

Personalized Image Descriptions from Attention Sequences

Ruoyu Xue¹, Hieu Le², Jingyi Xu¹, Sounak Mondal¹, Abe Leite¹,
 Gregory Zelinsky¹, Minh Hoai³, Dimitris Samaras¹

¹Stony Brook University, USA ²UNC-Charlotte, USA ³The University of Adelaide, Australia

Abstract

People can view the same image differently: they focus on different regions, objects, and details in varying orders and describe them in distinct linguistic styles. This leads to substantial variability in image descriptions. However, existing models for personalized image description generation focus on linguistic style alone, with no prior work leveraging individual viewing patterns. We address this gap by explicitly modeling personalized viewing behavior as a core factor in description generation. Our method, DEPER (DEscription-PERception persona encoder), learns a subject embedding that captures both linguistic style and viewing behavior, guided by an auxiliary attention-prediction task. A lightweight adapter aligns these embeddings with a frozen vision-language model, enabling few-shot personalization without retraining. Across four datasets spanning diverse viewing tasks and both short and detailed descriptions, DEPER achieves a 24% average improvement, showing that modeling personalized attention produces more human-aligned and high-quality descriptions. We posit that understanding how people see helps predict what they say; modeling human diversity in perception can improve both performance and human alignment in multi-modal systems. Code is available at: <https://github.com/cvlab-stonybrook/Personalized-Image-Description>

1. Introduction

People can view the same image differently: they focus on different regions, objects, and details in different orders, and describe them in their own distinct linguistic styles. These individual differences naturally lead to variability in image descriptions. The personalized image description task aims to capture these differences by conditioning generation on *who* is describing the image, rather than only on *what* is depicted [10]. Such personalization has broad applications, including accessibility for people with low vision [21, 53], media summarization [50], education [25], and preference-based product descriptions in advertisements [7]. Beyond practical benefits, it enables more diverse, human-aligned

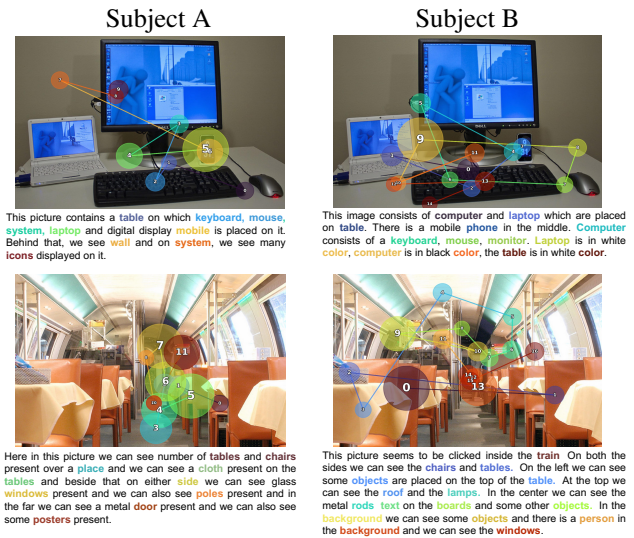


Figure 1. People have distinct viewing habits [9, 65] that shape how they describe an image. Subject A moves between major objects, while Subject B inspects them in detail and in a different order. Our method models these patterns to produce personalized image descriptions. The images and annotations are from [46].

descriptions that enrich multimodal representations [24].

Existing personalization models focus mainly on different linguistic style attributes such as vocabulary or tone [10, 30, 41, 61, 64]. However, beyond personal linguistic style, each person also exhibits a consistent and distinctive pattern in how they view images [65]. Cognitive studies have shown that visual attention patterns shape what people describe [11, 16, 17], and numerous deep-learning models have further demonstrated their critical role in image description generation task [1, 8, 12, 54, 72]. (see Fig. 1). Yet, no prior work has incorporated attention patterns into personalization.

To this end, we propose DEPER (DEscription-PERception persona encoder), a framework that grounds personalization in human attention. DEPER explicitly learns how individuals perceive and explore visual scenes: where they look, in what order, and for how long, and uses these attention patterns to guide personalized image descriptions.

There are two key challenges in leveraging personalized attention into description generation: (1) Modeling personalized human attention is difficult because human attention is noisy, continuous, and behaviorally diverse. Moreover, since attention is strongly conditioned on image content, disentangling stable personalized attention traits from content-specific cues is non-trivial [9, 65]. (2) Adapting image description models to new subjects with few support examples easily leads to overfitting because they are usually parameter-heavy. In real-world deployment, the model must rapidly adapt to and capture user-specific preferences from limited interaction data for new subjects.

The central component of DEPER is a content-invariant subject embedding that captures a viewer’s characteristic patterns of visual attention and linguistic expression. Specifically, DEPER employs a dual-context encoder that jointly models visual trajectories and the corresponding image descriptions. A subject embedding extractor distills this multi-modal information into the subject embedding, while an auxiliary attention trajectory reconstruction objective encourages the embedding to retain key attention dynamics. To generate personalized outputs, an adapter maps the subject embedding into the embedding space of a pretrained vision-language model (VLM). This allows us to prompt the VLM with “Write a description of this photo in the style of <sub j>.”, where <sub j> is the subject-specific token.

During inference, DEPER can produce personalized image descriptions without requiring attention data, since the model has already internalized individual attention patterns within the subject embedding. Additionally, a contrastive loss disentangles subject embeddings from image content, enabling efficient adaptation to new subjects from only a few examples, without additional fine-tuning, while remaining memory- and time-efficient for real-time applications. We evaluate on four datasets [19, 23, 46] that measure attention either via mouse movements or human gaze collected by an eye-tracker, covering three distinct viewing tasks and including either concise or detailed image descriptions. Across these settings, DEPER consistently outperforms baselines, improving BLEU-4 [40] by 12% and CIDEr [56] by 20% on average. In summary, our **contributions** are:

- We develop DEPER, the first method to personalize image description by leveraging human attention.
- Through extensive experiments, we show that DEPER reliably achieves superior performance versus baselines.
- We demonstrate that DEPER can adapt to new subjects with few supporting samples (few-shot personalization).

2. Related Work

Personalized Image Captioning. Personalized image captioning was introduced by CSMN [10], which predicts captions from Instagram posts by building a user-context

memory via TF-IDF [49] over a user’s prior posts. This treats the most frequently used words of an individual in the past as the representation of their personality. Subsequent work [30, 41, 61, 64] largely retains this TF-IDF representation with architectural improvements. MHTN [71] and UMCap [36] add short- and long-term user representations or key-value maps to capture literal preference rather than TF-IDF, but their representation of personality still remains grounded in frequently used words within recurring contexts. A related task [15, 20, 51, 58, 60, 62, 70] involves style-controlled caption generation via explicit text conditioning, such as “sweet” or “dramatic”.

However, prior approaches, with or without explicit style control, generally *ignore human visual attention*, which is a key driver of what gets mentioned, in what order, and the amount of details. Further, they mostly target short captions, while detailed, personalized descriptions remain underexplored. Our work is the first to treat personalized human attention as a key signal for image description generation, and we demonstrate consistent gains for both short and detailed description generation tasks.

Human Attention in Image Description. Human attention has been used to improve description models [8, 12, 54, 72], typically as an auxiliary signal predicted by pretrained models [8, 12, 72], sourced from ground-truth gaze [54], or approximated by noun-aligned boxes [1]. But these signals reflect population-level viewing tendencies, not individual traits. Attention-controlled methods [14, 32, 60, 63, 66] generate descriptions conditioned on fixation cues, yielding spatial-temporal interpretations aligned with the given gaze. Yet they do not learn subject-specific preferences that generalize across images and therefore require attention input at inference, which is impractical for large-scale use. Overall, prior work treats attention as an image-specific signal, not as a transferable, subject-level preference. In contrast, we learn personalization jointly from image, description, and attention signals, capturing individual viewing styles that persist across images (e.g., preference for background vs. foreground, or people vs. actions).

Identity-Based Personalization in Vision-Language Models. A common form of VLM personalization learns “identity tokens”, embeddings that encode a subject identity. Subject-driven text-to-image models like DreamBooth [47] map a person or object’s appearance into such tokens to generate images containing this subject, and personalized retrieval systems like [68] use such tokens to localize the subject across new contexts. Recent personalized VLM assistants [2–4, 18, 37, 38, 43, 44, 48] similarly ground subject identity so the model can recognize and answer questions about one or more specific depicted entities. For instance, Yo’LLaVA [37] learns latent tokens representing a pet like “bo”, enabling subject-specific dialogue. These

methods equate personalization with identity recognition, emphasizing appearance and explicit attributes. In contrast, we personalize to a viewer’s latent viewing patterns as well as their linguistic tendencies.

3. Method

We propose DEPER, a novel method for personalized image description generation. DEPER learns to represent each person as a *subject embedding*, a vector that summarizes how they tend to look at and describe images. This embedding is learned from triplets of (I, D_s, T_s) (the image, the subject’s description, and their attention trajectory) to capture stable, subject-specific patterns of how individuals describe images. This is achieved through our DEPER network (Sec. 3.1). An adapter then projects the subject embedding into the frozen VLM space (Sec. 3.2), enabling description generation in that person’s style without requiring gaze input at test time. To our knowledge, this is the first framework to embed individual attention dynamics into a transferable subject representation for vision–language generation. The overall architecture is illustrated in Fig. 2.

3.1. DEscription-PERception Persona Encoder

DEPER is designed to learn a subject embedding \mathbf{z}_s that serves as a personalized representation capturing both how a person perceives and describes visual scenes. We require \mathbf{z}_s to be consistent across images for a single viewer, distinctive across viewers, and structured in a way that reflects real behavioral tendencies. To achieve this, DEPER integrates three complementary modules: A **dual-context encoder** (Sec. 3.1.1) fuses visual, linguistic, and attentional streams to capture how perception and expression interact. A **subject embedding extractor** (Sec. 3.1.2) distills this fused representation into the compact, subject-specific \mathbf{z}_s under discriminative supervision. A **trajectory decoder** (Sec. 3.1.3) reconstructs each subject’s viewing sequence conditioned on \mathbf{z}_s , reinforcing its capacity to capture personalized attention dynamics.

3.1.1. Dual-Context Encoder

The dual-context encoder computes a representation capturing both perception and expression via two interacting streams that repeatedly exchange information through cross-attention.

Given an image I , description D_s , and trajectory $T_s = \{(b_i, \tau_i)\}_{i=1}^M$ (where each \mathbf{b}_i is a bounding box with duration τ_i . we extract image patch features \mathbf{V} , text token features \mathbf{L}_0 , and trajectory features \mathbf{T}_0 (See supplementary for the details of obtaining trajectory features). To better capture visual dynamics, especially the scan order of scene exploration, and the duration of attention on each region, we apply sinusoidal positional encoding [13] to encode the box duration and position (the index of boxes), and add them to

trajectory features, yielding multi-scale temporal features. The encoder alternates self- and cross-attention across the two streams: text tokens attend to image and trajectory context, and trajectory tokens attend to image and text context:

$$\mathbf{T}_{\ell+1} = \text{FFN}_{\ell_T}(\text{Cross}_{\ell_T}(\text{Self}_{\ell_T}(\mathbf{T}_\ell), [\mathbf{V}; \mathbf{L}_\ell])), \quad (1)$$

$$\mathbf{L}_{\ell+1} = \text{FFN}_{\ell_L}(\text{Cross}_{\ell_L}(\text{Self}_{\ell_L}(\mathbf{L}_\ell), [\mathbf{V}; \mathbf{T}_\ell])). \quad (2)$$

Here, $\text{Self}(\cdot)$, $\text{Cross}(\cdot)$ and $\text{FFN}(\cdot)$ denote self-attention layer following the design in [55], $[\cdot; \cdot]$ is concatenation. We repeat it for ℓ layers. This design yields a fused representation $\mathbf{Z}_{\text{dual}} = [\mathbf{L}'; \mathbf{T}']$ that inherently captures each subject’s linguistic style and attentional behavior.

3.1.2. Subject Embedding Extractor

We design the subject embedding extractor to produce a representation that captures a person’s consistent perceptual–linguistic behavior. It consists a learnable subject query q_s that attends to \mathbf{Z}_{dual} through a cross-attention layer. The subject query is optimized to selectively aggregate features that capture a specific subject’s distinctive viewing and descriptive patterns across images, resulting in a stable subject embedding \mathbf{z}_s . To this end, we supervise the extractor with a joint objective that integrates subject classification and supervised contrastive learning [22]. This joint formulation strengthens representational discriminability and prevents collapse into a shared subspace.

Subject embedding Losses. A classification head predicts the subject ID from \mathbf{z}_s to enforce inter-subject discrimination, optimized using a standard cross-entropy loss (\mathcal{L}_{cls}). An additional contrastive loss (\mathcal{L}_{con} , SupCon [22]) pulls embeddings of the same subject closer and pushes those of different subjects apart. See details in the supplementary.

3.1.3. Trajectory Decoder

A key component of our framework is the integration of personal viewing patterns into the subject embedding. To enforce this, we introduce trajectory decoder to reconstruct the attention trajectory. This decoder takes as input an instance-specific trajectory latent \mathbf{z}_{traj} , and reconstructs that instance’s attention trajectory T_s while being conditioned on the visual features \mathbf{V} , description features \mathbf{L}_0 , and subject embedding \mathbf{z}_s in a personality-aware manner. Thus, \mathbf{z}_s learns stable, individual viewing patterns that guide trajectory reconstruction, without collapsing into instance-specific memorization.

Decoding block. We initialize a learnable trajectory query \mathbf{q}_{traj} that extracts instance-specific attention dynamics from \mathbf{Z}_{dual} through cross-attention, yielding the trajectory latent \mathbf{z}_{traj} . We then initialize a sequence of M box queries $Q_0 = \{\mathbf{q}_i\}_{i=1}^M$ and broadcast \mathbf{z}_{traj} to each query as a global prior. Each layer first applies self-attention within the box queries, followed by cross-attention to \mathbf{z}_{traj} , providing step-wise latent control throughout decoding:

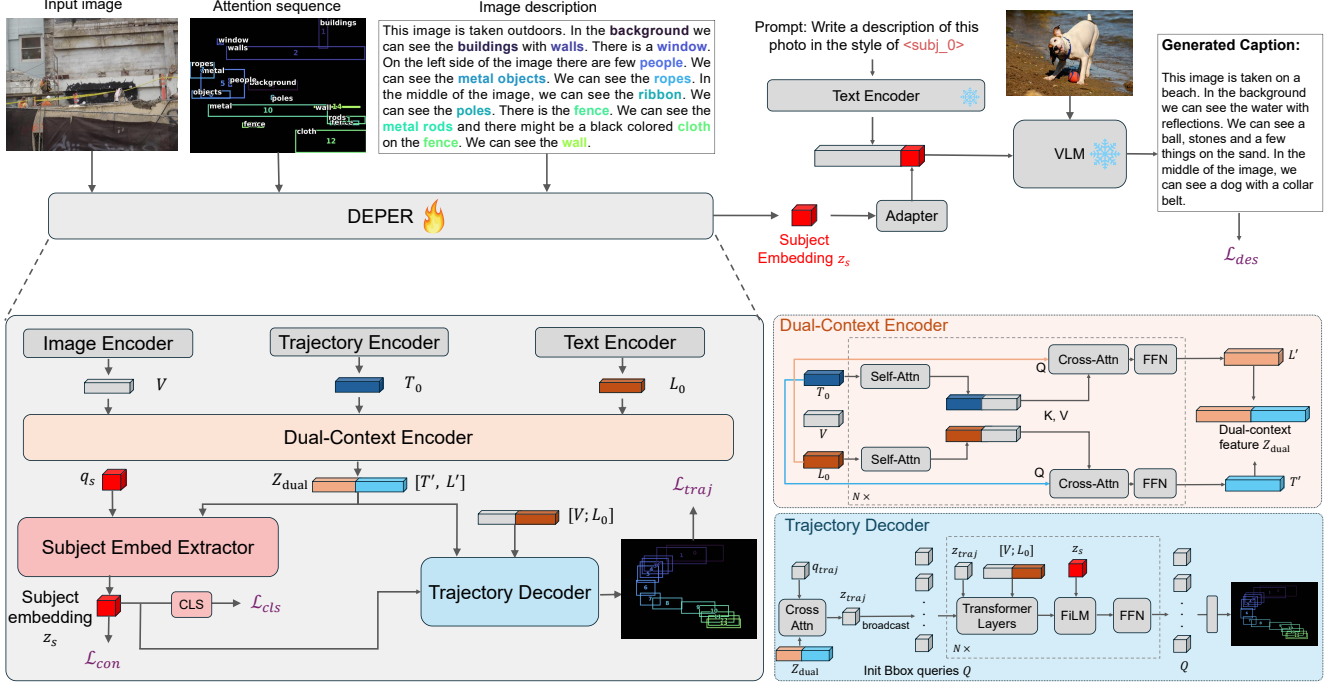


Figure 2. **Overview of DEPER and Description Generation:** DEPER extracts a subject embedding z_s from a triplet (I, D_s, T_s) , capturing the personalized viewing patterns and linguistic style. z_s conditions a VLM to produce subject-aligned image descriptions. A Dual-Context Encoder aligns perceptual and linguistic features into Z_{dual} . A Subject Embedding Extractor then distills Z_{dual} to z_s , yielding personalized attention–linguistic traits. z_s is distinctive across subjects yet consistent across images, enforced by classification and contrastive losses. A trajectory decoder further encourages Z_{dual} to capture viewing dynamics, and helps z_s capture a subject’s exploration behavior.

$$\mathbf{Q}_\ell = \text{Cross}_\ell(\text{Self}_\ell(\mathbf{Q}_\ell), \mathbf{z}_{traj}). \quad (3)$$

We then condition the reconstruction on the visual and linguistic context via cross-attention to the image features \mathbf{V} and description features \mathbf{L}_0 . We use FiLM [42] with the subject embedding z_s so that it modulates each decoder block in a personality-aware manner. Finally, we apply a position-wise feed-forward network and repeat it for ℓ layers to get final box embeddings \mathbf{Q}' :

$$\mathbf{Q}_{\ell+1} = \text{FFN}_\ell(\text{FiLM}(\text{Cross}_\ell(\mathbf{Q}_\ell, [\mathbf{V}; \mathbf{L}_0]), z_s)). \quad (4)$$

Last, we apply a linear head to the \mathbf{Q}' to obtain the box coordinates, yielding the predicted box sequence and their corresponding validity $\hat{\mathbf{B}} \in [0, 1]^{T \times 4}$ and $\hat{\mathbf{V}} \in [0, 1]^T$.

Training Objective. We use a smooth L1 loss (SL_1) on the box predictions $\hat{\mathbf{B}}$, and a binary cross-entropy loss (BCE) on $\hat{\mathbf{V}}$, for an overall loss $\mathcal{L}_{traj} = \mathcal{L}_{box} + \mathcal{L}_{valid}$:

$$\mathcal{L}_{box} = \left(\sum_{i=1}^M v_i \text{SL}_1(\hat{\mathbf{b}}_i, \mathbf{b}_i) \right) / \left(\sum_{i=1}^M v_i + \varepsilon \right), \quad (5)$$

$$\mathcal{L}_{valid} = \frac{1}{M} \sum_{t=1}^M \text{BCE}(\hat{v}_t, v_t).$$

3.2. Image Description Generation

VLM conditioning. Given a subject embedding, our next goal is to make the VLM generate descriptions in that person’s style. We achieve this by injecting the embedding directly into the model’s prompt space. First, we add subject tokens $\langle \text{subj}_x \rangle$ into VLM’s vocabulary (x denotes different subjects). Then a small adapter (single linear layer, as in LLaVA [29]) maps the DEPER embedding into the VLM’s token dimension, and then this adapted vector replaces the embedding of a dedicated subject token in the prompt. The rest of the VLM remains frozen; it simply treats this subject vector as part of the input sequence and adapts its generation accordingly. To be precise, in the prompt “Write a description of this photo in the style of $\langle \text{subj}_x \rangle$.”, we replace the token $\langle \text{subj}_x \rangle$ with the adapted subject vector at the embedding layer. This gives the model a direct, continuous representation of the subject’s style while keeping the prompt and VLM architecture unchanged. To avoid information leakage since DEPER takes descriptions as input, we condition the VLM on a *different* image–description pair (I', D'_s) from the same subject as the (I, D_s, T_s) used by DEPER.

Captioning Loss. We use Supervised Fine-Tuning [39]: the VLM generates \hat{D}'_s of length $N_{D'}$ conditioned on

prompt $P(\mathbf{z}_s)$:

$$\mathcal{L}_{\text{des}} = - \sum_{t=1}^{N_{D'}} \log p_{\phi}(d_t | I, P(\mathbf{z}_s), d_{1 \dots t-1}) \quad (6)$$

3.3. Training

We train DEPER in two stages. In Stage 1, we train with $\mathcal{L}_{\text{stage1}} = \lambda \mathcal{L}_{\text{con}} + \mathcal{L}_{\text{traj}} + \mathcal{L}_{\text{cls}}$. \mathcal{L}_{des} and \mathcal{L}_{cls} supervise DEPER to optimize subject embeddings to be image-independent and personality-aware; $\mathcal{L}_{\text{traj}}$ forces the dual-context encoder to learn visual dynamics.

In Stage 2, we freeze the dual-context encoder as it has been trained to a robust representation. We train the subject embedding extractor and the VLM adapter to utilize the pretrained knowledge in Stage 1, and align subject embeddings to VLM space. The training loss is $\mathcal{L}_{\text{Stage2}} = \mathcal{L}_{\text{des}} + \lambda \mathcal{L}_{\text{con}} + \mathcal{L}_{\text{cls}}$. The classification and contrastive losses in Stage 2 ensure that subject embeddings remain distinct and stable in the subject space.

4. Experiments

4.1. Dataset Setting

We evaluate on four datasets. (1) Localized Narratives [46]: participants describe images while moving the mouse over mentioned regions, yielding paired text–trace annotations on public image sets such as COCO [28] and Flickr30k [45]. We denote them as COCO-LN and Flickr30k-LN, with 108 and 37 subjects respectively. (2) He et al. [19]: 5 subjects view 1000 images and simultaneously speak a one-sentence caption. (3) Kollenda et al. [23]: participants observe 100 natural scenes for 3 s, then briefly describe them post-viewing, with gaze recorded during observation. For each dataset, we split training subjects into seen and unseen splits. Each unseen subject has 5 samples as the few-shot support set. We repeat this sampling by 5 times and report the average scores for reliability. Detailed statistics are shown in Tab. 1. Results on both seen and unseen are reported on the official test set when available; otherwise, a validation subset is sampled from the training data and the original validation set is used for testing.

4.2. Experiment Setting

Baselines: MITR[32]: A population-level (without personalization) image description model to generate image descriptions on Localized Narratives with train-time attention trajectory supervision. Qwen Zero-shot: Qwen2-VL-2B[59] off-the-shelf. It generates descriptions for each image, without personalization. CSMN [10]: to our knowledge, the only personalized image captioning model with publicly released code. Qwen+PT: Qwen with prompt tuning [26]. We add user-specific words ($\langle \text{subj}_1 \rangle$, $\langle \text{subj}_2 \rangle$, ...) to Qwen2-VL-2B[59]’s vocabulary and update only the corresponding weights in the input embedding

Table 1. **Dataset statistics.** The datasets are different sizes, but all have low human consistency, showing high variability among subjects. [#train/test]: number of annotated images; [#seen/unseen]: number of subjects in train and few-shot sets; [length] avg. description length; [HC] human consistency (m-BLEU-4).

Dataset	#train/test	#seen/unseen	length	HC
COCO-LN[46]	134272 / 8573	89 / 19	41	0.037
Flik30k-LN[46]	30546 / 1023	27 / 10	57	0.061
Kollenda et al.[23]	1950 / 450	22 / 8	16	0.054
He et al.[19]	3998 / 1000	5 / 0	8	0.054

layer that maps these tokens to embeddings, keeping all other parameters frozen. MITR-FT: MITR[32] fine-tuned separately for each subject using all of that subject’s data. Qwen few-shot: We prompt Qwen using the query image and the support set’s descriptions and attention trajectories, guiding it to generate a description consistent with the stylistic patterns demonstrated in the examples. See supplementary for the prompt we use.

Training and configuration details. The hidden dimension of DEPER is set to 384. The dual-context encoder, subject embedding extractor, and trajectory decoder use 2, 1, and 4 layers, respectively. We set $\lambda = 0.1$ for the contrastive loss for both stages. DEPER is trained for 40 epochs in Stage 1 and 15 epochs in Stage 2 with a learning rate of 0.0005. Training is conducted on two RTX A6000 GPUs with a batch size of 16 and gradient accumulation step of 4. We use Qwen2-VL-2B-Instruct [59] as our image description backbone, and use DINOv3 (ConvNeXt-Tiny) [52] as our image encoder. See supplementary for more details.

Inference. On the seen split, we estimate a subject representation by sampling up to $K = 100$ (see supplementary for details) image–description–trajectory triplets from the subject’s training set. For the unseen split, we use the support set (Sec. 4.1) to compute a DEPER embedding for each triplet and average them to obtain a single subject embedding. This design removes the need for trajectories at test time, simplifying deployment in practice.

4.3. Evaluation Metrics

We adopt two different kinds of metrics to conduct a comprehensive evaluation.

Widely adopted captioning metrics. We use BLEU-1 & 4[40], METEOR [5], ROUGE_L [27], and CIDEr [56]. We also use Polos [57], a metric using a VLM trained on large-scale human preference data for caption evaluation.

Personalization-centric evaluation. Following [6, 33, 34, 67, 69], we propose the Object Sequence Score (OSS): an object-level metric that compares *which* objects are mentioned and their *order*, capturing personalized narrative alignment. We extract ordered nouns from the prediction

Table 2. Quantitative results on four datasets of **seen** subjects split. The first two baselines are population-level image description models, and the others are personalization models. [B1]:BLEU-1, [B4]:BLEU-4, [M]:METEOR, [R]:ROUGE_L, [C]:CIDEr, [P]:Polos. Best per dataset in **bold**.

Method	COCO-LN [46]								Flickr30k-LN [46]							
	B1	B4	M	R	C	P	OSS	CLS	B1	B4	M	R	C	P	OSS	CLS
MITR [32]	0.415	0.142	0.182	0.336	0.139	0.512	0.218	–	0.296	0.076	0.160	0.281	0.077	0.319	0.200	–
Qwen Zero-shot [59]	0.161	0.023	0.165	0.181	0.007	0.584	0.115	–	0.169	0.024	0.166	0.187	0.004	0.612	0.133	–
MITR-FT [32]	0.437	0.176	0.209	0.341	0.142	0.589	0.246	0.415	0.302	0.101	0.169	0.298	0.094	0.320	0.224	0.427
CSMN [10]	0.295	0.086	0.156	0.299	0.086	0.177	0.133	0.443	0.071	0.010	0.042	0.083	0.003	0.142	0.070	0.459
Qwen+PT	0.304	0.145	0.205	0.432	0.587	0.612	0.340	0.623	0.271	0.135	0.207	0.426	0.498	0.654	0.320	0.563
Qwen+DEPER (Ours)	0.510	0.264	0.240	0.482	0.726	0.638	0.392	0.686	0.542	0.312	0.272	0.518	0.789	0.671	0.408	0.796

Method	Kollenda <i>et al.</i> [23]								He <i>et al.</i> [19]							
	B1	B4	M	R	C	P	OSS	CLS	B1	B4	M	R	C	P	OSS	CLS
Qwen Zero-shot [59]	0.267	0.047	0.135	0.291	0.363	0.559	0.317	–	0.430	0.132	0.209	0.422	1.144	0.578	0.377	–
CSMN [10]	0.019	0.001	0.012	0.034	0.002	0.100	0.012	0.021	0.025	0.008	0.032	0.003	0.004	0.091	0.015	0.157
Qwen+PT	0.344	0.067	0.176	0.325	0.504	0.527	0.328	0.056	0.475	0.174	0.211	0.446	1.515	0.589	0.448	0.262
Qwen+DEPER (Ours)	0.442	0.135	0.201	0.382	0.871	0.594	0.351	0.083	0.506	0.207	0.234	0.486	1.822	0.603	0.470	0.307

and reference, then align them via Needleman–Wunsch [35] with weighted matches of exact, stem and synonym. We also report top-1 classification accuracy (**CLS**) to assess whether a subject’s generated description is distinguishable from others for the same image. For each image, we rank all subjects’ generated descriptions under a given metric (using images with ≥ 3 subjects; Flickr30k-LN uses ≥ 2 due to limited annotations). A hit occurs when the same-subject description ranks first. We report the mean hit rate over BLEU-4, METEOR, ROUGE-L, and CIDEr.

4.4. Main Results

Human Consistency. To quantify how differently humans describe the same image, we compute Human Consistency (HC) using *m-BLEU-4* [31], which measures the average similarity among descriptions of the same image, with units comparison to BLEU-4. As shown in Tab. 1, human descriptions vary substantially. Identical captions yield HC = 1, while lower scores indicate greater diversity.

Performance on seen subjects. As shown in Tab. 2, our method achieves the best performance on all metrics, showing improvements (averaged over datasets) of 62% on BLEU-4, 28% on CIDEr. The improvement from Qwen+PT to Qwen+DEPER shows that, beyond prompt tuning, DEPER learns subject embeddings that more faithfully capture each author’s viewing patterns and linguistic style, producing higher-quality personalized descriptions.

Our method also yields strong gains on personalization-focused metrics. The 13.0% OSS gain reflects better modeling of the three key components of human attention: which objects are mentioned, in what order, and to what amount of detail. The 15.4% CLS gain shows that our generated descriptions effectively capture subjects’ distinctive features.

Table 3. Personalized description generation performance on **unseen** subjects. [C][46]:COCO-LN, [F][46]:Flickr30k-LN[46], [K][23]:Kollenda *et al.* [23]. Best per dataset in **bold**.

Dataset	Method	B4	M	R	C	P	OSS	CLS
C[46]	MITR-FT[41]	0.139	0.179	0.334	0.135	0.491	0.201	0.326
	CSMN[10]	0.042	0.105	0.133	0.026	0.149	0.198	0.312
	Qwen few-shot[59]	0.071	0.106	0.203	0.077	0.397	0.142	0.406
	Qwen+PT	0.058	0.141	0.265	0.317	0.434	0.138	0.317
	Ours	0.164	0.184	0.389	0.453	0.597	0.330	0.445
F[46]	MITR-FT[41]	0.074	0.126	0.261	0.104	0.209	0.187	0.415
	CSMN[10]	0.007	0.021	0.053	0.001	0.102	0.029	0.427
	Qwen few-shot[59]	0.122	0.148	0.241	0.068	0.386	0.150	0.416
	Qwen+PT	0.074	0.167	0.337	0.338	0.587	0.278	0.479
	Ours	0.202	0.232	0.410	0.382	0.610	0.329	0.625
K[23]	CSMN[10]	0.003	0.007	0.015	0.000	0.092	0.004	0.085
	Qwen few-shot[59]	0.063	0.157	0.288	0.538	0.535	0.272	0.151
	Qwen+PT	0.019	0.148	0.237	0.209	0.507	0.220	0.111
	Ours	0.143	0.207	0.398	1.053	0.583	0.380	0.157

The 30-way classification on Kollenda *et al.* is hard, yet our method achieves a $\sim 150\%$ improvement over chance. CLS is not computed for population-level models, which generate identical descriptions per image.

Qualitative Results on seen subjects. In Fig. 3, we show one example per dataset, each with two subjects and their brief or detailed descriptions. These cases illustrate components of subject-specific variation that our method captures: (1) object-centric vs. scene-centric openings (Example 1); (2) description granularity, where one subject offers fine-grained details while the other is terse (Examples 2 and 4); and (3) object-of-interest selection, e.g., background-oriented vs. sign-focused attention (Example 3).

Subject 1		Subject 2	
<p>(GT1) In this picture there are some elephants who are standing on the floor in front of the water behind there are some trees.</p> <p>(Ours) This is the picture of some elephants which are standing in the water and around there are some trees.</p> <p>(Q+PT) This is the picture of some elephants standing in the water and some of them are standing on the land.</p>		<p>(GT2) This is a picture taken in the outdoors. It is sunny. There are group of elephants drinking water in the river. Behind the elephants there are trees.</p> <p>(Ours) This is a picture taken of the outdoors. It is sunny. We can see there are elephants standing on the path and in front of the elephants there is water. Behind the elephants there are trees.</p> <p>(Q+PT) This is a picture taken of the outdoors. We can see elephants are standing on the ground and water.</p>	
<p>(GT1) In this image there are people. There is a ball in the air. There are trees, pillar, metal structure. We can see the shadows of the people on the grass surface.</p> <p>(Ours) In this image there are two people playing a game. There is a ball. There are trees. There is grass on the surface.</p> <p>(Q+PT) In this image, we can see two persons playing a game. At the bottom, there is grass on the ground. In the background, there are trees and we can see a pole. At the top right, there is a goal post.</p>		<p>(GT2) In this image, we can see people playing on the ground and they are wearing sports dress, shoes, socks and one of them is wearing a headband. We can see a ball in the air. There are trees, railings and there is a pole. There are shadows on the ground.</p> <p>(Ours) In this image, we can see people playing a game and one of them is wearing a headband. In the background, there are trees and poles. At the bottom, there is a ground.</p> <p>(Q+PT) In this picture we can see two women playing soccer, we can see ball, in the background of the picture there are some trees.</p>	
<p>(GT1) A seagull stands on a post. You can see breaking waves behind her.</p> <p>(Ours) A seagull sits on a pole in the foreground and looks at the camera. In the background, the waves roll in.</p> <p>(Q+PT) A seagull sits on a pole in front of a sign that says "surfing area".</p>		<p>(GT2) A seagull sits on a beach on a bollard, next to it is a sign on which "Surf Area" and a pictogram with a surf.</p> <p>(Ours) A seagull sits on a pole in the surf, there is a sign next to it that says "Surfing Area".</p> <p>(Q+PT) A seagull sits on a pole in front of the ocean.</p>	
		<p>(GT2) a steam train in a station.</p> <p>(Ours) a train engine at a station.</p> <p>(Q+PT) a steam train at a station.</p>	
			<p>(GT1) a black steam train with blue and yellow boards on it.</p> <p>(Ours) a black train with a yellow and blue stripe. (Q+PT) a train engine at a station.</p>

Figure 3. **Qualitative Results** show one example per dataset, each with two subject-specific descriptions (subjects 1 and 2). From top to bottom and left to right: COCO-LN [46], Flickr30k-LN [46], Kollenda *et al.* [23], and He *et al.* [19]. Subject-distinct content is highlighted in red. Qwen+PT is denoted as Q+PT.

Performance on unseen subjects. In Tab. 3, our method shows strong few-shot performance across all datasets. On Kollenda *et al.*, the unseen split exhibits comparable results to the seen split. Few-shot personalization is harder on COCO-LN and Flickr30k-LN because their support images differ from the seen-subject training data; yet, our performance remains stable across all five metrics. The higher CLS scores indicate that DEPER can infer distinct subject embeddings for previously unseen individuals, demonstrating the effective transfer of subject-specific cues. Notably, our method extracts subject embeddings without per-subject fine-tuning, enabling real-time adaptation and avoiding separate models for each subject. We omit unseen evaluation on He *et al.* due to its small subject count (N=5). We repeat the experiment 5 times with different sampled support sets; all baselines use the same sets.

Qualitative results on DEPER’s outputs. Fig. 4 visualizes DEPER’s learned subject embeddings and reconstructed attention trajectories for both seen and unseen subjects. The embedding space forms clear, subject-specific clusters, while the ground-truth and reconstructed trajectories illustrate that DEPER captures visual dynamics well. The coherent clustering of unseen subjects and the close match between reconstructed and true trajectories highlight strong generalization to new individuals. We present further results in supplementary.

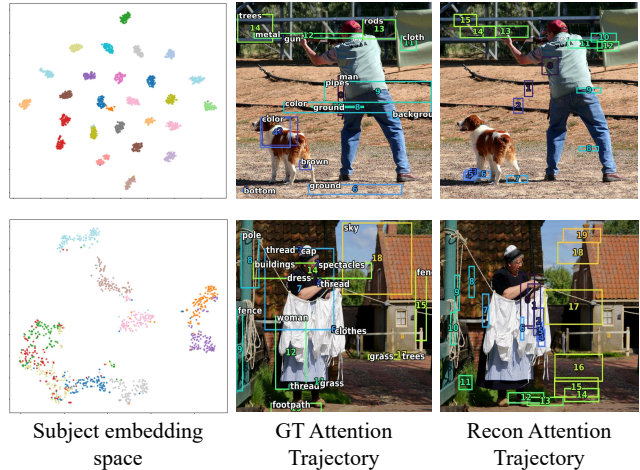


Figure 4. **Qualitative results of DEPER’s outputs.** We show DEPER’s outputs on seen and unseen splits (first and second rows) of Flickr30k-LN. The first column visualizes DEPER’s subject embeddings, where colors denote subjects and each point represents an image–description–trajectory triplet. The second column shows ground-truth attention trajectories with their corresponding nouns and orders; the third column shows reconstructed trajectories from the test set after Stage-2 training.

4.5. Ablations

Ablation on attention trajectory. To evaluate the role of human attention in shaping subject embeddings, we ablate three components in Tab. 4: adding attention-trajectory fea-

Table 4. **Ablation on DEPER’s components.** *Text Input* and *Traj Input* denote the input modalities provided to DEPER, *Traj Dyn* denotes the duration and index of each bounding box. *Traj Recon* refers to the presence of the trajectory–reconstruction decoder module. \checkmark indicates the inclusion of the corresponding module. The results show the importance of attention trajectories in image description generation at different stages. Results are reported on the Flickr30k-LN.

Text Input	Traj Input	Traj Dyn	Traj Recon	B4	M	R	C	OSS	CLS
\checkmark	–	–	–	0.222	0.247	0.500	0.770	0.379	0.649
\checkmark	\checkmark	–	\checkmark	0.276	0.266	0.514	0.748	0.378	0.731
\checkmark	\checkmark	\checkmark	–	0.230	0.256	0.499	0.774	0.381	0.724
\checkmark	\checkmark	\checkmark	\checkmark	0.312	0.272	0.518	0.789	0.408	0.796

tures, adding visual dynamics (fixation duration and order), and adding the trajectory reconstruction objective. Removing attention trajectories entirely causes a clear drop in performance, showing that human attention is critical to modeling subject-specific perception.

Results show that multiple attention components are required for full performance. Removing trajectory dynamics (duration and order) reduces performance to mid-way between the no-attention variant and the full model, highlighting their importance. Removing the trajectory-reconstruction objective degrades performance further, underscoring the need to preserve attentional information in \mathbf{Z}_{dual} and \mathbf{z}_s beyond linguistic style. See the supplementary for alternatives when human-annotated attention is unavailable.

Table 5. **Ablation on the effect of different modules.** We report performance under key removals on Flickr30k-LN: w/o Traj Latent (directly use subject embedding to reconstruct trajectory), w/o Dual-Context (dual-context encoder removed), w/o Contrast (contrastive loss disabled in both stages), w/o FiLM (no subject-embedding FiLM modulation in trajectory decoder).

Modules	B4	M	R	C	OSS	CLS
w/o Dual-Context	0.229	0.252	0.495	0.729	0.380	0.731
w/o Traj Latent	0.272	0.261	0.505	0.745	0.391	0.750
w/o Contrast	0.228	0.259	0.498	0.743	0.386	0.722
w/o FiLM	0.270	0.262	0.506	0.723	0.394	0.768
Ours	0.312	0.272	0.518	0.789	0.408	0.796

Ablation on modules. We further ablate modules in Tab. 5. Removing any component causes a significant drop, showing their complementarity: the Dual-Context Encoder captures cross-modal correspondences; the trajectory latent \mathbf{z}_{traj} reduces the need for the subject embedding to encode instance-specific reconstruction details; the contrastive loss improves subject discrimination; and FiLM injects subject-aware signals into trajectory reconstruction.

Table 6. **Performance with varying training sizes.** # Samples denotes the number of training examples per subject. The results show DEPER’s data efficiency. Results are reported on the Flickr30k-LN dataset.

# samples	B4	M	R	C	OSS	CLS
100	0.171	0.200	0.402	0.400	0.381	0.730
200	0.268	0.255	0.497	0.664	0.383	0.750
500	0.292	0.271	0.513	0.754	0.389	0.758
801(Full)	0.312	0.272	0.518	0.789	0.408	0.796

4.6. Analysis

Data efficiency analysis. To assess data efficiency, we vary the number of training samples per subject and evaluate on the seen split. As shown in Tab. 6, performance drops only slightly with 62% of the data. Even with 100 samples per subject (2,700 total), DEPER remains comparable to baselines trained on the full dataset. These results demonstrate that DEPER learns robust subject embeddings from limited supervision, reducing data requirements for VLM fine-tuning and enabling applications in data-scarce domains such as healthcare and assistive vision.

5. Conclusion and Discussion

We study how personalized human attention shapes personalized image descriptions, noting that people describe the image differently not only in what they say and the words they use, but also in the order and detail in which they explore a scene. However, modeling this implicit and consistent viewing behavior is challenging, especially when a system must adapt to new subjects with only a few examples. To address this, we develop a unified pipeline that learns a subject-specific representation capturing both a person’s visual exploration behavior and linguistic style, and uses this representation to guide a frozen VLM to generate personalized descriptions. DEPER consistently outperforms prior methods on seen subjects and achieves strong few-shot generalization to unseen individuals. In ablations, we demonstrate that attention signals are critical to this performance. Overall, our approach sets a new state-of-the-art for personalized image description generation task and highlights the value of richer, behavior-aware subject representations for future research. Future work may extend DEPER to learn personality representations beyond image description. Attention–language alignment may also benefit tasks such as visual question answering and robotics, where understanding implicit human behavior is important.

6. Acknowledgement

This project is supported by US National Science Foundation grant IIS-2123920.

References

- [1] Rehab Alahmadi and James Hahn. Improve image captioning by estimating the gazing patterns from the caption. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1025–1034, 2022. 1, 2
- [2] Yuval Alaluf, Elad Richardson, Sergey Tulyakov, Kfir Aberman, and Daniel Cohen-Or. Myvlm: Personalizing vlms for user-specific queries. In *European Conference on Computer Vision*, pages 73–91. Springer, 2024. 2
- [3] Ruichuan An, Sihan Yang, Ming Lu, Renrui Zhang, Kai Zeng, Yulin Luo, Jiajun Cao, Hao Liang, Ying Chen, Qi She, Shanghang Zhang, and Wentao Zhang. Mc-llava: Multi-concept personalized vision-language model, 2025.
- [4] Ruichuan An, Sihan Yang, Renrui Zhang, Zijun Shen, Ming Lu, Gaole Dai, Hao Liang, Ziyu Guo, Shilin Yan, Yulin Luo, Bocheng Zou, Chaoqun Yang, and Wentao Zhang. Unictokens: Boosting personalized understanding and generation via unified concept tokens, 2025. 2
- [5] Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005. 5
- [6] Souradeep Chakraborty, Ruoyu Xue, Rajarsi Gupta, Oksana Yaskiv, Constantin Friedman, Natallia Sheuka, Dana Perez, Paul Friedman, Won-Tak Choi, Waqas Mahmud, et al. Measuring and predicting where and when pathologists focus their visual attention while grading whole slide images of cancer. *Medical Image Analysis*, page 103752, 2025. 5
- [7] Qibin Chen, Junyang Lin, Yichang Zhang, Hongxia Yang, Jingren Zhou, and Jie Tang. Towards knowledge-based personalized product description generation in e-commerce. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 3040–3050, 2019. 1
- [8] Shi Chen and Qi Zhao. Boosted attention: Leveraging human attention for image captioning. In *Proceedings of the European conference on computer vision (ECCV)*, pages 68–84, 2018. 1, 2
- [9] Xianyu Chen, Ming Jiang, and Qi Zhao. Beyond average: Individualized visual scanpath prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25420–25431, 2024. 1, 2
- [10] Cesc Chunseong Park, Byeongchang Kim, and Gunhee Kim. Attend to you: Personalized image captioning with context sequence memory networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 895–903, 2017. 1, 2, 5, 6
- [11] Moreno I Coco and Frank Keller. Scan patterns predict sentence production in the cross-modal processing of visual scenes. *Cognitive science*, 36(7):1204–1223, 2012. 1
- [12] Marcella Cornia, Lorenzo Baraldi, Giuseppe Serra, and Rita Cucchiara. Paying more attention to saliency: Image captioning with saliency and context attention. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 14(2):1–21, 2018. 1, 2
- [13] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3
- [14] Yuhu Feng, Keisuke Maeda, Takahiro Ogawa, and Miki Haseyama. Human-centric image retrieval with gaze-based image captioning. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 3828–3832. IEEE, 2022. 2
- [15] Chuang Gan, Zhe Gan, Xiaodong He, Jianfeng Gao, and Li Deng. Stylenet: Generating attractive visual captions with styles. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3137–3146, 2017. 2
- [16] Lila R Gleitman, David January, Rebecca Nappa, and John C Trueswell. On the give and take between event apprehension and utterance formulation. *Journal of memory and language*, 57(4):544–569, 2007. 1
- [17] Zenzi M Griffin and Kathryn Bock. What the eyes say about speaking. *Psychological science*, 11(4):274–279, 2000. 1
- [18] Haoran Hao, Jiaming Han, Changsheng Li, Yu-Feng Li, and Xiangyu Yue. Rap: Retrieval-augmented personalization for multimodal large language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 14538–14548, 2025. 2
- [19] Sen He, Hamed R Tavakoli, Ali Borji, and Nicolas Pugeault. Human attention in image captioning: Dataset and analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8529–8538, 2019. 2, 5, 6, 7
- [20] Mehrdad Hosseinzadeh and Yang Wang. Few-shot personality-specific image captioning via meta-learning. In *2023 20th Conference on Robots and Vision (CRV)*, pages 320–327. IEEE, 2023. 2
- [21] Lucy Jiang, Crescentia Jung, Mahika Phutane, Abigale Stangl, and Shiri Azenkot. “it’s kind of context dependent”: Understanding blind and low vision people’s video accessibility preferences across viewing scenarios. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–20, 2024. 1
- [22] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673, 2020. 3
- [23] Diana Kollenda, Anna-Sophia Reher, and Benjamin de Haas. Individual gaze predicts individual scene descriptions. *Scientific Reports*, 15(1):9443, 2025. 2, 5, 6, 7
- [24] Samuel Lavoie, Polina Kirichenko, Mark Ibrahim, Mahmoud Assran, Andrew Gordon Wilson, Aaron Courville, and Nicolas Ballas. Modeling caption diversity in contrastive vision-language pretraining. *arXiv preprint arXiv:2405.00740*, 2024. 1
- [25] Maurizio Leotta, Fabrizio Mori, and Marina Ribaudò. Evaluating the effectiveness of automatic image captioning for web accessibility. *Universal access in the information society*, 22(4):1293–1313, 2023. 1
- [26] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021. 5

- [27] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004. 5
- [28] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 5
- [29] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023. 4
- [30] Cuirong Long, Xiaoshan Yang, and Changsheng Xu. Cross-domain personalized image captioning. *Multimedia Tools and Applications*, 79(45):33333–33348, 2020. 1, 2
- [31] Shweta Mahajan and Stefan Roth. Diverse image captioning with context-object split latent spaces. *Advances in Neural Information Processing Systems*, 33:3613–3624, 2020. 6
- [32] Zihang Meng, Licheng Yu, Ning Zhang, Tamara L Berg, Babak Damavandi, Vikas Singh, and Amy Bearman. Connecting what to say with where to look by modeling human attention traces. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12679–12688, 2021. 2, 5, 6
- [33] Sounak Mondal, Seoyoung Ahn, Zhibo Yang, Niranjan Balasubramanian, Dimitris Samaras, Gregory Zelinsky, and Minh Hoai. Look hear: Gaze prediction for speech-directed human attention. In *European Conference on Computer Vision*, pages 236–255. Springer, 2024. 5
- [34] Sounak Mondal, Naveen Sendhilnathan, Ting Zhang, Yue Liu, Michael Proulx, Michael Louis Iuzzolino, Chuan Qin, and Tanya R Jonker. Gaze-language alignment for zero-shot prediction of visual search targets from human gaze scanpaths. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2738–2749, 2025. 5
- [35] Saul B. Needleman and Christian D. Wunsch. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *Journal of Molecular Biology*, 48(3):443–453, 1970. 6
- [36] Dang-Man Nguyen, Xuan-Thang Tran, Duc-Vinh Vo, and Van-Nam Huynh. Umcap: User memory augmented method for personalized image descriptions. In *International Symposium on Knowledge and Systems Sciences*, pages 139–154. Springer, 2024. 2
- [37] Thao Nguyen, Haotian Liu, Yuheng Li, Mu Cai, Utkarsh Ojha, and Yong Jae Lee. Yo’llava: Your personalized language and vision assistant. *Advances in Neural Information Processing Systems*, 37:40913–40951, 2024. 2
- [38] Thao Nguyen, Krishna Kumar Singh, Jing Shi, Trung Bui, Yong Jae Lee, and Yuheng Li. Yo’chameleon: Personalized vision and language generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 14438–14448, 2025. 2
- [39] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022. 4
- [40] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002. 2, 5
- [41] Cesc Chunseong Park, Byeongchang Kim, and Gunhee Kim. Towards personalized image captioning via multi-modal memory networks. *IEEE transactions on pattern analysis and machine intelligence*, 41(4):999–1012, 2018. 1, 2, 6
- [42] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, 2018. 4
- [43] Chau Pham, Hoang Phan, David Doermann, and Yunjie Tian. Plvm: A tuning-free approach for personalized large vision-language model. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 3632–3641, 2025. 2
- [44] Renjie Pi, Jianshu Zhang, Tianyang Han, Jipeng Zhang, Rui Pan, and Tong Zhang. Personalized visual instruction tuning. *arXiv preprint arXiv:2410.07113*, 2024. 2
- [45] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649, 2015. 5
- [46] Jordi Pont-Tuset, Jasper Uijlings, Soravit Changpinyo, Radu Soricut, and Vittorio Ferrari. Connecting vision and language with localized narratives. In *European conference on computer vision*, pages 647–664. Springer, 2020. 1, 2, 5, 6, 7
- [47] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22500–22510, 2023. 2
- [48] Fiona Ryan, Josef Sivic, Fabian Caba Heilbron, Judy Hoffman, James M Rehg, and Bryan Russell. Improving personalized search with regularized low-rank parameter updates. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 19748–19757, 2025. 2
- [49] Gerard Salton and Chris Buckley. Term-weighting approaches in automatic text retrieval. *Information Processing & Management*, 24(5):513–523, 1988. 2
- [50] Nikshep Shetty and Yongmin Li. Detailed image captioning and hashtag generation. *Future Internet*, 16(12):444, 2024. 1
- [51] Kurt Shuster, Samuel Humeau, Hexiang Hu, Antoine Bordes, and Jason Weston. Engaging image captioning via personality. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12516–12526, 2019. 2
- [52] Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafranec, Seungeun Yi, Michaël Ramamonjisoa,

- Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski. DINOv3, 2025. 5
- [53] Abigale Stangl, Nitin Verma, Kenneth R Fleischmann, Meredith Ringel Morris, and Danna Gurari. Going beyond one-size-fits-all image descriptions to satisfy the information wants of people who are blind or have low vision. In *Proceedings of the 23rd international ACM SIGACCESS conference on computers and accessibility*, pages 1–15, 2021. 1
- [54] Ece Takmaz, Sandro Pezzelle, Lisa Beinborn, and Raquel Fernández. Generating image descriptions via sequential cross-modal alignment guided by human gaze. *arXiv preprint arXiv:2011.04592*, 2020. 1, 2
- [55] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 3
- [56] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015. 2, 5
- [57] Yuiga Wada, Kanta Kaneda, Daichi Saito, and Komei Sugiura. Polos: Multimodal Metric Learning from Human Feedback for Image Captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 5
- [58] Leiquan Wang, Xiaoliang Chu, Weishan Zhang, Yiwei Wei, Weichen Sun, and Chunlei Wu. Social image captioning: Exploring visual attention and user attention. *Sensors*, 18(2):646, 2018. 2
- [59] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution, 2024. 5, 6
- [60] Teng Wang, Jinrui Zhang, Junjie Fei, Hao Zheng, Yunlong Tang, Zhe Li, Mingqi Gao, and Shanshan Zhao. Caption anything: Interactive image description with diverse multimodal controls. *arXiv preprint arXiv:2305.02677*, 2023. 2
- [61] Xuan Wang, Guan hong Wang, Wenhao Chai, Jiayu Zhou, and Gaoang Wang. User-aware prefix-tuning is a good learner for personalized image captioning. In *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*, pages 384–395. Springer, 2023. 1, 2
- [62] Yuhang Wang, Canlong Zhang, Zhiwen Wang, and Zhixin Li. Image captioning according to user’s intention and style. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–9. IEEE, 2022. 2
- [63] Zeyu Wang, Yuanchun Shi, Yuntao Wang, Yuchen Yao, Kun Yan, Yuhan Wang, Lei Ji, Xuhai Xu, and Chun Yu. G-voila: gaze-facilitated information querying in daily scenarios. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 8(2):1–33, 2024. 2
- [64] Kun Xiong, Liu Jiang, Xuan Dang, Guolong Wang, Wenwen Ye, and Zheng Qin. Towards personalized aesthetic image caption. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2020. 1, 2
- [65] Ruoyu Xue, Jingyi Xu, Sounak Mondal, Hieu Le, Greg Zelinsky, Minh Hoai, and Dimitris Samaras. Few-shot personalized scanpath prediction. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 13497–13507, 2025. 1, 2
- [66] Kun Yan, Lei Ji, Huaishao Luo, Ming Zhou, Nan Duan, and Shuai Ma. 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)*, pages 2014–2025, 2021. 2
- [67] Zhibo Yang, Sounak Mondal, Seoyoung Ahn, Ruoyu Xue, Gregory Zelinsky, Minh Hoai, and Dimitris Samaras. Unifying top-down and bottom-up scanpath prediction using transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1683–1693, 2024. 5
- [68] Chun-Hsiao Yeh, Bryan Russell, Josef Sivic, Fabian Caba Heilbron, and Simon Jenni. Meta-personalizing vision-language models to find named instances in video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19123–19132, 2023. 2
- [69] Gregory J Zelinsky, Seoyoung Ahn, Zhibo Yang, Yupei Chen, Sounak Mondal, Minh Hoai, and Dimitrios Samaras. Reward maps predict target-present and target-absent visual search. *Journal of Vision*, 23(9):5161–5161, 2023. 5
- [70] Wenhuan Zeng, Abulikemu Abuduweili, Lei Li, and Pengcheng Yang. Automatic generation of personalized comment based on user profile. *arXiv preprint arXiv:1907.10371*, 2019. 2
- [71] Wei Zhang, Yue Ying, Pan Lu, and Hongyuan Zha. Learning long-and short-term user literal-preference with multimodal hierarchical transformer network for personalized image caption. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9571–9578, 2020. 2
- [72] Lian Zhou, Yuejie Zhang, Yu-Gang Jiang, Tao Zhang, and Weiguo Fan. Re-caption: Saliency-enhanced image captioning through two-phase learning. *IEEE Transactions on Image Processing*, 29:694–709, 2019. 1, 2