

SGDiff: Scene Graph Guided Diffusion Model for Image Collaborative SegCaptioning

Xu Zhang¹, Jin Yuan^{1*}, Hanwang Zhang², Guojin Zhong¹, Yongsheng Zang¹,
Jiacheng Lin¹, Zhiyong Li^{1*}

¹Hunan University, China

²Nanyang Technological University, Singapore

{xuzhang1211, yuanjin, gjzhong, pishuai, jcheng_lin, zhiyong.li}@hnu.edu.cn, hanwangzhang@ntu.edu.sg

Abstract

Controllable image semantic understanding tasks, such as captioning or segmentation, necessitate users to input a prompt (e.g., text or bounding boxes) to predict a unique outcome, presenting challenges such as high-cost prompt input or limited information output. This paper introduces a new task “Image Collaborative Segmentation and Captioning” (SegCaptioning), which aims to translate a straightforward prompt, like a bounding box around an object, into diverse semantic interpretations represented by (caption, masks) pairs, allowing flexible result selection by users. This task poses significant challenges, including accurately capturing a user’s intention from a minimal prompt while simultaneously predicting multiple semantically aligned caption words and masks. Technically, we propose a novel Scene Graph Guided Diffusion Model that leverages structured scene graph features for correlated mask-caption prediction. Initially, we introduce a Prompt-Centric Scene Graph Adaptor to map a user’s prompt to a scene graph, effectively capturing his intention. Subsequently, we employ a diffusion process incorporating a Scene Graph Guided Bi-modal Transformer to predict correlated caption-mask pairs by uncovering intricate correlations between them. To ensure accurate alignment, we design a Multi-Entities Contrastive Learning loss to explicitly align visual and textual entities by considering inter-modal similarity, resulting in well-aligned caption-mask pairs. Extensive experiments conducted on two datasets demonstrate that SGDiff achieves superior performance in SegCaptioning, yielding promising results for both captioning and segmentation tasks with minimal prompt input.

Introduction

The rapid advancement of deep learning has markedly accelerated the evolution of image semantic understanding, targeting the extraction of vital semantic information from images through visual or textual representations (Li et al. 2017; Tang et al. 2020). Classic image semantic understanding tasks encompass vision-based object detection (Dai et al. 2021), image segmentation (Xie et al. 2021; Cheng et al. 2022), and image captioning (Li et al. 2022; Luo et al. 2023) etc., where a given image is translated to semantic representations by using boxes, masks, and captions, respectively.

*Corresponding author.

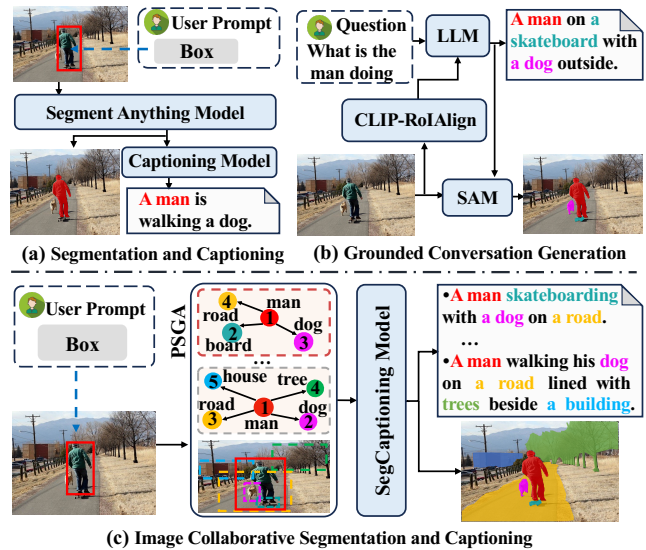


Figure 1: An example to compare different controllable image semantic understanding tasks, where our method uses a box to simultaneously predict multiple aligned bimodal results.

Recently, in pursuit of implementing controllable image semantic expression, researchers have explored enabling users to input prompts to articulate their demands, aiding models in accurately capturing the desired semantic expressions. For instance, referring image segmentation models (Wang et al. 2022) attempt to generate a referring mask based on a sentence. Similarly, controllable image captioning models (Cornia et al. 2019) can translate an image into a specific caption according to ordered bounding boxes. Building on the progress of large models, (Kirillov et al. 2023; Radford et al. 2021), advanced approaches have attempted to generate bimodal outcomes with richer semantic representations guided by prompts. For example, some methods predict a mask and a caption sequentially based on a bounding box (Wang et al. 2023; Huang et al. 2024) (see Figure 1a), while others generate a caption along with corresponding masks guided by a question (Rasheed et al. 2024) (see Figure 1b). While controllable input facilitates the accurate expression of user needs and encourages models to generate precise results, it also introduces several challenges. Firstly, gener-

ating prompt inputs (e.g., a sentence or multiple boxes) is often costly or time-consuming, significantly burdening users. Secondly, existing image semantic understanding methods typically support only unique result generation, which may not align with a user’s requirements. Thirdly, although the sequential generation of bimodal representations in Figures 1a and 1b offers richer semantic representations, it also suffers from severe error propagation.

To address these issues, we speculate that it may be feasible for users to input a simple prompt to generate multiple bimodal representations, and propose a novel task titled “Image Collaborative Segmentation and Captioning (SegCaptioning)”. This task concurrently generates intricate bimodal representations for an image, guided by a user’s prompt (see Figure 1c). Unlike methods that sequentially predicting bimodal outcomes, SegCaptioning produces coherent bimodal results (a caption alongside related masks) in parallel, leveraging their intricate correlations to mitigate error propagation and enhance performance. Moreover, our task simplifies user input requirements by only necessitating a basic prompt, such as a bounding box to denote the primary object, enabling the system to automatically infer various potential bimodal outcomes, thereby reducing input costs as well as enriching the semantics of outcomes. However, this multifaceted functionality presents notable technological challenges, such as accurately translating a simple prompt into precise intention representations and effectively exploring and aligning correlated bimodal representations.

Towards this end, this paper introduces a pioneering “Scene Graph Guided Diffusion Model” (SGDiff), which harnesses structured scene graphs to facilitate intention mapping from a basic prompt to multiple intention representation subgraphs. Additionally, it employs a diffusion process integrated with a bimodal transformer to deeply explore the intricate correlations between bimodal representations. Specifically, when provided with an image and a user’s prompt (such as a bounding box around an object), we introduce a Prompt-centric Scene Graph Adaptor (PSGA) to dynamically generate visual scene graphs. This facilitates the accurate capture of users’ intentions while effectively filtering out redundant information for caption-mask prediction. Building upon this, we devise a diffusion model featuring a Scene Graph Guided Bimodal Transformer (SGBTrans) to predict caption-mask pairs. The diffusion model initially introduces noise to a caption and subsequently denoises it using SGBTrans, which integrates both caption and scene graph features by exploring the complex correlations between them. To ensure precise alignment between correlated captions and masks, we devise a Multi-Entities Contrastive Learning Loss (MECL), which explicitly aligns the generated captions and masks by pulling feature representations of positive words and masks closer while pushing away those of negative ones. The MECL operates on both inter-sample and intra-sample levels, yielding well-aligned bimodal results. Extensive evaluations conducted on MSCOCO and Flickr30k entities datasets validate the effectiveness of SGDiff, achieving state-of-the-art performance compared to existing methods. In summary, the main contributions of this work can be outlined as follows:

- We introduce a novel image semantic understanding task

called “Image Collaborative Segmentation and Captioning (SegCaptioning)”. This task enables users to input a straightforward prompt to predict semantic-aligned caption and mask pairs, empowering users to generate richer semantic representations based on their intentions.

- We propose a novel SGDiff for this new task, facilitating simultaneous generation of (caption, masks) pairs based on a simple prompt, which is translated into a structured scene graph by our proposed PSGA, effectively capturing users’ intentions and aiding our model in accurately generating the desired caption-mask pairs.
- We design SGBTrans to delve into the intricate correlations between scene graphs and captions throughout the denoising. Additionally, our MECL explicitly aligns each mask and caption word, facilitating effective instance alignment.

Related Work

Image Captioning

Image captioning translates visual content within a scene into a linguistic interpretation, effectively bridging the gap between computer vision and natural language understanding (Gan et al. 2017). Early efforts (Yang et al. 2019; Pan et al. 2020) focused on describing visual content from a holistic perspective, resulting in generating limited and fixed captions for any given image. Driven by practical demands, researchers have shifted towards exploring caption generation with local or specific semantics. This shift leads to the development of Controllable Image Captioning (CIC) (Deshpande et al. 2019; Wang et al. 2024), which utilizes control signals such as clicks (Pont-Tuset et al. 2020), bounding boxes (Cornia et al. 2019; Chen et al. 2021a) and masks (Guo et al. 2024) to generate controllable captions, including semantic controllability (Yan et al. 2021), length controllability (Deng et al. 2020), and sentiment controllability (Deshpande et al. 2019; Zhao, Wu, and Zhang 2020) etc. For instance, (Cornia et al. 2019) proposes generating captions based on an ordered sequence of bounding boxes. Meanwhile, (Yang et al. 2019) and (Zhong et al. 2020) employ scene graphs as input to generate image captions, which incorporate both target nodes and relationships to effectively capture users’ intentions. However, the above autoregressive captioning methods generate text tokens step-by-step, which usually suffers from severe error propagation. Consequently, recent work has explored semi-autoregressive (Zhou et al. 2021) and non-autoregressive diffusion methods (Guo et al. 2020; Li et al. 2022; Luo et al. 2023) to enable parallel token output. Recently, with the advancement of large models (Li et al. 2023), CAT (Wang et al. 2023) introduces a universal image captioning approach that integrates various types of prompts (clicks, bounding boxes, scribbles). It utilizes the Segment Anything Model (SAM) (Kirillov et al. 2023) to first extract local visual features conditioned on visual prompts. These visual features are then input into GPT-3 (Brown et al. 2020) for caption generation. Instead of adopting the training-free scheme used by CAT, SCA (Huang et al. 2024) introduces a learned module, a “query-based feature mixer”, to effectively align region-specific visual features with the embedding space of language

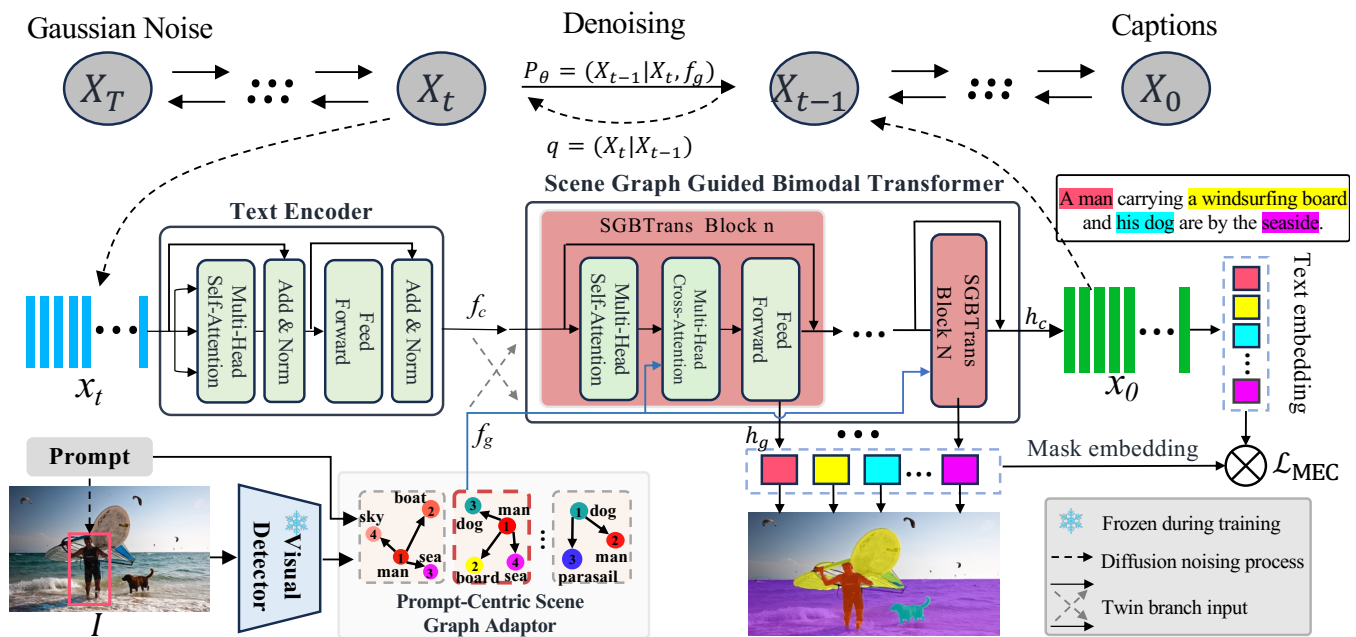


Figure 2: An example to illustrate the training process of our Scene Graph Guided Diffusion Model.

models for caption generation, yielding improved captioning performance.

Image Segmentation

Traditional image segmentation aims to localize accurate masks concerning semantic objects in an image (Kirillov et al. 2023). To generate controllable results, researchers have introduced various user prompts (Bai and Wu 2014), broadly categorized into two types: visual prompts, like scribbles (Wang et al. 2018), boxes (Zhang et al. 2020), and clicks (Chen et al. 2021b; Lin et al. 2023); and textual prompts (Wang et al. 2022). For instance, (Lin et al. 2020) optimizes the use of initial clicks to enhance outcomes, while (Zou et al. 2024) uses scribbles to provide richer and more precise information to capture users’ intentions, albeit at a higher labeling cost. Compared to visual prompts, language prompts can increase user burden but more accurately express users’ intentions, thereby helping models better predict referring masks (Luo et al. 2020; Liu, Ding, and Jiang 2023). Inspired by the complementarity of various prompts, SAM (Kirillov et al. 2023) first unifies various types of user prompts—including a set of clicks, a region box/mask, or free-form text—to establish a versatile foundational model for segmentation. This offers a flexible prompt interface for users. Additionally, it is the first visual large model for image segmentation. Building on this foundation, GROUNDHOG (Zhang et al. 2024) leverages SAM to obtain region mask proposals through visual prompts and then uses Multimodal Large Language Models (MLLMs) to generate grounded captions based on regional semantic features. GLaMM (Rasheed et al. 2024) aligns visual regions with text tokens using SAM and GPT-4 (OpenAI 2023), first generating a caption by querying GPT-4 and then producing corresponding region masks with SAM guided

by the caption. Although the bimodal outcomes enrich the content representations, the sequential predictions usually suffer from severe error propagation.

Different from (Yang et al. 2022; Zou et al. 2023; Ren et al. 2024; Huang et al. 2024), our approach only requires a simple visual prompt to infer user needs. This prompt guides the generation of multiple bimodal outcomes, enriching the output content. Moreover, our method predicts both masks and captions in parallel, alleviating the error propagation issues associated with the sequential prediction methods.

Methodology

Overview

Given an image $I \in \mathbb{R}^{H \times W \times 3}$ with height H and width W , along with a user’s prompt P_o such as a box bounding an object o , SegCaptioning aims to generate multiple (caption, masks) pairs. Here, we use S to denote a caption centered around o , and M to represent several masks related with S . To achieve this goal, we design a novel “Scene Graph Guided Diffusion” (SGDiff) model tailored for this specific task. As illustrated in Figure 2, our proposed model comprises three key components: a textual encoder that extracts textual features for captions, a Prompt-Centric Scene Graph Adaptor (PSGA) that transforms a prompt-guided image to a scene graph for effective intention capturing, and a diffusion model with a Scene Graph Guided Bimodal Transformer (SGBT) that predicts caption-mask pairs. Furthermore, our model introduces a MECL to explicitly align masks with their corresponding captions. This ensures that the bimodal results are well-aligned and coherent, enhancing the quality and relevance of the generated content.

Prompt-Centric Scene Graph Adaptor

Given an image I along with a caption-mask pair (S, M) , the proposed PSGA receives a user’s prompt and then transforms I into a scene graph G_i to encapsulate the caption-mask pair. Specifically, our approach first generates a global scene graph G following (Zellers et al. 2018). Subsequently, it predicts a subgraph G_i within G according to a user’s prompt, facilitating the subsequent prediction on (S, M) . Concretely, when a user inputs a prompt P_o referring to the center object o in (S, M) , PSGA aims to search for a subgraph G_i from G to represent (S, M) , which is mathematically expressed as:

$$\arg \max_{G_i} P(G_i = (N, E) | G, o). \quad (1)$$

To implement this, our approach initially generates a coarse subgraph $G_i = (N, E)$, where N and E denote the node and edge sets connected to o , respectively. As the numbers of nodes and edges in G_i typically exceed the desired ones in (S, M) , our approach further constructs an adaptor to refine G_i . Specifically, the adaptor comprises several self-attention blocks and one mapping layer. Initially, we feed the node feature $e_o \in \mathbb{R}^{L \times D}$ of G_i into the self-attention blocks to produce an updated feature $f_o \in \mathbb{R}^{L \times D}$, where L is the number of nodes and D is the feature dimension. Subsequently, the mapping layer predicts node relevance scores $P(G_i) \in \mathbb{R}^L$, representing the node relevance concerning (S, M) . We train the adaptor using the Binary Cross-Entropy (BCE) loss to minimize $P(G_i)$ and its ground truth $\hat{P}(G_i)$:

$$\mathcal{L}_{\text{adaptor}} = \text{BCE}(P(G_i), \hat{P}(G_i)). \quad (2)$$

For each object in (S, M) , we compute its Intersection over Union (IOU) with the corresponding node in G_i . If the IOU exceeds a threshold, we assign the score in $\hat{P}(G_i)$ as 1, otherwise 0. Consequently, the adaptor is able to effectively predict relevant objects in G_i while filtering out irrelevant ones.

Furthermore, the random order of objects in the node features f_o may not align with the appearance order in the caption S , potentially impacting caption generation, as discussed in (Cornia et al. 2019). Hence, we consider measuring the order consistency between the node feature and the caption. To address this, we employ a multi-layer fully connected network to map f_o into a permutation matrix $R \in \mathbb{R}^{L \times K}$, where the L rows represent the L objects in f_o with a random order, and the K columns record the ordered objects from the caption S . The element r_{lk} indicates the probability of the l -th object in f_o semantically consistent with the k -th object in S . We design a ranking loss $\mathcal{L}_{\text{ranking}}$ to minimize the distance between the predicted matrix R and the ground truth matrix \hat{R} using a Cross-Entropy (CE) loss:

$$\mathcal{L}_{\text{ranking}} = \text{CE}(R, \hat{R}). \quad (3)$$

After filtering out irrelevant nodes, we adjust the node order to match that of the caption, resulting in a refined feature f'_o :

$$f'_o = \text{Ranking}(f_o \times \delta(P(G_i))), \quad (4)$$

$$\delta(x) = \begin{cases} x, & \text{if } x \geq \theta \\ 0, & \text{otherwise,} \end{cases} \quad (5)$$

where θ is a threshold to filter irrelevant objects. $\text{Ranking}(\cdot)$ rearranges the row vectors in f_o so that its feature order is consistent with the object order in the caption. On this basis, our approach fuses f'_o and the edge feature f_e following (Guo et al. 2019), generating the graph feature f_g .

The PSGA is trained using a composite loss $\mathcal{L}_{\text{SGadaptor}} = \mathcal{L}_{\text{adaptor}} + \mathcal{L}_{\text{ranking}}$, offering several advantages: Firstly, it effectively maps a user’s prompt to a concise and accurate subgraph consistent with the to-be-predicted caption, retaining only relevant information to facilitate subsequent predictions. Secondly, the generated graph feature maintains a consistent object order with the caption for performance improvement. In the testing phrase, we use the PSGA to generate several related subgraphs to predict multiple (caption, masks) pairs.

Scene Graph Guided Diffusion Model

Similar to Bit Diffusion (Chen, Zhang, and Hinton 2022), our SGDDiff adds Gaussian noise into a caption sample \mathbf{x} from timestep $t = 0$ to T during the forward process and then performs a reverse denoising process to recover the caption. Differently, we introduce a scene graph feature f_g to strengthen the denoising to simultaneously predict a caption with related masks. Specifically, given a noised sample \mathbf{x}_t and f_g , SGDDiff iteratively generates noise-free text descriptions x_0 by progressively denoising x_t , and each denoising transition $x_t \rightarrow x_{t-1}$ is parametrized by a model $p_\theta(x_{t-1} | x_t, f_g)$:

$$p_\theta(\mathbf{x}_{0:T}) := p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t, f_g). \quad (6)$$

We design a scene graph guided bimodal transformer to implement Eq. (6) with a training loss:

$$\mathcal{L}_{\text{bit}} = \mathbb{E}_{t \sim \mathcal{U}(0, T), \epsilon \sim \mathcal{N}(0, I)} \|f(x_t, t, f_g) - x_0\|^2, \quad (7)$$

where $t \sim \mathcal{U}(0, T)$, $\epsilon \sim \mathcal{N}(0, I)$ are continuous variables.

Scene Graph Guided Bimodal Transformer. At each time step t , SGBTrans first receives a noised caption feature f_c and the graph feature f_g , and then predicts bimodal features including a denoised text feature and a visual feature. Specifically, SGBTrans comprises B blocks, each containing a multi-head self-attention layer, a multi-head cross-attention layer, and a feed-forward layer. Unlike uni-modal generation transformers, our task involves predicting a caption with masks. Therefore, we devise a cross-training scheme for bimodal generation. To elaborate, for caption generation, we utilize f_c as a query and f_g as value and key into the blocks:

$$\begin{aligned} f'_c &= \text{MultiHead}_s(f_c, f_c, f_c), \\ u_c &= \text{FFN}(f'_c + \text{MultiHead}_c(f'_c, f_g, f_g)), \\ h_c &= \phi_c(u_c + f_c). \end{aligned} \quad (8)$$

where MultiHead_s and MultiHead_c represent the multi-head self-attention and cross-attention layers, respectively, FFN denotes the feed-forward layer, and $\phi(\cdot)$ indicates the mlp layer. The resulting h_c represents a refined caption feature, which is then either passed to the subsequent block or utilized for caption prediction. Compared to the global visual feature, the scene graph feature provides a more precise visual description, aligning better with the caption to be predicted and thereby enhancing caption generation. Regarding

mask generation, we input f_g as a query and f_c as both value and key into the blocks as follows:

$$\begin{aligned} f'_g &= \text{MultiHead}_s(f_g, f_g, f_g), \\ u_g &= \text{FFN}(f'_g + \text{MultiHead}_c(f'_g, f_c, f_c)), \\ h_g &= \phi_g(u_g + f_g). \end{aligned} \quad (9)$$

Guided by the caption feature, h_g could accurately offer visual features for mask prediction. Unlike (Yang et al. 2022; Sun et al. 2023), Our method distinguishes itself by using PSGA to convert a visual prompt into object-centered scene graphs, guiding the generation of multiple bimodal outputs.

Caption and Mask Generation. We utilize the denoised caption feature h_c to predict the probability distribution for each word in the caption. Further details can be found in (Luo et al. 2023), where the loss $\mathcal{L}_{\text{caption}}$ contains a diffusion loss in Eq. (7) and a cross-entropy loss \mathcal{L}_{CE} with equal weights. For mask generation, we adopt the Mask2Former (Cheng et al. 2022), which takes the global visual feature f_v extracted by Swin-Transformer (Liu et al. 2021) and the graph h_g in Eq. (9) as input to produce class-agnostic binary region masks. The mask loss L_{Mask} employs both Dice loss and BCE loss.

Multi-Entities Contrastive Learning

We propose a MECL to explicitly explore cross-modal representation correlations between visual entities and caption words. Our motivation is that matching mask-word pairs should have similar contextual representations. Given a predicted mask-caption pair (M, S) , it's important to note that the relationship between masks and words is not strictly one-to-one. For instance, the word *dog* may refer to multiple mask entities, and conversely, a mask may contain objects indicated by multiple words. Therefore, the i -th mask embedding \mathbf{m}_i is expected to be close to the matched word embedding sets \mathbf{s}_i^+ while being distant from other word embeddings \mathbf{s}_j . In other words, the mutual information between mask entities and words that are matched should be maximized. Since directly optimizing mutual information is infeasible, we design an intra-sample contrastive loss $\mathcal{L}_{\text{intra}}$ inspired by InfoNCE (Oord, Li, and Vinyals 2018), representing a lower bound of mutual information. This loss is formulated as follows:

$$\begin{aligned} \mathcal{L}_{\text{intra}} &= - \sum_{i=1}^N \frac{1}{|\{\mathbf{s}_i\}^+|} \sum_{\mathbf{s}_i \in \{\mathbf{s}_i\}^+} \log \frac{\exp(\mathbf{m}_i^\top \mathbf{s}_i / \tau)}{\sum_{j \neq i} \exp(\mathbf{m}_i^\top \mathbf{s}_j / \tau)} \\ &\quad - \sum_{i=1}^L \frac{1}{|\{\mathbf{m}_i\}^+|} \sum_{\mathbf{m}_i \in \{\mathbf{m}_i\}^+} \log \frac{\exp(\mathbf{s}_i^\top \mathbf{m}_i / \tau)}{\sum_{j \neq i} \exp(\mathbf{s}_i^\top \mathbf{m}_j / \tau)}, \end{aligned} \quad (10)$$

where τ is a learnable parameter, N and L are the number of mask and word tokens. Likewise, we impose that the i -th word embedding \mathbf{s}_i is similar to the matched mask embedding sets $\{\mathbf{m}_i\}^+$ and dissimilar to the irrelevant \mathbf{m}_j .

Aligning mask and word embeddings solely at the intra-sample level may be insufficient as it presents challenges in exploring the discriminability of visual and textual representations across different samples. Consequently, we propose inter-sample contrastive learning to capture the discriminative information of multiple mask-caption pairs. Specifically, we define the global matching score for each mask-caption pair $g(\mathbf{m}, \mathbf{s})$ as the weighted average of their mask-word pair

scores. This allows us to consider the relationships between masks and words across multiple samples:

$$g(\mathbf{m}, \mathbf{s}) = \frac{1}{L} \sum_{j=1}^L \sum_{i=1}^N \rho(\mathbf{m}, \mathbf{s})_{i,j} \cdot \langle \mathbf{m}_i, \mathbf{s}_j \rangle, \quad (11)$$

where $\langle \cdot, \cdot \rangle$ is the dot product of two vectors and $\rho(\mathbf{m}, \mathbf{s})_{i,j} = \frac{\exp(\mathbf{m}_i, \mathbf{s}_j)}{\sum_{i'=1}^N \exp(\mathbf{m}_{i'}, \mathbf{s}_j)}$. Then the inter-sample contrastive loss $\mathcal{L}_{\text{inter}}$ can be computed by the random pairs between mask and caption in the same minibatch B as follows:

$$\begin{aligned} \mathcal{L}_{\text{inter}} &= - \frac{1}{B} \sum_{i=1}^B \log \frac{\exp(g(\mathbf{m}^i, \mathbf{s}^i) / \tau)}{\sum_{j \neq i} \exp(g(\mathbf{m}^i, \mathbf{s}^j) / \tau)} \\ &\quad - \frac{1}{B} \sum_{i=1}^B \log \frac{\exp(g(\mathbf{s}^i, \mathbf{m}^i) / \tau)}{\sum_{j \neq i} \exp(g(\mathbf{s}^i, \mathbf{m}^j) / \tau)}, \end{aligned} \quad (12)$$

where the matched mask-caption pairs with the same index will be pulled closer in the shared semantic representation space and $\mathcal{L}_{\text{inter}}$ can be viewed as an optimization of the global mutual information between visual entities and caption. Toward this objective, the MECL can be expressed as $\mathcal{L}_{\text{MEC}} = \mathcal{L}_{\text{inter}} + \mathcal{L}_{\text{intra}}$, designed to complement representation learning and further facilitate joint multi-modal learning within the diffusion process. Finally, it's worth noting that we train the entire SGDiff model using either $\mathcal{L}_{\text{Caption}}$ or $\mathcal{L}_{\text{Mask}}$, alongside the proposed \mathcal{L}_{MEC} and $\mathcal{L}_{\text{SGadaptor}}$ with balancing weights $\lambda_{\{1,2\}}$. The total loss is calculated as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{Caption}} + \mathcal{L}_{\text{Mask}} + \lambda_1 \mathcal{L}_{\text{SGadaptor}} + \lambda_2 \mathcal{L}_{\text{MEC}}. \quad (13)$$

Experiments

Datasets and Metrics

MSCOCO (Lin et al. 2014) comprises 123, 287 images, each annotated with five captions. For each caption, we utilize all the nouns to identify the corresponding masks provided in (Lin et al. 2014), resulting in a (caption, masks) pair for each image across 133 categories. We adhere to the Karpathy split (Karpathy and Fei-Fei 2015) to allocate 113, 287 images for training, 5, 000 for validation, and 5, 000 for testing.

Flickr30k Entities (Plummer et al. 2015) is an extension of Flickr30k (Young et al. 2014), consisting of 31, 000 images with five captions for each one. For each caption, noun entities are linked to one or more boxes. We employ SAM (Kirillov et al. 2023) to approximately convert each box into a binary mask, thereby matching each caption with several masks. We follow the split suggested by Karpathy (Cornia et al. 2019), designating 29, 000 images for training, and 1, 000 for validation and testing, respectively.

We evaluate our model on three aspects: caption quality, mask quality, and caption-mask matching. Caption quality is assessed using Blue-4, METEOR, SPICE and CIDEr metrics. Mask quality is evaluated with Mask IoU. Caption-mask matching is measured by calculating the mAP of mask classification for objects in the caption. Since our method supports the generation of multiple results, we select the top 5 results and calculate the highest captioning score as the final result for comparison.

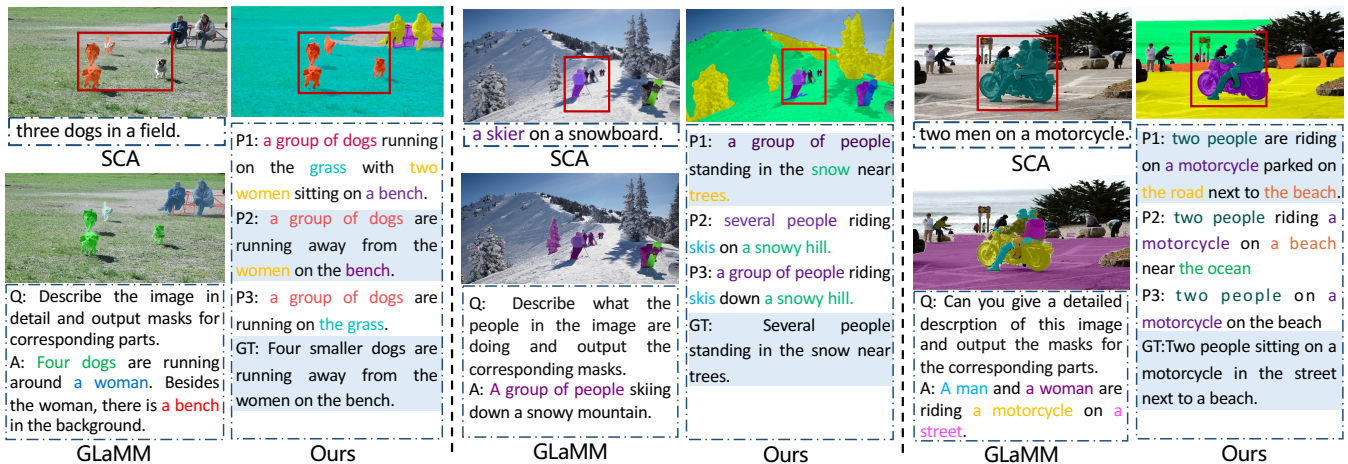


Figure 3: Quantitative results of our SGDiff, SCA, and GLaMM on MSCOCO, where color indicates the matching relationship between masks and words.

Model	SEG	CAP	SPICE	CIDEr	mIoU	mAP
Sub-GC <small>ECCV20</small>	-	✓	20.2	115.3	-	-
SATIC <small>ICCV21</small>	-	✓	22.3	127.2	-	-
X-Transformer <small>CVPR20</small>	-	✓	23.4	132.8	-	-
CMAL <small>IJCAI20</small>	-	✓	21.8	124.0	-	-
SCD-Net <small>CVPR23</small>	-	✓	23.0	131.6	-	-
OpenSeg <small>ECCV22</small>	✓	-	-	-	38.1	-
OpenSeeD <small>ICCV23</small>	✓	-	-	-	64.0	47.6
ODISE <small>CVPR23</small>	✓	-	-	-	65.2	46.0
X-Decoder <small>CVPR23</small>	✓	✓	-	122.3	62.4	41.3
CAT <small>arXiv23</small>	✓	✓	17.9	90.1	24.9	8.2
SCA <small>CVPR24</small>	✓	✓	18.8	85.0	32.4	11.3
GLaMM <small>CVPR24</small>	✓	✓	23.3	130.9	36.2	14.8
SGDiff (ours)	✓	✓	26.1	137.4	66.3	47.2

Table 1: Performance comparison with several advanced methods on MSCOCO. SEG and CAP denote Segmentation and Captioning, respectively.

Implementation Details

To train our model, we employ large-scale jittering on each input image with a size of 1024^2 . We approximately use the first noun in each caption as the central object, and its corresponding bounding box as a user’s prompt. Our training consists of two stages. First, we train the model for captioning by deleting $\mathcal{L}_{\text{Mask}}$ and \mathcal{L}_{MEC} in Eq. (13), where the SGDiff network is optimized using the Adam optimizer with a learning rate of 0.0001 and a weight decay of 0.05 on two A6000 GPUs. The optimization process spans 60 epochs with a batch size of 16. Second, we jointly train the mask and caption generation with the complete loss. The optimization process spans 60 epochs with a batch size of 16.

Comparison with State-of-the-Art Methods

This is the first study to support generating multiple caption-mask pairs. We benchmark SGDiff against advanced methods for generating masks, captions, or both, as shown in Tables 1

Model	Bleu-4	METEOR	SPICE	CIDEr
SCT <small>CVPR19</small>	12.5	16.8	23.5	84.0
Sub-GC <small>ECCV20</small>	11.2	15.9	22.2	73.7
VSR <small>CVPR21</small>	10.7	18.0	21.9	97.5
ComPro <small>IJCV24</small>	11.9	17.3	23.9	89.4
CAT <small>arXiv23</small>	30.2	22.7	15.6	63.4
SCA <small>CVPR24</small>	20.8	17.9	13.4	35.4
GLaMM <small>CVPR24</small>	35.4	26.4	18.8	95.3
GROUNDHOG <small>CVPR24</small>	36.7	26.5	20.4	91.3
SGDiff (ours)	38.2	27.4	24.5	98.6

Table 2: Performance comparison with several advanced methods on Flickr30k Entities.

and 2. Compared to captioning methods like Sub-GC, SCA, and CAT, SGDiff outperforms them due to its precise capture of user intentions through the PSGA and its effective prediction capabilities enabled by the scene graph-guided diffusion process. For segmentation, we compare SGDiff with several open-vocabulary approaches (OpenSeg, OpenSeeD, and ODISE) on the MSCOCO dataset. Our SGDiff excels due to its thorough exploration and strict alignment between masks and words. It also surpasses large models like CAT, SCA, and GLaMM in both tasks. Notably, SGDiff facilitates collaborative prediction in both mask and caption generation, which significantly mitigates error propagation and boosts performance. Figure 3 shows quantitative results comparing SGDiff, SAC, and GLaMM on MSCOCO. SAC is limited to generating a text description with masks confined to a bounding box and lacks the capability to capture semantic and visual representations outside of it, resulting in a marked performance decline (as shown in Table 1). GLaMM, which requires users to type a sentence, shows relatively low segmentation performance since its SAM struggles to accurately predict masks guided by text. In contrast, our SGDiff supports the generation of multiple bimodal results, offering diverse yet precise outcomes guided by a simple prompt, thus

Filtering	Ranking	SPICE	CIDEr	mIOU	mAP
-	-	24.2	127.4	62.3	43.2
✓	-	24.9	133.7	64.2	45.4
-	✓	25.4	135.8	65.5	46.3
✓	✓	26.1	137.4	66.3	47.2

Table 3: Performance comparison of PSGA by using different components.

λ_1	SPICE	CIDEr	mIOU	mAP
0	24.2	127.4	62.3	43.2
1	24.9	135.7	65.7	46.0
2	26.1	137.4	66.3	47.2
5	25.8	137.6	65.9	46.6
10	25.3	136.5	65.4	45.8

Table 4: Performance change of the loss $\mathcal{L}_{\text{SGadaptor}}$ with different weights on MSCOCO.

SGB-Trans	SPICE	CIDEr	mIOU	mAP
without cross-attention	24.5	133.7	64.5	46.0
with cross-attention	26.1	137.4	66.3	47.2

Table 5: Performance comparison by using SGBTrans with or without cross-attention on MSCOCO.

$\mathcal{L}_{\text{Mask}}$	$\mathcal{L}_{\text{Caption}}$	SPICE	CIDEr	mIOU	mAP
-	✓	25.2	135.8	-	-
✓	-	-	-	63.7	44.5
✓	✓	26.1	137.4	66.3	47.2

Table 6: Performance comparison of our SGDiff with individual tasks on MSCOCO.

reducing user burden and enhancing overall performance.

Analysis of PSGA

Our PSGA operates on a coarse subgraph by performing node filtering (see Eq. (5)) and ranking (see Eq. (4)). This experiment confirms the effectiveness of these two components, as illustrated in Table 3. Node filtering removes irrelevant objects, while node ranking ensures a consistent order between graph nodes and caption words, both significantly improving performance. Table 4 details the performance changes with respect to different weights λ_1 in Eq. (13) under $\lambda_2 = 1$. When $\lambda_1 = 0$, ignoring PSGA training, performance is lowest. Increasing λ_1 from 0 to 2 consistently improves performance, demonstrating the effectiveness of PSGA. However, larger λ_1 values lead to performance declines, as an excessively large value may diminish the utility of other losses.

Analysis of SGBTrans

Our SGBTrans takes a noised caption feature and a scene graph feature, facilitating deep interaction for caption-mask prediction. This experiment compares the performance of SGBTrans with and without the cross-attention layer (see Eq. (8) and Eq. (9)), as shown in Table 5. Incorporating the cross-

L_{intra}	L_{inter}	SPICE	CIDEr	mIOU	mAP
-	-	24.8	134.8	64.2	42.7
✓	-	25.4	136.4	65.6	46.2
-	✓	25.6	136.7	65.3	46.0
✓	✓	26.1	137.4	66.3	47.2

Table 7: Performance comparison of the \mathcal{L}_{MEC} with different components on MSCOCO.

λ_2	SPICE	CIDEr	mIOU	mAP
0	24.8	134.8	64.2	42.7
1	26.1	137.4	66.3	47.2
2	25.6	137.1	66.1	47.0
5	25.4	136.8	65.8	46.4
10	25.0	136.2	65.5	46.1

Table 8: Performance comparison of the \mathcal{L}_{MEC} with different weights on MSCOCO.

attention operation between caption and graph features significantly enhances performance since the cross-modal interaction compensates for the deficiencies of each feature, thereby guiding the generation of accurate captions and masks.

Analysis of Collaborative SegCaptioning

This experiment verifies the effectiveness of collaborative segmentation and captioning, as shown in Table 6. Our SGDiff employs $\mathcal{L}_{\text{Mask}}$, $\mathcal{L}_{\text{Caption}}$, or both. It is evident that collaborative training of both tasks achieves superior performance by effectively exploring the intricate correlations between captions and masks. Furthermore, our approach introduces the Multi-Entities Contrastive Loss (\mathcal{L}_{MEC}), comprising $\mathcal{L}_{\text{intra}}$ and $\mathcal{L}_{\text{inter}}$ for mask-caption alignment, as shown in Table 7. Both components enhance performance, particularly mAP, by aligning masks and words at intra-sample and inter-sample levels. Table 8 further demonstrates the overall effectiveness of \mathcal{L}_{MEC} with different weights in Eq. (13), where the optimal performance is achieved when $\lambda_2 = 1$. Increasing λ_2 beyond this point results in a performance decline, as it tends to diminish the utilities of the other losses.

Conclusion

This paper introduces a novel ‘‘Scene Graph Guided Diffusion Model’’ for the newly defined task of ‘‘Image Collaborative SegCaptioning’’. It represents the first exploration of this topic, providing valuable insights and guidance for future research. Technically, we develop an innovative PSGA that automatically translates a simple prompt into multiple related scene graphs, effectively inferring a user’s intention. Building on this, we establish an effective diffusion pipeline that integrates a SGBTrans to explore the complex interplay between scene graph and text inputs, as well as a MECL to explicitly align each mask with its corresponding caption word, producing diverse yet precise caption-mask pairs that meet user requirements. Extensive experiments on two benchmark datasets demonstrate the effectiveness of our approach.

Acknowledgments

The authors would like to thank the reviewers for their comments that help improve the manuscript. This research is supported by the National Natural Science Foundation of China (Grants No. 62272157, U21A20518, and U23A20341), the RIE2025 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) (Award I2301E0026) administered by A*STAR, and by Alibaba Group and NTU Singapore through Alibaba-NTU Global e-Sustainability CorpLab (ANGEL).

References

- Bai, J.; and Wu, X. 2014. Error-tolerant scribbles based interactive image segmentation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 392–399.
- Brown, T.; et al. 2020. Language models are few-shot learners. *Proc. Adv. Neural Inf. Process. Syst.*, 33: 1877–1901.
- Chen, L.; Jiang, Z.; Xiao, J.; and Liu, W. 2021a. Human-like controllable image captioning with verb-specific semantic roles. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 16846–16856.
- Chen, T.; Zhang, R.; and Hinton, G. 2022. Analog bits: Generating discrete data using diffusion models with self-conditioning. *arXiv preprint arXiv:2208.04202*.
- Chen, X.; Zhao, Z.; Yu, F.; Zhang, Y.; and Duan, M. 2021b. Conditional diffusion for interactive segmentation. In *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 7345–7354.
- Cheng, B.; Misra, I.; Schwing, A. G.; Kirillov, A.; and Girdhar, R. 2022. Masked-attention mask transformer for universal image segmentation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 1290–1299.
- Cornia, M.; Baraldi, L.; Cucchiara, R.; et al. 2019. Show, control and tell: A framework for generating controllable and grounded captions. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 8307–8316.
- Dai, X.; Chen, Y.; Yang, J.; Zhang, P.; Yuan, L.; and Zhang, L. 2021. Dynamic detr: End-to-end object detection with dynamic attention. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2988–2997.
- Deng, C.; Ding, N.; Tan, M.; and Wu, Q. 2020. Length-controllable image captioning. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 712–729. Springer.
- Deshpande, A.; Aneja, J.; Wang, L.; Schwing, A. G.; and Forsyth, D. 2019. Fast, diverse and accurate image captioning guided by part-of-speech. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 10695–10704.
- Gan, C.; Gan, Z.; He, X.; Gao, J.; and Deng, L. 2017. Stylenet: Generating attractive visual captions with styles. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 3137–3146.
- Guo, L.; Liu, J.; Tang, J.; Li, J.; Luo, W.; and Lu, H. 2019. Aligning linguistic words and visual semantic units for image captioning. In *ACM International Conference on Multimedia*, 765–773.
- Guo, L.; Liu, J.; Zhu, X.; He, X.; Jiang, J.; and Lu, H. 2020. Non-autoregressive image captioning with counterfactuals-critical multi-agent learning. *Int. Joint Conf. Artif. Intell. (IJCAI)*.
- Guo, Q.; De Mello, S.; Yin, H.; Byeon, W.; Cheung, K. C.; Yu, Y.; Luo, P.; and Liu, S. 2024. Regionopt: Towards region understanding vision language model. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 13796–13806.
- Huang, X.; Wang, J.; Tang, Y.; Zhang, Z.; Hu, H.; Lu, J.; Wang, L.; and Liu, Z. 2024. Segment and caption anything. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 13405–13417.
- Karpathy, A.; and Fei-Fei, L. 2015. Deep visual-semantic alignments for generating image descriptions. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 3128–3137.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 4015–4026.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, 19730–19742. PMLR.
- Li, X.; Thickstun, J.; Gulrajani, I.; Liang, P. S.; and Hashimoto, T. B. 2022. Diffusion-lm improves controllable text generation. *Proc. Adv. Neural Inf. Process. Syst.*, 35: 4328–4343.
- Li, Y.; Ouyang, W.; Zhou, B.; Wang, K.; and Wang, X. 2017. Scene graph generation from objects, phrases and region captions. In *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 1261–1270.
- Lin, J.; et al. 2023. Click-Pixel Cognition Fusion Network With Balanced Cut for Interactive Image Segmentation. *IEEE Trans. Image Process.*, 33: 177–190.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 740–755. Springer.
- Lin, Z.; Zhang, Z.; Chen, L.-Z.; Cheng, M.-M.; and Lu, S.-P. 2020. Interactive image segmentation with first click attention. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 13339–13348.
- Liu, C.; Ding, H.; and Jiang, X. 2023. GRES: Generalized referring expression segmentation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 23592–23601.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 10012–10022.
- Luo, G.; Zhou, Y.; Sun, X.; Cao, L.; Wu, C.; Deng, C.; and Ji, R. 2020. Multi-task collaborative network for joint referring expression comprehension and segmentation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 10034–10043.
- Luo, J.; Li, Y.; Pan, Y.; Yao, T.; Feng, J.; Chao, H.; and Mei, T. 2023. Semantic-conditional diffusion networks for image captioning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 23359–23368.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- OpenAI. 2023. GPT-4 Technical Report. arXiv:2303.08774.
- Pan, Y.; Yao, T.; Li, Y.; and Mei, T. 2020. X-linear attention networks for image captioning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 10971–10980.
- Plummer, B. A.; Wang, L.; Cervantes, C. M.; Caicedo, J. C.; Hockenmaier, J.; and Lazebnik, S. 2015. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2641–2649.
- Pont-Tuset, J.; Uijlings, J.; Changpinyo, S.; Soricut, R.; and Ferrari, V. 2020. Connecting vision and language with localized narratives. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 647–664. Springer.
- Radford, A.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.

- Rasheed, H.; Maaz, M.; Shaji, S.; Shaker, A.; Khan, S.; Cholakkal, H.; Anwer, R. M.; Xing, E.; Yang, M.-H.; and Khan, F. S. 2024. Glamm: Pixel grounding large multimodal model. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 13009–13018.
- Ren, Z.; Huang, Z.; Wei, Y.; Zhao, Y.; Fu, D.; Feng, J.; and Jin, X. 2024. Pixellm: Pixel reasoning with large multimodal model. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 26374–26383.
- Sun, Z.; Zhou, Y.; He, H.; and Mok, P. 2023. Sgdiff: A style guided diffusion model for fashion synthesis. In *ACM International Conference on Multimedia*, 8433–8442.
- Tang, K.; Niu, Y.; Huang, J.; Shi, J.; and Zhang, H. 2020. Unbiased scene graph generation from biased training. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 3716–3725.
- Wang, G.; et al. 2018. DeepIGeoS: a deep interactive geodesic framework for medical image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(7): 1559–1572.
- Wang, T.; et al. 2023. Caption anything: Interactive image description with diverse multimodal controls. *arXiv preprint arXiv:2305.02677*.
- Wang, Z.; Xiao, J.; Zhuang, Y.; Gao, F.; Shao, J.; and Chen, L. 2024. Learning combinatorial prompts for universal controllable image captioning. *International Journal of Computer Vision*, 1–22.
- Wang, Z.; et al. 2022. Cris: Clip-driven referring image segmentation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 11686–11695.
- Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and efficient design for semantic segmentation with transformers. *Proc. Adv. Neural Inf. Process. Syst.*, 34: 12077–12090.
- Yan, K.; Ji, L.; Luo, H.; Zhou, M.; Duan, N.; and Ma, S. 2021. Control image captioning spatially and temporally. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2014–2025.
- Yang, L.; Huang, Z.; Song, Y.; Hong, S.; Li, G.; Zhang, W.; Cui, B.; Ghanem, B.; and Yang, M.-H. 2022. Diffusion-based scene graph to image generation with masked contrastive pre-training. *arXiv preprint arXiv:2211.11138*.
- Yang, X.; Tang, K.; Zhang, H.; and Cai, J. 2019. Auto-encoding scene graphs for image captioning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 10685–10694.
- Young, P.; Lai, A.; Hodosh, M.; and Hockenmaier, J. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Trans. Assoc. Comput. Linguist.*, 2: 67–78.
- Zellers, R.; Yatskar, M.; Thomson, S.; and Choi, Y. 2018. Neural motifs: Scene graph parsing with global context. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 5831–5840.
- Zhang, S.; Liew, J. H.; Wei, Y.; Wei, S.; and Zhao, Y. 2020. Interactive object segmentation with inside-outside guidance. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 12234–12244.
- Zhang, Y.; Ma, Z.; Gao, X.; Shakiah, S.; Gao, Q.; and Chai, J. 2024. GROUNDHOG: Grounding Large Language Models to Holistic Segmentation. *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*.
- Zhao, W.; Wu, X.; and Zhang, X. 2020. Memcap: Memorizing style knowledge for image captioning. In *Proc. AAAI Conf. Artif. Intell.*, 12984–12992.
- Zhong, Y.; Wang, L.; Chen, J.; Yu, D.; and Li, Y. 2020. Comprehensive image captioning via scene graph decomposition. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 211–229. Springer.
- Zhou, Y.; Zhang, Y.; Hu, Z.; and Wang, M. 2021. Semi-autoregressive transformer for image captioning. In *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 3139–3143.
- Zou, X.; Dou, Z.-Y.; Yang, J.; Gan, Z.; Li, L.; Li, C.; Dai, X.; Behl, H.; Wang, J.; Yuan, L.; et al. 2023. Generalized decoding for pixel, image, and language. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 15116–15127.
- Zou, X.; Yang, J.; Zhang, H.; Li, F.; Li, L.; Wang, J.; Wang, L.; Gao, J.; and Lee, Y. J. 2024. Segment everything everywhere all at once. *Proc. Adv. Neural Inf. Process. Syst.*, 36.