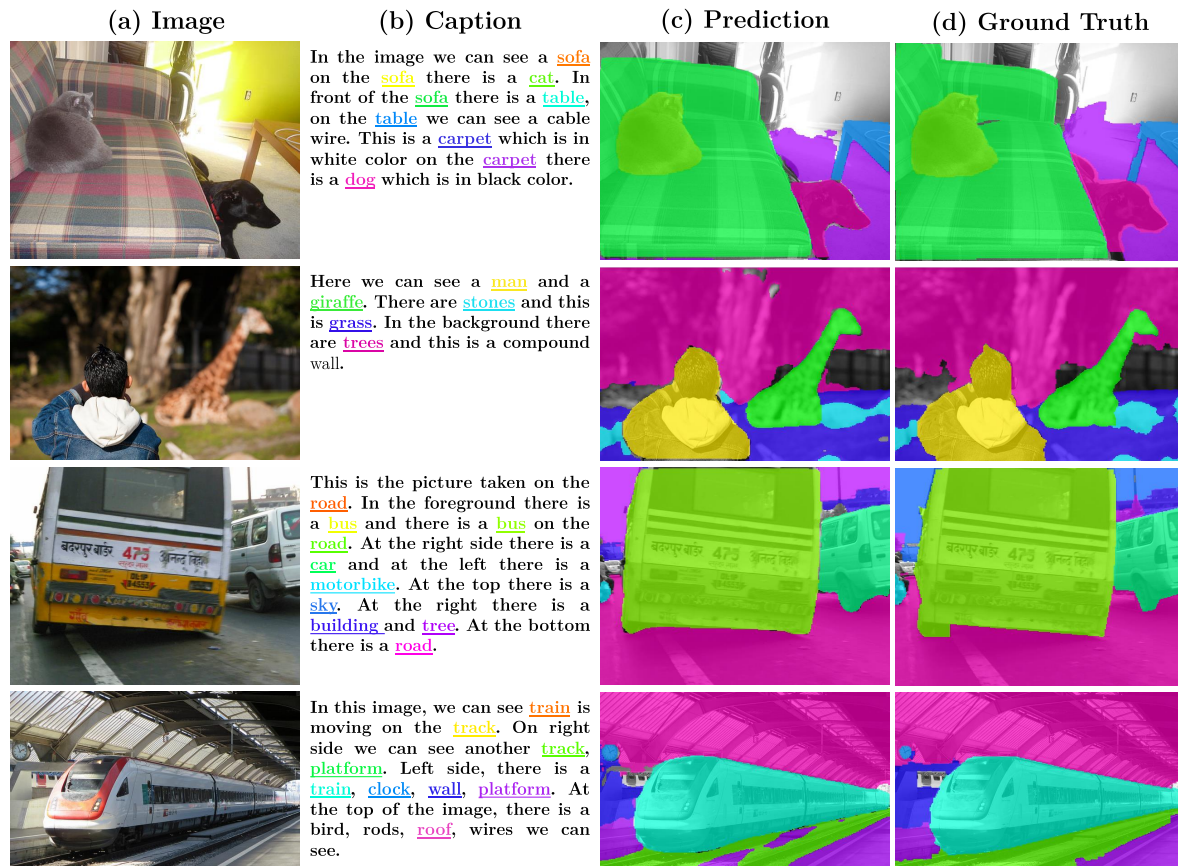


Figure 6: **Qualitative results for Panoptic Narrative Grounding.** Example predictions of our baseline in the validation split of our benchmark. The inputs are the image (a) and the caption without highlighted noun phrases (b). The outputs are a set of noun phrases in the caption (b), each with a corresponding region in the predicted segmentation (c). (d) shows the ground-truth panoptic segmentations. More qualitative results can be found in the Supplementary Material.

model fails to disambiguate between the instances of that category and select the one to which the noun phrase refers in the caption’s context.

We also disaggregate the performance on singular and plural noun phrases (Tab. 2b, Fig. 5c). The drop in performance for plural noun phrases indicates that our model does not retrieve all the instances of the object to which language refers. We attribute this limitation to the trace not exhaustively intersecting all the instances in many cases, which spreads to our annotations and ends up affecting the results.

Figure 6 shows some qualitative results of our baseline method compared to the ground-truth panoptic segmentations of our benchmark. More qualitative results can be found in the Supplementary Material.

6. Conclusions

In this work, we present Panoptic Narrative Grounding, a new formulation for the natural language visual ground-

ing problem that aims at producing a panoptic segmentation that grounds densely each noun phrase in the caption. This version of the problem (i) maintains the finest granularity over language by visually grounding noun phrases, while (ii) including a spatially detailed visual grounding with segmentations, and (iii) incorporating all panoptic categories to leverage the intrinsic semantic information in visual scenes. We establish a strong experimental framework for the study of this task, including new annotations and evaluation metrics. We design an algorithm to transfer Localized Narratives’ visual grounding to specific regions in the panoptic segmentations of the MS COCO dataset. Furthermore, we propose a strong baseline method to serve as stepping stone for future work. The formulation of Panoptic Narrative Grounding, its benchmark and experimental framework will push the envelope further in the development of fine and general natural language visual grounding methods. In turn, advances in this task will impact solutions for other problems on the intersection of vision and language.

References

- [1] Edward H Adelson. On seeing stuff: the perception of materials by humans and machines. In *Human vision and electronic imaging VI*, volume 4299, pages 1–12. International Society for Optics and Photonics, 2001. 1
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. VQA: Visual question answering. In *ICCV*, 2015. 1, 2, 6
- [3] Steven Bird, Ewan Klein, and Edward Loper. *Natural Language Processing with Python*. O’Reilly Media, first edition, 2009. 4
- [4] Ding-Jie Chen, Songhao Jia, Yi-Chen Lo, Hwann-Tzong Chen, and Tyng-Luh Liu. See-through-text grouping for referring image segmentation. In *ICCV*, 2019. 3
- [5] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO captions: Data collection and evaluation server. *arXiv:1504.00325*, 2015. 1, 2, 3, 6
- [6] Yi-Wen Chen, Yi-Hsuan Tsai, Tiantian Wang, Yen-Yu Lin, and Ming-Hsuan Yang. Referring expression object segmentation with caption-aware consistency. *arXiv:1910.04748*, 2019. 3
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv:1810.04805*, 2018. 6
- [8] David A Forsyth, Jitendra Malik, Margaret M Fleck, Hayit Greenspan, Thomas Leung, Serge Belongie, Chad Carson, and Chris Bregler. Finding pictures of objects in large collections of images. In *International workshop on object representation in computer vision*, 1996. 1
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [10] Jeremy Heitz and Daphne Koller. Learning spatial context: Using stuff to find things. In *ECCV*, 2008. 1
- [11] Ronghang Hu, Marcus Rohrbach, and Trevor Darrell. Segmentation from natural language expressions. In *ECCV*, 2016. 3
- [12] Zhiwei Hu, Guang Feng, Jiayu Sun, Lihe Zhang, and Huchuan Lu. Bi-directional relationship inferring network for referring image segmentation. In *CVPR*, 2020. 3
- [13] Shaofei Huang, Tianrui Hui, Si Liu, Guanbin Li, Yunchao Wei, Jizhong Han, Luoqi Liu, and Bo Li. Referring image segmentation via cross-modal progressive comprehension. In *CVPR*, 2020. 3
- [14] Tianrui Hui, Si Liu, Shaofei Huang, Guanbin Li, Sansi Yu, Faxi Zhang, and Jizhong Han. Linguistic structure guided context modeling for referring image segmentation. In *ECCV*, 2020. 3
- [15] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *EMNLP*, pages 787–798, 2014. 1, 2, 3
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014. 6
- [17] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *CVPR*, 2019. 1, 3, 6
- [18] Jonathan Krause, Justin Johnson, Ranjay Krishna, and Li Fei-Fei. A hierarchical approach for generating descriptive image paragraphs. *arXiv:1611.06607*, 2016. 1, 2, 3
- [19] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 123(1):32–73, 2017. 1, 2, 3
- [20] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The Open Images Dataset V4. *IJCV*, pages 1–26, 2020. 3
- [21] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. VisualBERT: A simple and performant baseline for vision and language. *arXiv:1908.03557*, 2019. 3
- [22] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. What does bert with vision look at? In *ACL*, 2020. 3
- [23] Ruiyi Li, Kaican Li, Yi-Chun Kuo, Michelle Shu, Xiaojuan Qi, Xiaoyong Shen, and Jiaya Jia. Referring image segmentation via recurrent refinement networks. In *CVPR*, 2018. 3
- [24] Zhenyang Li, Ran Tao, Efstratios Gavves, Cees GM Snoek, and Arnold WM Smeulders. Tracking by natural language specification. In *CVPR*, 2017. 3
- [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 *ECCV*, 2014. 1, 2, 3, 6
- [26] Chenxi Liu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, and Alan Yuille. Recurrent multimodal interaction for referring image segmentation. In *ICCV*, 2017. 3
- [27] Daqing Liu, Hanwang Zhang, Feng Wu, and Zheng-Jun Zha. Learning to assemble neural module tree networks for visual grounding. In *ICCV*, 2019. 3
- [28] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. ViL-BERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *arXiv:1908.02265*, 2019. 3
- [29] Gen Luo, Yiyi Zhou, Rongrong Ji, Xiaoshuai Sun, Jinsong Su, Chia-Wen Lin, and Qi Tian. Cascade grouped attention network for referring expression segmentation. In *ACM Multimedia*, 2020. 3
- [30] Gen Luo, Yiyi Zhou, Xiaoshuai Sun, Liujuan Cao, Chenglin Wu, Cheng Deng, and Rongrong Ji. Multi-task collaborative network for joint referring expression comprehension and segmentation. In *CVPR*, 2020. 3, 7
- [31] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L. Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. *arXiv:1511.02283*, 2015. 1, 3
- [32] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation

- and comprehension of unambiguous object descriptions. In *CVPR*, 2016. 1, 2, 3
- [33] Edgar Margffoy-Tuay, Juan C Pérez, Emilio Botero, and Pablo Arbeláez. Dynamic multimodal instance segmentation guided by natural language queries. In *ECCV*, 2018. 3
- [34] George A Miller. *WordNet: An electronic lexical database*. MIT press, 1998. 2, 4
- [35] Bryan A. Plummer, Liwei Wang, Christopher M. Cervantes, Juan C. Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k Entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. *IJCV*, 123(1):74–93, 2017. 1, 3
- [36] Jordi Pont-Tuset, Jasper Uijlings, Soravit Changpinyo, Radu Soricut, and Vittorio Ferrari. Connecting vision and language with localized narratives. In *ECCV*, 2020. 1, 3, 5
- [37] Shuang Qiu, Yao Zhao, Jianbo Jiao, Yunchao Wei, and Shikui Wei. Referring image segmentation by generative adversarial learning. *IEEE Transactions on Multimedia*, 22(5):1333–1344, 2019. 3
- [38] Xuejian Rong, Chucai Yi, and Yingli Tian. Unambiguous scene text segmentation with referring expression comprehension. 29:591–601, 2019. 3
- [39] Ramprasaath R Selvaraju, Stefan Lee, Yilin Shen, Hongxia Jin, Shalini Ghosh, Larry Heck, Dhruv Batra, and Devi Parikh. Taking a hint: Leveraging explanations to make vision and language models more grounded. In *ICCV*, 2019. 2
- [40] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*, 2018. 1, 2, 3
- [41] Hengcan Shi, Hongliang Li, Fanman Meng, and Qingbo Wu. Key-word-aware network for referring expression image segmentation. In *ECCV*, 2018. 3
- [42] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. VL-BERT: Pre-training of generic visual-linguistic representations. *arXiv:1908.08530*, 2020. 3
- [43] Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. A corpus of natural language for visual reasoning. In *ACL*, 2017. 2, 6
- [44] Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Hua-jun Bai, and Yoav Artzi. A corpus for reasoning about natural language grounded in photographs. *arXiv:1811.00491*, 2018. 2, 6
- [45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 3
- [46] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 6
- [47] Linwei Ye, Zhi Liu, and Yang Wang. Dual convolutional lstm network for referring image segmentation. *IEEE Transactions on Multimedia*, 22(12):3224–3235, 2020. 3
- [48] Linwei Ye, Mrigank Rochan, Zhi Liu, and Yang Wang. Cross-modal self-attention network for referring image segmentation. In *CVPR*, 2019. 3
- [49] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *TACL*, 2:67–78, 2014. 1, 3
- [50] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *TACL*, 2:67–78, 2014. 3
- [51] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. MATTNet: Modular attention network for referring expression comprehension. In *CVPR*, 2018. 3
- [52] Chen Zheng, Quan Guo, and Parisa Kordjamshidi. Cross-modality relevance for reasoning on language and vision. *arXiv:2005.06035*, 2020. 2, 3, 6
- [53] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, 2017. 3