

CapeNext: Rethinking and refining dynamic support information for category-agnostic pose estimation

Yu Zhu^{1,2}, Dan Zeng^{1*}, Shuiwang Li³, Qijun Zhao⁴, Qiaomu Shen⁵, Bo Tang²

¹School of Artificial Intelligence, Sun Yat-sen University, Zhuhai 510275, China

²Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen 518055, China

³College of Computer Science and Engineering, Guilin University of Technology, Guilin 541004, China

⁴College of Computer Science, Sichuan University, Chengdu 610065, China

⁵Aerospace and Informatics Domain, Beijing Institute of Technology, Zhuhai 519088, China

Abstract

Recent research in Category-Agnostic Pose Estimation (CAPE) has adopted fixed textual keypoint description as semantic prior for two-stage pose matching frameworks. While this paradigm enhances robustness and flexibility by disentangling the dependency of support images, our critical analysis reveals two inherent limitations of static joint embedding: (1) polysemy-induced cross-category ambiguity during the matching process (e.g., the concept "leg" exhibiting divergent visual manifestations across humans and furniture), and (2) insufficient discriminability for fine-grained intra-category variations (e.g., posture and fur discrepancies between a sleeping white cat and a standing black cat). To overcome these challenges, we propose a new framework that innovatively integrates hierarchical cross-modal interaction with dual-stream feature refinement, enhancing the joint embedding with both class-level and instance-specific cues from textual description and specific images. Experiments on the MP-100 dataset demonstrate that, regardless of the network backbone, CapeNext consistently outperforms state-of-the-art CAPE methods by a large margin.

1 Introduction

Category-Agnostic Pose Estimation (CAPE), as a fundamental task in computer vision, has demonstrated significant values in application ranging from action recognition to human-computer interaction for arbitrary category (Minderer et al. 2022; Zhang et al. 2023; Zareian et al. 2021). POMNet (Xu et al. 2022) reformulates category-agnostic pose estimation as a keypoint matching problem, where the model learns to transfer keypoint knowledge from annotated support images to query images of unseen categories. Subsequent researches (Shi et al. 2023; Hirschorn and Avidan 2024; Chen et al. 2024; Ren et al. 2024) train CAPE models in a two-stage framework, generating similarity-aware position proposals and then refining them through a transformer decoder. However, the model performance in this paradigm

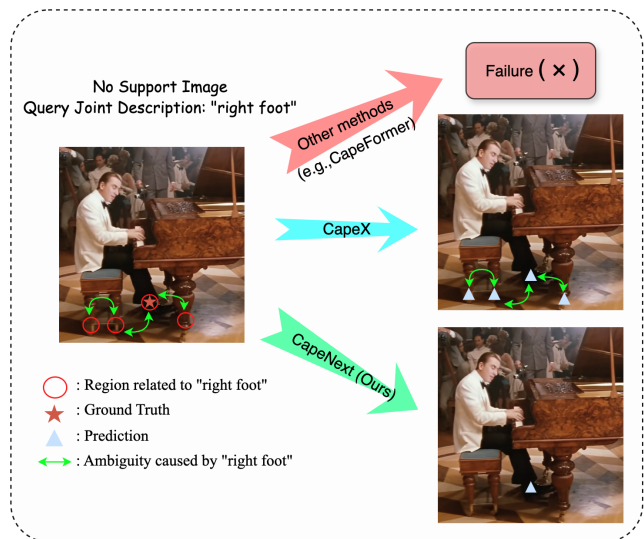


Figure 1: Conventional CAPE methods can not work when there is no support image but only textual keypoint descriptions, CapeX will be affected by the textual ambiguity to match the wrong candidate points, while our method can better avoid this problem.

is affected to a greater extent by the quality of the support image, and even when there are invisible keypoints in the support image, it fails to provide effective keypoint knowledge for the query image. To address this issue, CapeX (Rusanovsky, Hirschorn, and Avidan 2025) replaces unstable keypoint knowledge from support images with stable text feature of keypoint descriptions encoded by LLM.

However, compared to support images, relying solely on textual keypoint descriptions can lead to inherent ambiguity in complex scenarios. For example, in Figure 1, when searching for "right foot", there are multiple candidates that match the query, which cannot be effectively distinguished using text alone. Let's take a step back and consider how humans approach this task. When searching for the "right foot", we subconsciously focus on the pianist's shoes, mak-

*Corresponding author: zengd8@mail.sysu.edu.cn

Codes and supplementary materials of this paper are available at <https://github.com/yzrs/CapeNext>

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

ing it unlikely for us to mistake the chair’s leg for a foot. That’s because we instinctively associate “right foot” with concepts like “human”, “leg”, and “shoes” using them as additional contextual cues to guide our final decision while naturally disregarding the chair’s foot. However, obtaining rich relevant cues for CAPE is both cumbersome and impractical for model training. Beyond additional associative information, humans naturally perceive the overall context of an image, even when not specifically searching for keypoints.

Inspired by this, we propose leveraging the perceptual understanding of the query image as contextual support and utilizing easily accessible category information as associative information. The former introduces the visual features of the query image and reduces the difficulty of matching between the support text and the query image in the case of large intra-category instance variance. It adds the detail information from query images to keypoint features to help the keypoints dynamically adapt to different query images. The latter provides us with class information that can help filter the noise information of other non-target objects in the query image, enhance the feature expression of the subject and diminish polysemy-induced cross-category ambiguity. It is worth noting that we are the first to incorporate the query image itself as additional “support” information to construct keypoint embeddings.

However, effectively integrating the query image, class description, and keypoint description to improve CAPE is challenging, primarily due to the modality differences between class descriptions and query image features. Direct feature fusion can also introduce noise, adversely affecting model performance. This is because the query image contains relevant target information and extraneous noise. Moreover, the detailed features in class descriptions may conflict with specific attributes of the query image, further complicating effective fusion.

To this end, we propose **CapeNext** to elevate CAPE’s performance to the **Next** level by the strengths of different modalities while mitigating potential feature conflicts. Specially, CapeNext contains two innovative modules: Hierarchical Cross-Modal Interaction (HCMI) and Dual-Stream Feature Refinement (DSFR). These two modules align the feature from both the specific image and the general class representation, leveraging the images’ strengths in capturing fine-grained, instance-level details and text’s ability to encode broad, coarse-grained semantics, refining the keypoint feature from different aspects. CapeNext introduces multi-modal input and innovative feature refinement modules to dynamically adopt the joint embedding of the fixed keypoint descriptions to be closer to the specific query image.

To sum up, this paper makes the following contributions:

- We rethink the support information in CAPE and propose CapeNext, which is the first to incorporate query information as “support” to facilitate the dynamic feature learning of keypoints.
- We introduce two novel modules, Hierarchical Cross-Modal Interaction (HCMI) and Dual-Stream Feature Refinement (DSFR), which effectively leverage the class-level and instance-specific cues to produce constant text-

modal joint embedding that resolves polysemy-induced cross-category ambiguity and enhances discriminability for fine-grained intra-category variations.

- Our experiments on the MP-100 dataset demonstrate that, regardless of the network backbone, CapeNext consistently outperforms state-of-the-art CAPE methods.

2 Related Work

Category-Agnostic Pose Estimation (CAPE) removes category restrictions during testing, enabling pose prediction for arbitrary targets, different from conventional pose estimation methods (Cao et al. 2019; Yang et al. 2023; Sun et al. 2019; Xiao, Wu, and Wei 2018; Ding et al. 2022; Li et al. 2021; Andriluka et al. 2018). Its generalization aligns with real-world needs, where models must adapt to diverse unseen categories without category-specific training data.

POMNet (Xu et al. 2022), the first category-agnostic method, frames pose estimation as embedding-space semantic matching between support keypoint and query image features but suffers from unreliable results due to visually similar irrelevant regions (e.g., mirror keypoints) and insufficient similarity under large pose/texture/style differences. CapeFormer (Shi et al. 2023) pioneers a two-stage transformer framework: generating similarity-aware initial keypoint proposals via improved matching, then refining them by aggregating context through Transformer Decoder. X-Pose (Yang et al. 2024) leverages CLIP (Radford et al. 2021) vision-language info via support visual prompts and structured text prompts (“An [style] photo of [object]/[part]/[keypoint]”), generating semantic embeddings to boost generalization to unseen categories via flexible prompt use. SDPNet (Ren et al. 2024) models joint dependencies via skeleton graphs and GCN (with self-attention for adjacency matrices) but has suboptimal efficiency due to separate GCN training. GraphCape (Hirschorn and Avidan 2024) replaces transformer feed-forward layers with GCN for end-to-end training (reducing overhead). CapeX (Rusanovsky, Hirschorn, and Avidan 2025) uses CLIP text prompts to match query features, cutting support image dependency and improving accuracy.

These existing methods share a paradigm of enriching support image keypoint representations via skeleton graph text prompts to boost CAPE accuracy, but overlook the query image’s rich semantic information that can refine support data. Our method fully leverages query image information and dynamically adjusts support sample keypoint features accordingly.

3 Our Method: CapeNext

3.1 Preliminaries

Category-Agnostic Pose Estimation aims to estimate desired keypoints on the query image given a support image annotated with keypoints for arbitrary category. Formally, under the 1-shot setting (one support image), given a support set $S = \{I_s, K_s\}$ where $I_s \in \mathbb{R}^{H \times W \times 3}$ denotes the support image and $K_s = \{k_s^j\}_{j=1}^N$ represents its corresponding N keypoint coordinates, CAPE predicts the target keypoint set $\hat{K}_q = \{\hat{k}_q^j\}_{j=1}^N$ from a query image $I_q \in \mathbb{R}^{H \times W \times 3}$. The

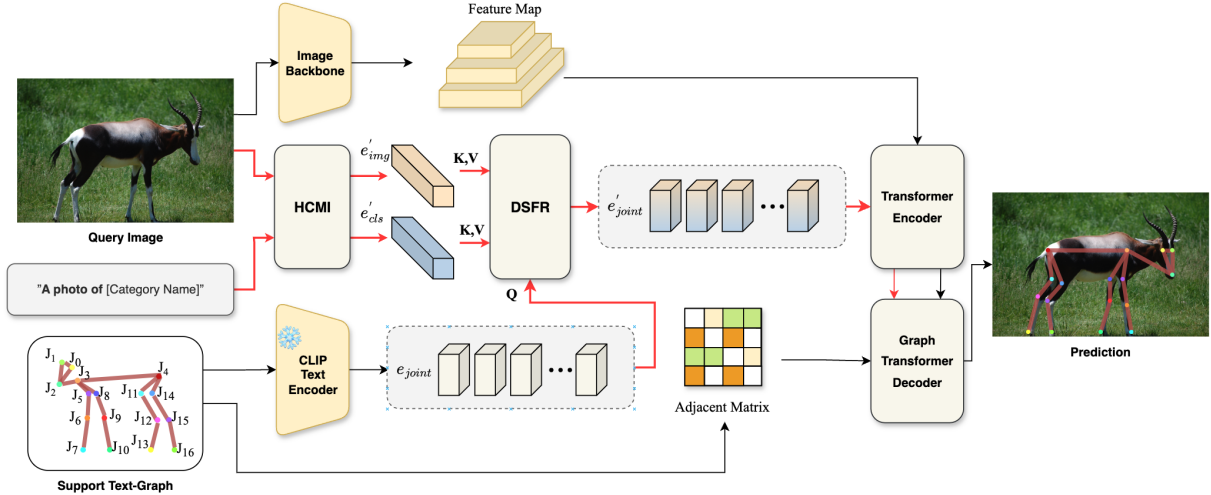


Figure 2: **Framework Overview of CapeNext.** CapeNext employs a dual-stream architecture to process multimodal inputs.

core idea is to leverage explicitly defined keypoint features from support samples to guide the model’s search for similar feature patterns in query images, thereby disentangling category-specific prior knowledge dependencies. This process could be formulated as follows:

$$K_q = M_\theta(I_q, \{I_s, K_s\}), \quad (1)$$

where K_q indicates the predicted keypoints and M_θ is the CAPE model used here.

CapeX propose a novel way for CAPE by replacing support keypoint feature with LLM-encoded text embedding, guiding the model to search for visual feature that are similar to text embedding in query images. Formally, it can be described as follows:

$$K_q = M_\theta(I_q, \{K_s^{text}\}), \quad (2)$$

where $K_s^{text} = \{K_s^{text_j}\}_{j=1}^N$ indicates the support joint embedding obtained from LLM encodings.

This paradigm relies solely on textual keypoint descriptions rather than specific support images, eliminating dependency on support image quality. It is useful for unseen categories where annotated support samples are typically unavailable, demonstrating superior practicality in real-world applications. Simultaneously, benefiting from the semantic extensibility of textual descriptions, this paradigm based on textual keypoint representations also demonstrates great flexibility. Our method integrates this framework, inheriting the advantages of representational stability and input flexibility inherent in text-driven approaches.

3.2 Motivation

Previous methods benefit from the flexibility of using joint textual descriptions as support, but they inherit a fundamental drawback: text cannot capture fine-grained visual details as rich as images. For example, the textual representations “tail” or even “tail of a cat” can’t accurately match a black cat’s tail or a white cat’s tail at the same time without a more detailed visual representations.

In the CAPE domain, its reliance on unconstrained textual keypoint definitions also introduces two fundamental limitations:(1) **Polysemy-induced Cross-category Ambiguity:** identical textual descriptions (e.g., “left leg”) may refer to anatomically distinct parts across categories (e.g., human legs vs. chair legs), causing the model to conflate heterogeneous geometric structures due to their semantic overlap in the text embedding space.(2) **Intra-category Variance:** even within a single category, instances with diverse appearances (e.g., a standing black cat’s “right front paw” vs. a curled white cat’s “right front paw”) are forced to share the same text embedding, suppressing the model’s ability to recognize the pose-dependent visual variations.

As illustrated in Figure 2, our framework CapeNext consists of three parts: (1) multimodal feature extraction and refinement, (2) feature matching and graph-based structural reasoning, and (3) keypoint prediction. Our key contribution lies in the design and refinement of multimodal feature which enhance support joint feature with both class-level and instance-specific cues. Transformer encoder and graph transformer decoder share the same architectures as Graph-Cape and CapeX.

3.3 Method Overview

Given a query image, we first extract its multi-scale visual features using a network backbone (e.g., HRNet-w32, ViT-Base-16, DINOv2-ViT-S and Swinv2-T). These features capture both local part details and global contextual patterns. To overcome polysemy-induced cross-category ambiguity and intra-category instance variation issues in CAPE, we derive an enhanced joint embedding through our multimodal input: textual class descriptions, which are easy to obtain, along with the query image itself as complementary information for keypoint and textual class descriptions.

The effectiveness of multimodal input comes from two aspects. First, textual class descriptions can mitigate polysemy-induced cross-category ambiguity. By incorporating textual class descriptions, the model resolves seman-

tic ambiguities arising from identical textual representations across different categories. For example, when multiple objects in an image share identical keypoint descriptions but differ semantically, textual class descriptions can act as discriminative filters. These descriptions suppress interference from the regions irrelevant to the targets while constraining the semantic positioning space for valid keypoints. Second, visual feature integration from query images addresses the issue of intra-category instance variance. Consider the query images of black cat and white cat: Neither original joint embeddings nor textual class description encode the uncertain visual attributes such as the fur color. When the query image shifts between black cat and white cat, the fixed keypoint representations struggle to adapt to the changed visual feature. But the integration of query image embedding can compensate for the joint embedding inaccuracies caused by significant intra-category variations because the target feature is just in the query image itself. It provides complementary instance-aware details, which are missing in textual class descriptions, and thus effectively balancing semantic precision and visual fidelity.

We use pretrained CLIP to encode the textual keypoint descriptions to get the corresponding joint embedding and our multimodal input (textual class description and query image) to get our class embedding and image embedding. Both class embedding and image embedding contain the query image feature while the former is more related to the general category information and the latter is more related to the specific instance feature. To this end, we proposed **Hierarchical Cross-Modal Interaction and Dual-Stream Feature Refinement** to enhance the joint embedding for CAPE. First, Hierarchical Cross-Modal Interaction (HCMI) amplifies dominant instance patterns in the image embedding by the class embedding and enriches class embedding with instance-adaptive details from the image embedding. Then, Dual-Stream Feature Refinement (DSFR) performs dual cross-attention between the original joint embedding and the refined class/image embedding, yielding residual-enhanced embedding. The computation of HCMI and DSFR can be described as Eq.3 and 4, respectively. The final enhanced joint embedding is fed into the subsequent transformer encoder-decoder module.

$$e'_{img}, e'_{cls} = HCMI(e_{img}, e_{cls}), \quad (3)$$

$$e'_{joint} = e_{joint} + DSFR(e_{joint}, e'_{img}, e'_{cls}), \quad (4)$$

where e_{img} and e_{cls} represent the image embedding and class embedding encoded from CLIP, e'_{img} and e'_{cls} represent the image embedding and class embedding enhanced by HCMI. e_{joint} and e'_{joint} represent the joint embedding encoded from CLIP and the joint embedding refined by DSFR.

Once the refined joint embedding with class and instance info have been learned, we employ a graph transformer decoder, to propagate node features through graph connections. Updated graph node features are decoded into similarity heatmaps and offset maps through parallel CNN heads. We extract the heatmap peaks and apply offset correction to obtain the final joint coordinates.

3.4 Hierarchical Cross-Modal Interaction (HCMI)

Although a certain degree of feature alignment has been done between image embedding and class embedding during the pre-training process of CLIP, there are still differences in feature representation between these two. The image embedding e_{img} contains more information about the specific visual features of the query image at the lower hierarchy and introduces more image noise when used alone, while the class embedding e_{cls} at the higher hierarchy contains more general and comprehensive information about the query image of that class and cannot be adapted to the specific query image more flexibly when used alone.

To bridge the hierarchy and modality gap between these two visual and semantic embeddings, we propose to interactively refine the image embedding e_{img} and class embedding e_{cls} through a self-attention mechanism, as illustrated in Figure 3. This module processes two inputs: (1) one concrete query image from the visual modality and (2) universal category description from the textual modality. Both of them are independently encoded by CLIP to generate corresponding image embedding e_{img} and class embedding e_{cls} . They are concatenated and undergo cross-modal interaction through self-attention mechanism, followed by multilayer perceptron layers (MLP) to produce enhanced embedding. The output is spliced to produce enhanced image embedding e'_{img} and class embedding e'_{cls} . This interaction is formulated as follows:

$$e'_{img}, e'_{cls} = MLP(SelfAttn(concat(e_{img}, e_{cls}))), \quad (5)$$

where $SelfAttn(\cdot)$ denotes the self-attention operation that captures cross-modal dependencies, followed by a MLP module with ReLU-activated fully-connected layers. e'_{img} and e'_{cls} indicate the enhanced image embedding and class embedding, which have reduced inter-modality discrepancy through the attention alignment and mutually enriched semantic-visual information via cross-modal feature propagation. Their compatibilities for subsequent joint embedding construction are also enhanced in this interaction process.

3.5 Dual-Stream Feature Refinement (DSFR)

Building upon the enhanced image embedding and class embedding above, the dual-stream feature refinement module dynamically adapts the CLIP-generated joint embedding to enhance cross-modal alignment between the query image’s visual feature, the category feature and the support joint feature. As shown in Figure 3, DSFR module processes three inputs: (1) the text-modality joint embedding $e_{joint} \in \mathbb{R}^{N \times C}$ derived from textual keypoint descriptions, (2) the text-modality class embedding $e'_{cls} \in \mathbb{R}^C$ derived from image caption and refined by HCMI module, (3) the image-modality image embedding $e'_{img} \in \mathbb{R}^C$ encoded from CLIP and refined by HCMI module. After the cross-modal feature Refinement, DSFR module outputs an augmented joint embedding $e'_{joint} \in \mathbb{R}^{N \times C}$, where C denotes the embedding dimension of the LLM encoder and N represents the predefined keypoint number.

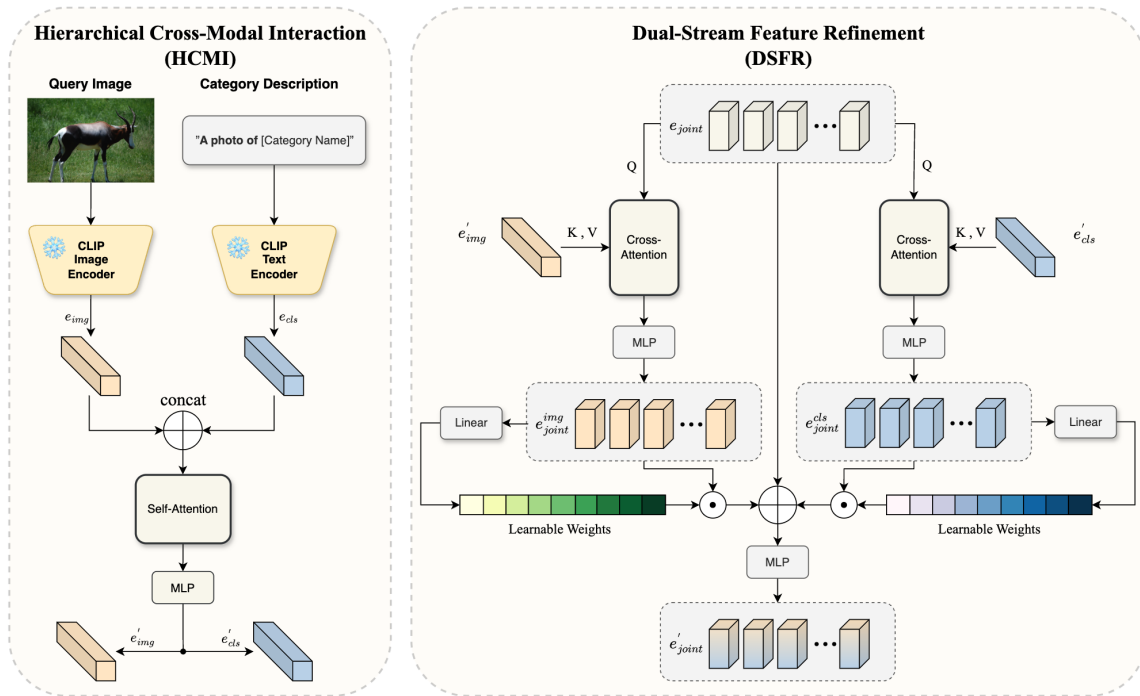


Figure 3: Left: HCMI mediates hierarchical and cross-modal discrepancies between image and class embeddings. Right: DSFR leverages the class-level feature and instance-aware attributes from class and image embeddings to refine the joint embedding.

The DSFR module employs a dual-stream cross-attention mechanism where original joint embedding serves as the query stream, while enhanced image embedding and enhanced class embedding respectively serves as the key-value stream for two parallel cross-attention computations. This is to enable original joint embedding, which is full of key-point feature information, to dynamically retrieve instance-specific visual features from enhanced image embedding and general class-aware attributes from enhanced class embedding. This refinement can bridge textual joint priors with specific query image feature through attention-based feature correlation effectively. We use ReLU-activated MLP layers to process the outputs of the cross-attention and get the refined joint embedding. This refined joint embedding conditioned on the query image embedding and class embedding is formally denoted as: $e^{img}_{joint} \in \mathbb{R}^{N \times C}$ and $e^{cls}_{joint} \in \mathbb{R}^{N \times C}$. This attention computation can be described as follows:

$$e^{img}_{joint} = \text{MLP}(\text{CrossAttn}(q = e_{joint}, k, v = e'_{img})) \quad (6)$$

$$e^{cls}_{joint} = \text{MLP}(\text{CrossAttn}(q = e_{joint}, k, v = e'_{cls})) \quad (7)$$

To incorporate the information from e^{img}_{joint} and e^{cls}_{joint} more flexibly, we employ two independent gating networks to adaptively weight the enhanced joint embeddings of e^{img}_{joint} and e^{cls}_{joint} . Gating networks produce joint-wise scores of the cross-modal feature using linear layers with sigmoid activation. The refined joint embedding is subsequently computed as follow:

$$e'_{joint} = \text{MLP}(\alpha \odot e^{img}_{joint} + \beta \odot e^{cls}_{joint} + e_{joint}), \quad (8)$$

where $\alpha, \beta \in [0, 1]^N$ denote learnable coefficients, \odot indicates Hadamard product and the MLP contains ReLU-activated hidden layers for nonlinear projection.

3.6 Loss Function

We use a heatmap loss $\mathcal{L}_{heatmap}$ and an offset loss \mathcal{L}_{offset} to train our model. $\mathcal{L}_{heatmap}$ supervises the similarity maps with the shape of $H \times W$ obtained from proposal generator while \mathcal{L}_{offset} supervises the decoder's localization output. They can be described as follows:

$$\mathcal{L}_{heatmap} = \frac{1}{N} \sum_{i=1}^N \frac{1}{H \cdot W} \|\text{sigmoid}(H_i) - \hat{H}_i\| \quad (9)$$

$$\mathcal{L}_{offset} = \frac{1}{L} \sum_{l=1}^L \sum_{i=1}^N |P_i^l - \hat{P}_i| \quad (10)$$

where H_i is the output similarity heatmap of the i -th joint from the proposal generator, \hat{H}_i is the corresponding ground truth heatmap, P_i^l is the output location of the l -th layer in the decoder with L layers and \hat{P}_i is the corresponding ground truth location.

The total loss for training can be formatted as follows:

$$\mathcal{L} = \lambda_{heatmap} \cdot \mathcal{L}_{heatmap} + \mathcal{L}_{offset} \quad (11)$$

where the loss weight $\lambda_{heatmap}$ for the heatmap loss is 2.

4 Experiment

4.1 Implementation Details

Dataset and Metric. We perform experiments on MP-100 dataset (Xu et al. 2022). The dataset consists of more than

Methods	Img Backbone	CLIP Backbone	Split1	Split2	Split3	Split4	Split5	Avg
POMNet	ResNet-50	-	84.23	78.25	78.17	78.68	79.17	79.70
CapeFormer	ResNet-50	-	89.45	84.88	83.59	83.53	85.09	85.31
ESCAPE	ResNet-50	-	86.89	82.55	81.25	81.72	81.32	82.74
MetaPoint+	ResNet-50	-	90.43	85.59	84.52	84.34	85.96	86.17
X-Pose	ResNet-50	ViT-Base-32	89.07	85.05	85.26	85.52	85.79	86.14
SDPNet	HRNet-32	-	91.54	86.72	85.49	85.77	87.26	87.36
GraphCape	Swinv2-T	-	91.19	87.81	85.68	85.87	85.61	87.23
CapeX	HRNet-w32	ViT-Base-32	89.1	85.0	81.9	84.4	85.4	85.2
CapeX	ViT-Base-16	ViT-Base-32	90.75	82.87	83.18	85.95	85.49	85.65
CapeX	DINOv2-ViT-S	ViT-Base-32	90.6	83.74	83.67	86.87	85.93	86.18
CapeX	Swinv2-T	ViT-Base-32	91.9	86.97	84.41	86.13	88.64	87.61
CapeNext	HRNet-w32	ViT-Base-32	90.2	86.0	82.9	85.4	87.1	86.3
CapeNext	ViT-Base-16	ViT-Base-32	90.84	86.73	86.5	82.44	87.91	86.88
CapeNext	DINOv2-ViT-S	ViT-Base-32	92.12	87.75	83.76	87.16	88.95	87.95
CapeNext	Swinv2-T	ViT-Base-32	92.44	87.31	85.44	86.47	90.17	88.37

Table 1: PCK@0.2 performance on MP-100 test dataset. CapeNext outperforms other methods in the 1-shot setting.

Settings	Split1	Split2	Split3	Split4	Split5	Avg
CapeNext	92.44	87.31	85.44	86.47	90.17	88.37
CapeNext w/o HCMI	91.43	83.4	84.97	87.27	86.98	86.81
CapeNext w/o DSFR	89.04	82.72	82.94	82.97	83.58	84.25
CapeNext w/o LW	92.59	86.8	85.03	86.52	90.5	88.28

Table 2: Ablation study for our module design.

18K images of 100 categories from 8 supercategories, including human hand, human face, human body, animal face, clothes, furniture, and vehicle. We also follow five splits for standard CAPE settings. Each split divides the dataset into non-overlap 70 training categories, 10 validation categories, and 20 testing categories. We also adopt the MP-100 dataset with unified skeleton definitions from GraphCape and keypoint text descriptions extended in CapeX. We primarily report probability of correct keypoint with a threshold of 0.2 (PCK@0.2) for evaluation. We also report the PCK performance of other thresholds (e.g., 0.05, 0.1, 0.15, 0.2, 0.25 and the average PCK performance of all thresholds) for comparison in our supplementary material.

Training Details. We use the pretrained Tiny-Swin-Transformer(Swinv2-T) (Liu et al. 2021) as our default feature extraction backbone and pretrained ViT-Base-32 (Dosovitskiy et al. 2020) as the frozen CLIP backbone. We also report results using ViT-Base-16 and HRNet-w32 as feature extraction backbones to validate the generalizability. CapeNext is trained end-to-end by Adam optimizer for 200 epochs with the batch size of 32 on one A100 GPU. The initial learning rate is 10^{-5} , reducing by a factor of 10 at the 160th and 180th epochs, following the conventional setting.

4.2 Main Results

We compare CapeNext with eight SOTA methods including POMNet, CapeFormer, ESCAPE (Nguyen, Li, and Lee

2024), MetaPoint, X-Pose (Yang et al. 2024), SDPNet, GraphCape and our baseline CapeX. As Table 1 shows, with Swinv2-T as backbone, CapeNext outperforms other methods by an average of 0.76% under the 1-shot setting, validating the efficacy of our HCMI module and DSFR module. More results with HRNet-w32, ViT-Base-16 and DINOv2-ViT-S as backbone are also provided.

4.3 Ablation Study

Effect of CapeNext Design. Table 2 represents the performance of CapeNext using HCMI, DSFR, and learnable weights (LW). Regarding “CapeNext w/o HCMI”, we remove the HCMI module entirely and directly feed the image embedding (obtained from CLIP image encoding) and the class embedding (obtained from CLIP text encoding) into the subsequent DSFR module. In this setting, original image embedding and original class embedding serve as the key and value, respectively, and compute cross-attention with the joint embedding as query. “CapeNext w/o HCMI” leads to a decrease of 1.56% in the average PCK@0.2 performance, further demonstrating its effectiveness on the reverse side. Regarding “CapeNext w/o DSFR”, we replace DSFR module with two Relu-activated linear layers that perform projection transformations. These layers project the image embedding and class embedding separately, and the results are added to the joint embedding. For learnable weights, we replace it with a fixed value of 1, which corresponds to “CapeNext w/o LW”. LW (learnable weights) is employed to balance the influence weights between image embeddings and class embeddings. We find that the LW lead to a modest improvement, although the effect is relatively small, with an average accuracy increase of only 0.09%.

Insights of Multimodal Input. To demonstrate the effectiveness of our multimodal input data, we conducted ablation studies on both the image embedding and class embedding. The baseline setting refers to the official implementation in CapeX. For image embedding, we added its compu-

Settings	Split1	Split2	Split3	Split4	Split5	Avg
Baseline	91.9	86.97	84.41	86.13	88.64	87.61
+img emb	92.77	87.36	84.55	86.31	89.76	88.15
+img & cls emb	92.44	87.31	85.44	86.47	90.17	88.37

Table 3: Ablation study for our multimodal input.

Settings	Split1	Split2	Split3	Split4	Split5	Avg
Ours w/o LW	92.23	86.18	84.64	84.96	89.91	87.58
Ours w/ LW	91.94	86.94	85.48	86.12	90.27	88.15(+0.67)

Table 4: PCK@0.2 results when using **noisy class prompts**.

Methods	keypoint typo	paw→foreleg	elbow→forearm
CapeX	64.10	86.77	87.37
CapeNext	66.35(+2.25)	87.30(+0.53)	88.05(+0.68)
Methods	eye→eyeball	paw→foot	shoulder→upper arm
CapeX	87.60	87.38	86.81
CapeNext	88.19(+0.59)	88.06(+0.68)	87.58(+0.77)

Table 5: PCK@0.2 results when using different **noisy keypoint prompts**.

tational contributions in both HCM1 and DSFR based on the baseline setting. A similar process was employed when we further incorporated the class embedding. As shown in Table 3, the introduction of image embedding significantly improved the accuracy of the model, enhancing 0.54% average performance of PCK@0.2 over the baseline. This demonstrates that our additional image embedding has a notable positive effect on optimizing the model. Furthermore, additional class embedding further improved the model, achieving an average accuracy of 88.37%, which further validates the effectiveness of our multimodal input.

Robustness to Noisy Prompts. We evaluate model under two noise conditions: (1) class prompt noise via random substitution with incorrect categories. With gating mechanism (LW), CapeNext exhibits strong class noise tolerance with merely 0.2% PCK@0.2 drop (88.3%→88.1%) as shown in Table 4. We also retrain the model without gating mechanism, average PCK@0.2 performance decrease to 87.5%, confirming the gate’s error suppression capability. (2) keypoint prompt noise using typo patterns (e.g., “left eey”, “nosse”) from CapeX. While both CapeNext and CapeX suffer performance degradation, CapeNext mitigates feature corruption, reducing error propagation from inaccurate keypoint features. To further quantify sensitivity, we conduct multiple cross-prompt-template evaluation by modifying keypoint descriptors for evaluation. For example, from “xx paw” to “xx foreleg” means models are trained with “xx paw” but tested with “xx foreleg”. PCK results in Table 5 show that both models suffer from this variant prompt template but CapeNext demonstrates relative robustness through cross-modal feature refinement.



Figure 4: **Qualitative results.** We visualize joint predictions on MP-100 test dataset. Zoom in for more details.

4.4 Qualitative Results.

Figure 4 presents qualitative comparisons between our method and the competing baseline CapeX. The qualitative results show that CapeNext outperforms the baseline, validating its effectiveness. Red circles highlight regions with significant differences. The examples cover animal face, animal body, furniture and clothes. However, due to the inherent complexity of the CAPE task, some joints exhibit inaccuracies (such as “top left side of the backseat” and “top right side of the backseat” of the chair in the forth case), highlighting the challenges in this domain.

5 Conclusion

In this paper, we revisit the role of support information in category-agnostic pose estimation and propose a novel framework, CapeNext, which is the first to simultaneously integrate query information and class descriptions to facilitate dynamic keypoint feature learning. We introduce two innovative modules, HCM1 and DSFR, designed to enhance keypoint embeddings by fully leveraging multiple modalities. Our CapeNext not only outperforms the state-of-the-art methods by a large margin, but also maintains remarkable robustness under noisy inputs, including noisy class prompts and typo keypoint prompts.

6 Acknowledgments

Bo Tang and Yu Zhu were partially supported by National Science Foundation of China (NSFC No. 62422206) and a research gift from AlayaDB Inc. Dan Zeng was supported by the National Natural Science Foundation of China (No. 62206123). This work is supported by the National Natural Science Foundation of China (62206123, 62466013, 62176170, 62176169) and the Sichuan Science and Technology Program (2025ZNSFSC0469).

References

- Andriluka, M.; Iqbal, U.; Insafutdinov, E.; Pishchulin, L.; Milan, A.; Gall, J.; and Schiele, B. 2018. PoseTrack: A benchmark for human pose estimation and tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5167–5176.
- Cao, Z.; Hidalgo, G.; Simon, T.; Wei, S.-E.; and Sheikh, Y. 2019. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1): 172–186.
- Chen, J.; Yan, J.; Fang, Y.; and Niu, L. 2024. Meta-Point Learning and Refining for Category-Agnostic Pose Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23534–23543.
- Ding, Y.; Deng, W.; Zheng, Y.; Liu, P.; Wang, M.; Cheng, X.; Bao, J.; Chen, D.; and Zeng, M. 2022. I²R-Net: intra-and inter-human relation network for multi-person pose estimation. *arXiv preprint arXiv:2206.10892*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Hirschorn, O.; and Avidan, S. 2024. A Graph-Based Approach for Category-Agnostic Pose Estimation. *arXiv:2311.17891*.
- Li, Y.; Zhang, S.; Wang, Z.; Yang, S.; Yang, W.; Xia, S.-T.; and Zhou, E. 2021. Tokenpose: Learning keypoint tokens for human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11313–11322.
- 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10012–10022.
- Minderer, M.; Gritsenko, A.; Stone, A.; Neumann, M.; Weissenborn, D.; Dosovitskiy, A.; Mahendran, A.; Arnab, A.; Dehghani, M.; Shen, Z.; et al. 2022. Simple open-vocabulary object detection. In *European Conference on Computer Vision*, 728–755. Springer.
- Nguyen, K. D.; Li, C.; and Lee, G. H. 2024. Escape: Encoding super-keypoints for category-agnostic pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23491–23500.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 8748–8763. PMLR.
- Ren, P.; Gao, Y.; Sun, H.; Qi, Q.; Wang, J.; and Liao, J. 2024. Dynamic support information mining for category-agnostic pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1921–1930.
- Rusanovsky, M.; Hirschorn, O.; and Avidan, S. 2025. CapeX: Category-Agnostic Pose Estimation from Textual Point Explanation. In *The Thirteenth International Conference on Learning Representations*.
- Shi, M.; Huang, Z.; Ma, X.; Hu, X.; and Cao, Z. 2023. Matching is not enough: A two-stage framework for category-agnostic pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7308–7317.
- Sun, K.; Xiao, B.; Liu, D.; and Wang, J. 2019. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5693–5703.
- Xiao, B.; Wu, H.; and Wei, Y. 2018. Simple baselines for human pose estimation and tracking. In *Proceedings of the European conference on computer vision*, 466–481.
- Xu, L.; Jin, S.; Zeng, W.; Liu, W.; Qian, C.; Ouyang, W.; Luo, P.; and Wang, X. 2022. Pose for everything: Towards category-agnostic pose estimation. In *European Conference on Computer Vision*, 398–416. Springer.
- Yang, J.; Zeng, A.; Liu, S.; Li, F.; Zhang, R.; and Zhang, L. 2023. Explicit box detection unifies end-to-end multi-person pose estimation. *arXiv preprint arXiv:2302.01593*.
- Yang, J.; Zeng, A.; Zhang, R.; and Zhang, L. 2024. X-pose: Detecting any keypoints. In *European Conference on Computer Vision*, 249–268. Springer.
- Zareian, A.; Rosa, K. D.; Hu, D. H.; and Chang, S.-F. 2021. Open-vocabulary object detection using captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14393–14402.
- Zhang, H.; Li, F.; Zou, X.; Liu, S.; Li, C.; Yang, J.; and Zhang, L. 2023. A simple framework for open-vocabulary segmentation and detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1020–1031.