

Modulating Bottom-Up and Top-Down Visual Processing via Language-Conditional Filters

Ilker Kesen^{1,2} Ozan Arkan Can³ Erkut Erdem^{1,4} Aykut Erdem^{1,2} Deniz Yuret^{1,2}

¹ Koç University, KUIS AI Center ² Koç University, Computer Engineering Department

³ Amazon Search ⁴ Hacettepe University, Computer Engineering Department

Abstract

How to best integrate linguistic and perceptual processing in multi-modal tasks that involve language and vision is an important open problem. In this work, we argue that the common practice of using language in a top-down manner, to direct visual attention over high-level visual features, may not be optimal. We hypothesize that the use of language to also condition the bottom-up processing from pixels to high-level features can provide benefits to the overall performance. To support our claim, we propose a U-Net-based model and perform experiments on two language-vision dense-prediction tasks: referring expression segmentation and language-guided image colorization. We compare results where either one or both of the top-down and bottom-up visual branches are conditioned on language. Our experiments reveal that using language to control the filters for bottom-up visual processing in addition to top-down attention leads to better results on both tasks and achieves competitive performance. Our linguistic analysis suggests that bottom-up conditioning improves segmentation of objects especially when input text refers to low-level visual concepts. Code is available at <https://github.com/ilkerkesen/bvpr>.

1. Introduction

As human beings, we can easily perceive our surroundings with our visual system and interact with each other using language. Since the work of Winograd [75], developing a system that understands human language in a situated environment has been one of the long-standing goals of artificial intelligence. Recent successes of deep learning studies in both language and vision domains have increased the interest in tasks that combine language and vision [2, 3, 33, 41, 69, 77]. However, how to best integrate linguistic and perceptual processing is still an important open problem. Towards this end, we investigate whether language should be used for conditioning bottom-up visual

processing as well as top-down attention.

In the human visual system, attention is driven by both *top-down* cognitive processes (e.g. focusing on a given shape) and *bottom-up* salient, behaviourally relevant stimuli (e.g. fast moving objects and contrasting colors) [13, 14, 71]. Studies on embodied language explore the link between linguistic and perceptual representations [21, 64, 73] and often assume that language has a *high-level* effect on perception and drives the *top-down* visual attention [5, 18, 36]. However, recent studies from cognitive science point out that language comprehension also affects low-level visual processing [6, 52, 59]. Motivated by this, we propose a model that can modulate either or both of *bottom-up* and *top-down* visual processing with language and compare different designs for language modulation.

Current deep learning systems for language-vision tasks typically start with low-level image processing, then connect the language representation with high-level visual features to control the visual focus. To integrate both modalities, concatenation [54], element-wise multiplication [40, 49], attention from language to vision [1, 48, 76, 83, 88] and transformers [17, 19, 70] are commonly used. These studies typically do not condition low-level visual features on language. Some methods [15, 63] do the opposite by conditioning only the bottom-up visual processing on language.

To evaluate language-modulation on the bottom-up and top-down visual branches independently, we develop an architecture that clearly separates these two branches (based on U-Net [66]) and allows us to experiment with modulating one or both branches with language. The bottom-up branch starts from low-level visual features and applies a sequence of contracting filters that result in successively higher level feature maps with lower spatial resolution. Following this, a top-down branch takes the final low resolution feature map and applies a sequence of expanding filters that eventually result in a map with the original image resolution. Information flows between branches through skip connections between contracting and expanding filters at the same level. Our proposed architecture is task-agnostic

and it can be used for various vision-language tasks involving dense prediction. We evaluate our model with different language-modulation settings on two different tasks: referring expression segmentation and language-guided image colorization.

In the *referring expression segmentation* (RES) task, given an image and a natural language description, the aim is to obtain the segmentation mask that marks the object(s) described. We can contrast this with pure image based object detection [24, 65] and semantic segmentation [9, 47] tasks which are limited to predefined semantic classes. The language input may contain various visual attributes (e.g. *shape*), spatial information (e.g. *in front of*), actions (e.g. *running*) and interactions/relations between different objects (e.g. *arm of the chair that the cat is sitting on*).

In *language-guided image colorization* (LIC) task, given a grayscale image and a description, the aim is to predict pixel color values. The absence of color information in the input images makes this problem interesting to experiment with because color words do not help in conditioning the bottom-up branch when the input image is grayscale.

We find that conditioning both branches leads to better results, achieving competitive performance on both tasks. Our experiments suggest that conditioning the bottom-up branch on language is important to ground low-level visual information. On RES, we find that modulating only the bottom-up branch performs significantly better than modulating only the top-down branch especially when color-dependent language is present in the input. Our findings on LIC show that when color information absent in input images, the bottom-up baseline naturally fails to predict and manipulate colors of target objects specified by input language. That said, conditioning the bottom-up branch still improves the colorization quality by helping our model to accurately segment and colorize the target objects as a whole.

The rest of the paper is structured as follows: We summarize related work and compare it to our approach in Section 2. We describe our model in detail in Section 3. We share the details of our experiments in Section 4. Section 5 summarizes our contributions.

2. Related Work

Referring Expression Comprehension (REC). In this problem, the goal is to locate a bounding box for the object(s) described in the input language. The proposed solutions can be divided into two categories: two-stage and one-stage methods. Two-stage methods [12, 29, 29, 46, 56, 61, 74, 78, 86, 86] rely on a pre-trained object detector [26, 65] to generate object proposals in the first stage. In the second stage, they assign scores to the object proposals depending on how much they match with input language. One-stage methods [17, 43, 51, 79, 81, 82] directly localize the referred

objects in one step. Most of these methods condition only the top-down visual processing on language, while some fuse language with multi-level visual representations.

Referring Expression Segmentation (RES). In this task, the aim is to generate a segmentation mask for the object(s) referred in the input language [30]. To accomplish this, multi-modal LSTMs [45, 58], ConvLSTMs [7, 45, 68, 84], word-level attention [7, 34, 67, 84], cross-modal attention [31, 32, 37, 50, 51, 85], and transformers [19, 72] have been used. Each one of these methods modulates only the top-down branch with language. As one exception, EFN [20] conditions the bottom-up branch on language, but does not modulate the top-down branch with language.

Language-guided Image Colorization (LIC). In this task, the aim is to predict colors for all the pixels of a given input grayscale image based on input descriptive text. Specifically, [55] inserts extra Feature-wise Linear Modulation (FiLM) layers [63] into a pre-trained ResNet to predict color values in LAB color space. Multi-modal LSTMs [45, 91] and generative adversarial networks [4, 8, 25] are also used in this context to colorize sketches. Similar to us, Tag2Pix [39] extends U-Net to perform colorization on line art data, but it modulates only the top-down visual processing with symbolic input using concatenation.

Language-conditional Parameters. Here we review methods that use input-text-dependent dynamic parameters to process visual features. To control a visual model with language, MODERN and FiLM [15, 63] used conditional batch normalization layers with language-conditioned coefficients rather than customized filters. Numerous methods [11, 22, 23, 42, 58, 60] generate language-conditional dynamic filters to convolve visual features. Some RES models [11, 58] also incorporate language-conditional filters into their top-down visual processing. To map instructions to actions in virtual environments, LingUNet [60] extends U-Net by adding language-conditional filters to the top-down visual processing only. Each one of these methods conditions either top-down or bottom-up branch only.

Comparison. To support our main research question, our architecture clearly separates bottom-up and top-down visual processing. This allows us to experiment with modulating either one branch or both branches with language and evaluate their individual contributions. The majority of related work conditions only the top-down visual processing on language. Other U-Net-based methods [39, 60] and most transformer models [17, 19] which implement cross-modal attention between textual and visual representations in top-down visual processing fall into this category. A few exceptions [15, 20, 63] do the opposite by conditioning only the bottom-up branch. Some methods [32, 51, 82] fuse language with a multi-level visual representation, which leads to good results, but this kind of fusion does not allow the evaluation of language conditioning on top-down vs bottom-up visual

processing. Our architecture allows language to control either or both of top-down and bottom-up branches. We show that (i) the bottom-up conditioning is important to ground language to low-level visual features, and (ii) conditioning both branches on language leads to best results.

3. The Model

Here, we describe our proposed model in detail. Figure 1 shows an overview of our proposed architecture. First, the model extracts a tensor of low-level features using a pre-trained convolutional neural network and encodes the given natural language expression to a vector representation using an LSTM [28]. Starting with the visual feature tensor, the model generates feature maps through a contracting and expanding path where the output head takes the final map and performs dense prediction, similar to U-Net [66]. Our proposed architecture modulates both of the contracting and expanding paths with language using convolutional filters generated from the given expression. It is important to emphasize that the previous works either have a language-guided top-down or a language-conditional bottom-up visual processing branch. As will be discussed in the next section, our experiments show that modulating both of these paths improves the performance dramatically.

3.1. Low-level Image Features

Given an input image I , we extract visual features I_0 by using a pre-trained convolutional network. In the RES task, we use DeepLab-v3+ backbones [10], and in the LIC task, we use ResNet-101 pre-trained on ImageNet [16, 27]. On the task of referring expression segmentation, we also concatenate 8-D location features to this feature map following the previous work [45, 85].

3.2. Language Representation

Consider a language input $S = [w_1, w_2, \dots, w_n]$ where each word w_i is represented with a 300-dimensional GloVe embedding [62]. We map the language input S to hidden states using an LSTM as $h_i = \text{LSTM}(h_{i-1}, w_i)$. We use the final hidden state of the LSTM as the language representation, $r = h_n$. Later on, we split this language representation into pieces to generate language-conditional filters.

3.3. Multi-modal Encoder

After generating image and language representations, our model generates a multi-modal feature map representing the input image and the given natural language expression. Our multi-modal encoder module extends U-Net by conditioning both contracting and expanding branches on language using language-conditional filters.

In the bottom-up branch, our model applies m convolutional modules to the image representation I_0 . Each mod-

ule, CNN_i , takes the concatenation of the previously generated feature map (F_{i-1}) and its convolved version with language-conditional filters K_i^F and produces an output feature map (F_i). Each CNN_i has a 2D convolution layer followed by batch normalization [35] and ReLU activation function [53]. The convolution layers have 5×5 filters with $\text{stride} = 2$ and $\text{padding} = 2$ halving the spatial resolution, and they all have the same number of output channels.

Similar to [60], we split the textual representation r to m equal parts (t_i), and then use each t_i to generate a language-conditional filter for i th bottom-up layer (K_i^F):

$$K_i^F = \text{AFFINE}_i^F(t_i) \quad (1)$$

Each AFFINE_i^F is an affine transformation followed by normalizing and reshaping to convert the output to convolutional filters. After obtaining the filters, we convolve them over the feature map obtained from the previous module (F_{i-1}) to relate expressions to image features:

$$G_i^F = \text{CONVOLVE}(K_i^F, F_{i-1}) \quad (2)$$

Then, the concatenation of these text-modulated features (G_i^F) for i th bottom-up layer and the previously generated features (F_{i-1}) is fed into module CNN_i for the next step.

In the top-down branch, we generate m feature maps starting from the final output of the contracting branch as:

$$G_i^H = \text{CONVOLVE}(K_i^H, F_i) \quad (3)$$

$$H_m = \text{DCNN}_i(G_m^H) \quad (4)$$

$$H_i = \text{DCNN}_i(G_i^H \oplus H_{i-1}) \quad (5)$$

Similar to the bottom-up branch, G_i^H is the modulated feature map with language-conditional filters defined as:

$$K_i^H = \text{AFFINE}_i^H(t_i) \quad (6)$$

where AFFINE_i^H is again an affine transformation followed by normalizing and reshaping for i th layer of the top-down branch. Here, we convolve the filter (K_i^H) over the feature maps from the contracting branch (F_i). Each upsampling module DCNN_i gets the concatenation (\oplus) of the text related features and the feature map (H_i) generated from the previous module. Only the first module operates on just convolved features. Each DCNN_i consists of a 2D deconvolution layer followed by a batch normalization and ReLU activation function. Final output H_1 becomes our joint feature map J representing the input image / language pair. The deconvolution layers have 5×5 filters with $\text{stride} = 2$ and $\text{padding} = 2$ doubling the spatial resolution, and they all have the same number of output channels.

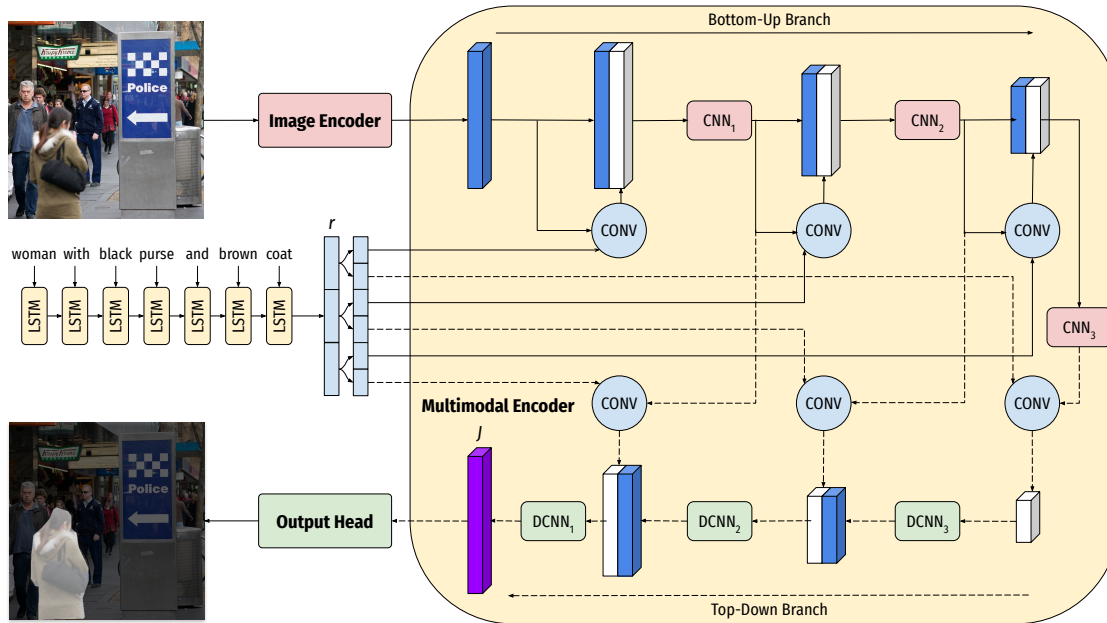


Figure 1. Overview of the proposed model on the task of referring expression segmentation.

3.4. Output Heads

As mentioned earlier, we develop our model as a generic solution which can be used to solve language-vision problems involving dense prediction. In this direction, we adapt our model to two different tasks by varying the output head: referring expression segmentation and language-guided image colorization.

Segmentation. In the referring expression segmentation problem, the goal is to generate segmentation mask for a given input image and language pair. After generating the joint feature map J , we apply a stack of layers (D_1, D_2, \dots, D_m) to map J to the exact image size. Similar to up-sampling modules, each D_k is a 2D deconvolution layer followed by batch normalization and the ReLU activation. Each D_k preserves the number of channels except for the last one which maps the features to a single channel for the mask prediction. We omit the batch normalization and the ReLU activation for the final module, instead we apply a sigmoid function to turn the final features into probabilities. Given these probabilities and ground-truth mask, we train our network by using binary cross entropy loss.

Colorization. In the language-guided image colorization task, the goal is to predict pixel color values for given input image with the guidance of language input. A convolutional layer by with 3×3 filters generates class scores for each spatial location of J . We apply bilinear upsampling to these predicted scores to match input image size. Given predicted scores and ground-truth color classes, we train our model by using a weighted cross entropy loss. To create compound LAB color classes and their weights, we follow the exactly

same process with [55, 89].

4. Experimental Analysis

This section contains the details of our experiments on referring expression segmentation (Section 4.1) and language-guided image colorization (Section 4.2) tasks¹.

4.1. Referring Expression Segmentation

Datasets. We evaluate our model on ReferIt (130.5k expressions, 19.9k images) [38], UNC (142k expressions, 20k images), UNC+ (141.5k expressions, 20k images) [87] and Google-Ref (G-Ref) (104.5k expressions, 26.7k images) [56] datasets. Unlike UNC, location-specific expressions are excluded in UNC+ through enforcing annotators to describe objects by their appearance. ReferIt, UNC, UNC+ datasets are collected through a two-player game and have short expressions (avg. 4 words). G-Ref have longer and richer expressions, its expressions are collected from Amazon Mechanical Turk instead of a two-player game. G-Ref does not contain a test split, and Nagaraja et al. [61] extends it by having separate splits for validation and test, which are denoted as val (U) and test (U).

Evaluation Metrics. Following previous work, we use intersection-over-union (IoU) and $p@X$ as evaluation metrics. Given the predicted segmentation mask and the ground truth, the IoU metric is the ratio between the intersection and the union of the two. There are two different ways to

¹We provide the implementation details, and present the complete ablation experiments in the supplementary material.

Method	UNC			UNC+			G-Ref			ReferIt
	val	testA	testB	val	testA	testB	val (G)	val (U)	test (U)	test
CMSA [85]	58.32	60.61	55.09	43.76	47.60	37.89	39.98	-	-	63.80
STEP [7]	60.04	63.46	57.97	48.19	52.33	40.41	46.40	-	-	64.13
BRINet [31]	61.35	63.37	59.57	48.57	52.87	42.13	48.04	-	-	63.46
CMPC [32]	61.36	64.53	59.64	49.56	53.44	43.23	49.05	-	-	65.53
LSCM [34]	61.47	64.99	59.55	49.34	53.12	43.50	48.05	-	-	66.57
EFN [20]	62.76	65.69	59.67	51.50	55.24	43.01	51.93	-	-	66.70
BUSNet [80]	63.27	66.41	61.39	51.76	56.87	44.13	50.56	-	-	-
Our model	64.63	67.76	61.03	51.76	56.77	43.80	50.88	52.12	52.94	66.01
MCN [†] [51]	62.44	64.20	59.71	50.62	54.99	44.69	-	49.22	49.40	-
CGAN [†] [50]	64.86	68.04	62.07	51.03	55.51	44.06	46.54	51.01	51.69	-
LTS [†] [37]	65.43	67.76	63.08	54.21	58.32	48.02	-	54.40	54.25	-
VLT [†] [19]	65.65	68.29	62.73	55.50	59.20	49.36	49.76	52.99	56.65	-
Our model [†]	67.01	69.63	63.45	55.34	60.72	47.11	53.51	55.09	55.31	57.09

Table 1. Comparison with the previous work by using the overall IoU metric. [†] denotes the corresponding method uses the mean IoU metric. "-" indicates that the model has not been evaluated on that dataset.

Method	$p@0.5$	$p@0.6$	$p@0.7$	$p@0.8$	$p@0.9$
CMSA	66.44	59.70	50.77	35.52	10.96
STEP	70.15	63.37	53.15	36.53	10.45
BRINet	71.83	65.05	55.64	39.36	11.21
LSCM	70.84	63.82	53.67	38.69	12.06
EFN	73.95	69.58	62.59	49.61	20.63
MCN	76.60	70.33	58.39	33.68	5.26
LTS	75.16	69.51	60.74	45.17	14.41
Our Model	76.67	71.77	64.76	51.69	22.73

Table 2. Comparison with the previous works on the val set of UNC dataset with $p@X$ metrics.

calculate IoU : the overall IoU calculates the total intersection over total union score throughout the entire dataset and the mean IoU calculates the mean of IoU scores of each individual example. For a fair comparison, we use both IoU metrics. The second metric, $p@X$, calculates the percentage of the examples that have IoU score higher than the threshold X .

Quantitative Results. Table 1 shows the comparison of our model with previous methods. Bold faces highlight the highest achieved scores. We evaluate our model by using both IoU metrics for a fair evaluation. Table 2 presents the comparison of our model with the state-of-the-art in terms of $p@X$. The difference between our model and the state-of-the-art increases as the threshold increases. This indicates that our model is better at segmenting the referred objects including smaller ones.

Qualitative Results. We visualize some of the segmentation predictions of our model to gain better insights about the trained model. Figure 2 shows some of the cases that our model segments correctly. These examples demon-

Method	Backbone	IoU
Top-down Baseline	ResNet-50	58.06
Bottom-up Baseline	ResNet-50	60.74
Our Model	ResNet-50	63.59
Our Model	ResNet-101	64.63

Table 3. Ablation study on the validation set of the UNC dataset with overall IoU metric.

strate that our model can learn a variety of language and visual reasoning patterns. For example, the first two examples of the fourth column show that our model learns to relate superlative adjectives (e.g., *taller*) with visual comparison. Examples including spatial prepositions (e.g., *near to*) demonstrate the spatial reasoning ability of the model. Our model can also learn a domain-specific nomenclature (e.g. *catcher*) that is present in the dataset. Lastly, we can see that the model can also detect certain actions (e.g., *sitting*).

Figure 3 shows failure cases from our model on the UNC test split. Examples in (a) show that our model tends to fail in case of typos. Our model segments the correct objects for these two examples when the typos are fixed (e.g. *pink* instead of *pick*). Examples in (b) show that some of the expressions are ambiguous, where there are multiple objects that could be referred to by the expression. In this case, the model seems to segment the most salient object. Some of the annotations contain incorrect or incomplete ground-truth mask (c). Finally, some of the examples (d) are hard to segment completely due to the lack of light or objects that mask the referred objects.

Ablation Study. We implemented 3 different models, the top-down baseline, the bottom-up baseline and our model,



Figure 2. Some predictions of our model on the UNC validation set.




Figure 3. Some incorrect predictions of our model on the UNC validation set. Each group (a-d) shows one pattern we observed within the predictions. The first row shows the ground truth mask, the second one is the prediction of our model.

to show the effect of modulating language in expanding and contracting visual branches. While the bottom-up baseline modulates language in bottom-up branch only, the top-down baseline modulates language in top-down branch only. Our model conditions language on both branches. Table 3 shows us the results. The bottom-up baseline outperforms the top-down one with ≈ 2.7 IoU improvement. Modulating language in both branches yields the best results by improving the bottom-up baseline with ≈ 2.85 IoU score.

Figure 4 visualizes the predictions of the different models on the same examples. The bottom-up baseline performs better when the description has color information as we show in the first three examples. The top-down-only

baseline also fails to detect object categories in some cases, and segments additional unwanted objects with similar category or appearance (e.g. *banana* vs. *orange*). Overall, our model which conditions both visual branches on language gives the best results.

Language-oriented Analysis. To analyze the effect of language on model performance, we divided UNC test splits into subsets depending on the different types of words (e.g. colors) and phrases (e.g. noun phrases with multiple adjectives) included in input expressions. Table 4 shows us the results of different models on these subsets. The first column stands for models, and the rest stand for different input expression categories. We exclude the categories which

Method	+C	-C	+IN	+IN,-C	+JJ*	+JJ*,-C
Top-down Baseline	52.59	59.59	54.84	55.66	57.66	61.16
Bottom-up Baseline	60.40	60.02	56.05	55.49	61.18	61.86
Our Model	62.98	63.57	59.60	59.27	64.45	65.56

Table 4. IoU performance of different setups depending on the input expression category for the UNC test splits.




Figure 4. Comparison of different architectural setups on the UNC test samples.

do not contribute to our analysis. We use DeepLab-v3+ ResNet-50 as visual backbone in each method. The notation of the categories are similar to Part of Speech (POS) tags [57], where we denote prepositions with IN, examples with adjectives with JJ*, and colors with C. Preceding plus and minus signs stand for inclusion and exclusion. For instance, +IN,-C column stands for the subset where each expression contains at least one preposition without any color words. Color words (e.g. red, darker) has the most impact on the performance in comparison to other types of words and phrases. Our model and the bottom-up baseline performs significantly better than top-down baseline on the subset that includes colors. In the opposite case, where input expressions with colors are excluded, the top-down baseline has performance similar to the bottom-up baseline, and our final model outperforms both single branch models. Since colors can be seen as low-level sensory information, low performance in the absence of the bottom-up branch can be expected. This demonstrates the importance of conditioning the bottom-up visual branch on language to capture low-level visual concepts.

Method	acc@1	acc@5	PSNR	LPIPS
FiLMed ResNet	23.70	60.50	-	-
FiLMed ResNet (ours)	20.22	49.57	20.89	0.1280
Top-down Baseline	22.83	51.85	21.29	0.1226
Bottom-up Baseline	21.85	51.34	20.98	0.1448
Our Model	23.38	54.27	21.42	0.1262
Our Model w/o balancing	33.74	67.83	22.75	0.1250

Table 5. Colorization results on the modified COCO validation split.

4.2. Language-guided Image Colorization

Datasets. Following the prior work [55], we use a modified version of COCO dataset [44] where the descriptions that do not contain color words are excluded. In this modified version, the training split has 66165, and the validation split has 32900 image / description pairs, and all images have a resolution of 224×224 pixels.

Evaluation Metrics. Following the previous work, we use pixel-level top-1 (acc@1) and top-5 accuracies (acc@5) in LAB color space, and additionally PSNR and LPIPS [90] in RGB for evaluation. A lower score is better for LPIPS, and a higher score is better for the rest.

Quantitative Results. We present the quantitative performance of our model in Table 5, and compare it with different design choices and previous work. FiLMed ResNet [55] uses FiLM [63] to perform language-conditional colorization. FiLMed ResNet (ours) denotes the results reproduced by the implementation provided by the authors. To show the effect of language modulation on different branches, we train 3 different models again: the top-down baseline, the bottom-up baseline and our model. We also re-train our model without class rebalancing and denote it as *Our Model w/o balancing*.

Contrary to the segmentation experiments, the top-down baseline performs better than the bottom-up baseline on the colorization task in all measures. Since, color information is absent in input images, bottom-up branch cannot encode low-level image features by modulating color-dependent language.

When we disable class rebalancing in the training phase, we observe a large improvement in acc@1 and acc@5 due to the imbalanced color distribution, where the model predicts the frequent colors exist in the backgrounds.

Qualitative Results. We visualize some of the colorization outputs of the trained models to analyze them in more de-

tail in Figure 5. FiLMed ResNet (ours) can understand all colorization hints, and it can manipulate object colors with some incorrectly predicted areas. The top-down baseline also performs similar to FiLMed ResNet (ours), where both models condition only the top-down branch on language.

In this task, since the models are blind to color, the bottom-up baseline loses its effectiveness to some degree, and starts to predict the most probable colors. This can be seen on the second and the last example, where the bottom-up baseline predicts red for the stop sign and blue for the sky. Although, the bottom-up baseline performs worse in this task, modulating the bottom-up branch with language still contributes to our final model to localize and recognize the objects present in the scene. This can be seen on the last two examples, where the top-down baseline mixes colors up in some object parts (e.g. the red parts in the motorcycle). Our model w/o balancing tends to predict more grayish colors (e.g. dog, sky).

Figure 6 highlights some of the failure cases we observed throughout the dataset. In the first two examples, our model is able to localize and recognize the target objects, but it fails to colorize them successfully by colorizing not only the targeted parts but also other parts. Models generally fail to colorize small objects since the data is imbalanced and it contains vast backgrounds and big objects frequently. The last two examples show that models fail to colorize reflective or transparent objects like glasses or water, these were also difficult in the language based segmentation task (see Figure 3 (d)).

5. Conclusion

In this work, we suggested that conditioning both top-down and bottom-up visual processing on language is beneficial for grounding language to vision. To support this claim, we proposed a generic architecture with explicit bottom-up and top-down visual branches for vision-language problems involving dense prediction. Our experiments on two different tasks demonstrated that conditioning both visual branches on language gives the best results. Our experiments on the referring expression segmentation task revealed that conditioning the bottom-up branch on language plays a vital role to process color-dependent input language. The language-guided image colorization experiments demonstrated similar conclusions, the bottom-up baseline failed to colorize the target objects since the color information is absent in the input images.

Limitations. We share common failure cases in Figure 3 and Figure 6. The performance of our model on both tasks decreases in the presence of transparent and/or reflective objects. Our model also fails to colorize small objects, mostly due to having an imbalanced color distribution. Finally, our current model is limited to integrated vision and language



Figure 5. Color manipulation performance of different models on COCO validation examples. Each column focuses on a different color conversion.



Figure 6. Failure cases of our model on the language-guided image colorization task.

tasks involving dense prediction, and we did not perform experiments on other vision and language problems.

Acknowledgements. This work was supported in part by an AI Fellowship to I. Kesen provided by the KUIS AI Center, GEBIP 2018 Award of the Turkish Academy of Sciences to E. Erdem, and BAGEP 2021 Award of the Science Academy to A. Erdem.

References

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, pages 6077–6086, 2018. 1
- [2] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton van den Hengel. Vision-and-Language Navigation: Interpreting Visually-Grounded Navigation Instructions in Real Environments. In *CVPR*, June 2018. 1
- [3] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. VQA: Visual Question Answering. In *ICCV*, Dec. 2015. 1
- [4] Hyojin Bahng, Seungjoo Yoo, Wonwoong Cho, David Keetae Park, Ziming Wu, Xiaojuan Ma, and Jaegul Choo. Coloring with words: Guiding image colorization through text-based palette generation. In *ECCV*, pages 431–447, 2018. 2
- [5] Paul Bloom. *How children learn the meanings of words*. MIT press, 2002. 1
- [6] Bastien Boutonnet and Gary Lupyan. Words jump-start vision: A label advantage in object recognition. *Journal of Neuroscience*, 35(25):9329–9335, 2015. 1
- [7] 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*, pages 7454–7463, 2019. 2, 5
- [8] Jianbo Chen, Yelong Shen, Jianfeng Gao, Jingjing Liu, and Xiaodong Liu. Language-based image editing with recurrent attentive models. In *CVPR*, pages 8721–8729, 2018. 2
- [9] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *TPAMI*, 40(4):834–848, 2017. 2
- [10] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, pages 801–818, 2018. 3
- [11] Y.-W. Chen, Y.-H. Tsai, T. Wang, Y.-Y. Lin, and M.-H. Yang. Referring Expression Object Segmentation with Caption-Aware Consistency. In *BMVC*, 2019. 2
- [12] Volkan Cirik, Taylor Berg-Kirkpatrick, and Louis-Philippe Morency. Using syntax to ground referring expressions in natural images. In *AAAI*, 2018. 2
- [13] Charles E Connor, Howard E Egeth, and Steven Yantis. Visual attention: bottom-up versus top-down. *Current biology*, 14(19):R850–R852, 2004. 1
- [14] Maurizio Corbetta and Gordon L Shulman. Control of goal-directed and stimulus-driven attention in the brain. *Nature reviews neuroscience*, 3(3):201–215, 2002. 1
- [15] Harm De Vries, Florian Strub, Jérémie Mary, Hugo Larochelle, Olivier Pietquin, and Aaron C Courville. Modulating early visual processing by language. In *NeurIPS*, pages 6594–6604, 2017. 1, 2
- [16] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 3
- [17] Jiajun Deng, Zhengyuan Yang, Tianlang Chen, Wengang Zhou, and Houqiang Li. Transvg: End-to-end visual grounding with transformers. In *ICCV*, pages 1769–1779, October 2021. 1, 2
- [18] Banchiamlack Dessalegn and Barbara Landau. More than meets the eye: The role of language in binding and maintaining feature conjunctions. *Psychological science*, 19(2):189–195, 2008. 1
- [19] Henghui Ding, Chang Liu, Suchen Wang, and Xudong Jiang. Vision-language transformer and query generation for referring segmentation. In *ICCV*, 2021. 1, 2, 5
- [20] Guang Feng, Zhiwei Hu, Lihe Zhang, and Huchuan Lu. Encoder fusion network with co-attention embedding for referring image segmentation. In *CVPR*, pages 15506–15515, 2021. 2, 5
- [21] Vittorio Gallese and George Lakoff. The brain’s concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive neuropsychology*, 22(3-4):455–479, 2005. 1
- [22] Peng Gao, Hongsheng Li, Shuang Li, Pan Lu, Yikang Li, Steven CH Hoi, and Xiaogang Wang. Question-guided hybrid convolution for visual question answering. In *ECCV*, pages 469–485, 2018. 2
- [23] Kirill Gavriljuk, Amir Ghodrati, Zhenyang Li, and Cees GM Snoek. Actor and action video segmentation from a sentence. In *CVPR*, pages 5958–5966, 2018. 2
- [24] Ross Girshick. Fast R-CNN. *ICCV*, Dec. 2015. 2
- [25] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *NeurIPS*, volume 27, 2014. 2
- [26] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask R-CNN. *ICCV*, Oct. 2017. 2
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 3
- [28] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 3
- [29] Ronghang Hu, Marcus Rohrbach, Jacob Andreas, Trevor Darrell, and Kate Saenko. Modeling Relationships in Referential Expressions with Compositional Modular Networks. *CVPR*, July 2017. 2
- [30] Ronghang Hu, Marcus Rohrbach, and Trevor Darrell. Segmentation from Natural Language Expressions. *Lecture Notes in Computer Science*, pages 108–124, 2016. 2
- [31] Zhiwei Hu, Guang Feng, Jiayu Sun, Lihe Zhang, and Huchuan Lu. Bi-Directional Relationship Inferring Network for Referring Image Segmentation. In *CVPR*, pages 4424–4433, 2020. 2, 5
- [32] 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*, pages 10488–10497, 2020. 2, 5

- [33] Drew A Hudson and Christopher D Manning. GQA: A New Dataset for Real-World Visual Reasoning and Compositional Question Answering. *CVPR*, 2019. 1
- [34] 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*, pages 59–75. Springer, 2020. 2, 5
- [35] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, pages 448–456. PMLR, 2015. 3
- [36] Ray Jackendoff and Ray S Jackendoff. *Foundations of language: Brain, meaning, grammar, evolution*. Oxford University Press, USA, 2002. 1
- [37] Ya Jing, Tao Kong, Wei Wang, Liang Wang, Lei Li, and Tieniu Tan. Locate then segment: A strong pipeline for referring image segmentation. In *CVPR*, pages 9858–9867, 2021. 2, 5
- [38] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *EMNLP*, pages 787–798, 2014. 4
- [39] Hyunsu Kim, Ho Young Jhoo, Eunhyeok Park, and Sungjoo Yoo. Tag2pix: Line art colorization using text tag with secat and changing loss. In *ICCV*, pages 9056–9065, 2019. 2
- [40] Jin-Hwa Kim, Sang-Woo Lee, Donghyun Kwak, Min-Oh Heo, Jeonghee Kim, Jung-Woo Ha, and Byoung-Tak Zhang. Multimodal residual learning for visual qa. In *NeurIPS*, pages 361–369, 2016. 1
- [41] 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
- [42] Zhenyang Li, Ran Tao, Efstratios Gavves, Cees GM Snoek, and Arnold WM Smeulders. Tracking by natural language specification. In *CVPR*, pages 6495–6503, 2017. 2
- [43] Yue Liao, Si Liu, Guanbin Li, Fei Wang, Yanjie Chen, Chen Qian, and Bo Li. A real-time cross-modality correlation filtering method for referring expression comprehension. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [44] 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*, pages 740–755, 2014. 7
- [45] Chenxi Liu, Zhe L. Lin, Xiaohui Shen, Jimei Yang, Xin Lu, and Alan L. Yuille. Recurrent Multimodal Interaction for Referring Image Segmentation. *ICCV*, pages 1280–1289, 2017. 2, 3
- [46] Daqing Liu, Hanwang Zhang, Feng Wu, and Zheng-Jun Zha. Learning to Assemble Neural Module Tree Networks for Visual Grounding. In *ICCV*, Oct. 2019. 2
- [47] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015. 2
- [48] Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *CVPR*, pages 375–383, 2017. 1
- [49] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. Hierarchical question-image co-attention for visual question answering. In *NeurIPS*, pages 289–297, 2016. 1
- [50] 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 *Proceedings of the 28th ACM International Conference on Multimedia*, pages 1274–1282, 2020. 2, 5
- [51] 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*, pages 10034–10043, 2020. 2, 5
- [52] Gary Lupyan and Andy Clark. Words and the world: Predictive coding and the language-perception-cognition interface. *Current Directions in Psychological Science*, 24(4):279–284, 2015. Publisher: Sage Publications Sage CA: Los Angeles, CA. 1
- [53] Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. Rectifier nonlinearities improve neural network acoustic models. In *ICML*, volume 30, page 3, 2013. 3
- [54] Mateusz Malinowski, Marcus Rohrbach, and Mario Fritz. Ask your neurons: A neural-based approach to answering questions about images. In *ICCV*, pages 1–9, 2015. 1
- [55] Varun Manjunatha, Mohit Iyyer, Jordan Boyd-Graber, and Larry Davis. Learning to color from language. In *NAACL-HLT, vol. 2*, pages 764–769, June 2018. 2, 4, 7
- [56] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan Yuille, and Kevin Murphy. Generation and Comprehension of Unambiguous Object Descriptions. *CVPR*, June 2016. 2, 4
- [57] Mitchell P Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. Building a large annotated corpus of english: the penn treebank. *Computational Linguistics*, 19(2):313–330, 1993. 7
- [58] Edgar Margffoy-Tuay, Juan C Pérez, Emilio Botero, and Pablo Arbeláez. Dynamic multimodal instance segmentation guided by natural language queries. In *ECCV*, pages 630–645, 2018. 2
- [59] Lotte Meteyard, Bahador Bahrami, and Gabriella Vigliocco. Motion detection and motion verbs: Language affects low-level visual perception. *Psychological Science*, 18(11):1007–1013, 2007. 1
- [60] Dipendra Misra, Andrew Bennett, Valts Blukis, Eyvind Niklasson, Max Shatkhin, and Yoav Artzi. Mapping instructions to actions in 3D environments with visual goal prediction. In *EMNLP*, pages 2667–2678, Oct.-Nov. 2018. 2, 3
- [61] Varun K. Nagaraja, Vlad I. Morariu, and Larry S. Davis. Modeling Context Between Objects for Referring Expression Understanding. *Lecture Notes in Computer Science*, pages 792–807, 2016. 2, 4
- [62] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543, 2014. 3
- [63] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron C. Courville. FiLM: Visual Reasoning with a General Conditioning Layer. In *AAAI*, 2018. 1, 2, 7
- [64] Friedemann Pulvermüller. Words in the brain’s language. *Behavioral and brain sciences*, 22(2):253–279, 1999. 1

- [65] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *TPAMI*, 39(6):1137–1149, June 2017. 2
- [66] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015. 1, 3
- [67] Hengcan Shi, Hongliang Li, Fanman Meng, and Qingbo Wu. Key-Word-Aware Network for Referring Expression Image Segmentation. In *ECCV*, Sept. 2018. 2
- [68] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun WOO. Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting. In *NeurIPS*, pages 802–810. 2015. 2
- [69] Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. A Corpus of Natural Language for Visual Reasoning. In *ACL vol. 2*, pages 217–223, Vancouver, Canada, July 2017. Association for Computational Linguistics. 1
- [70] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *EMNLP-IJCNLP*, pages 5100–5111, 2019. 1
- [71] Jan Theeuwes. Top-down and bottom-up control of visual selection. *Acta psychologica*, 135(2):77–99, 2010. 1
- [72] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, \textbackslashLukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017. 2
- [73] Gabriella Vigliocco, David P Vinson, William Lewis, and Merrill F Garrett. Representing the meanings of object and action words: The featural and unitary semantic space hypothesis. *Cognitive psychology*, 48(4):422–488, 2004. 1
- [74] Peng Wang, Qi Wu, Jiewei Cao, Chunhua Shen, Lianli Gao, and Anton van den Hengel. Neighbourhood Watch: Referring Expression Comprehension via Language-Guided Graph Attention Networks. *CVPR*, June 2019. 2
- [75] Terry Winograd. Understanding natural language. *Cognitive psychology*, 3(1):1–191, 1972. 1
- [76] Huijuan Xu and Kate Saenko. Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. In *ECCV*, pages 451–466. Springer, 2016. 1
- [77] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, pages 2048–2057, 2015. 1
- [78] Sibe Yang, Guanbin Li, and Yizhou Yu. Graph-structured referring expressions reasoning in the wild. 2020. 2
- [79] Sibe Yang, Guanbin Li, and Yizhou Yu. Propagating over phrase relations for one-stage visual grounding. In *ECCV*, pages 589–605. Springer, 2020. 2
- [80] Sibe Yang, Meng Xia, Guanbin Li, Hong-Yu Zhou, and Yizhou Yu. Bottom-up shift and reasoning for referring image segmentation. In *CVPR*, pages 11266–11275, 2021. 5
- [81] Zhengyuan Yang, Tianlang Chen, Liwei Wang, and Jiebo Luo. Improving one-stage visual grounding by recursive sub-query construction. In *ECCV*, pages 387–404. Springer, 2020. 2
- [82] Zhengyuan Yang, Boqing Gong, Liwei Wang, Wenbing Huang, Dong Yu, and Jiebo Luo. A fast and accurate one-stage approach to visual grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4683–4693, 2019. 2
- [83] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. Stacked Attention Networks for Image Question Answering. In *CVPR*, June 2016. 1
- [84] Linwei Ye, Zhi Liu, and Yang Wang. Dual convolutional lstm network for referring image segmentation. *TMM*, 22(12):3224–3235, 2020. 2
- [85] Linwei Ye, Mrigank Rochan, Zhi Liu, and Yang Wang. Cross-Modal Self-Attention Network for Referring Image Segmentation. *CVPR*, June 2019. 2, 3, 5
- [86] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L. Berg. MAttNet: Modular Attention Network for Referring Expression Comprehension. *CVPR*, June 2018. 2
- [87] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C. Berg, and Tamara L. Berg. Modeling Context in Referring Expressions. *Lecture Notes in Computer Science*, pages 69–85, 2016. 4
- [88] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From Recognition to Cognition: Visual Commonsense Reasoning. In *CVPR*, June 2019. 1
- [89] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, pages 649–666, 2016. 4
- [90] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 7
- [91] Changqing Zou, Haoran Mo, Chengying Gao, Ruofei Du, and Hongbo Fu. Language-based colorization of scene sketches. *TOG*, 38(6):1–16, 2019. Publisher: ACM New York, NY, USA. 2