

ProMotion: Prototypes As Motion Learners

Yawen Lu¹, Dongfang Liu^{2*}, Qifan Wang³, Cheng Han², Yiming Cui⁴, Zhiwen Cao⁵, Xueling Zhang²,
Yingjie Victor Chen¹, Heng Fan^{6*}

¹Purdue University

²Rochester Institute of Technology

³Meta AI

⁴Tiktok

⁵Adobe

⁶University of North Texas

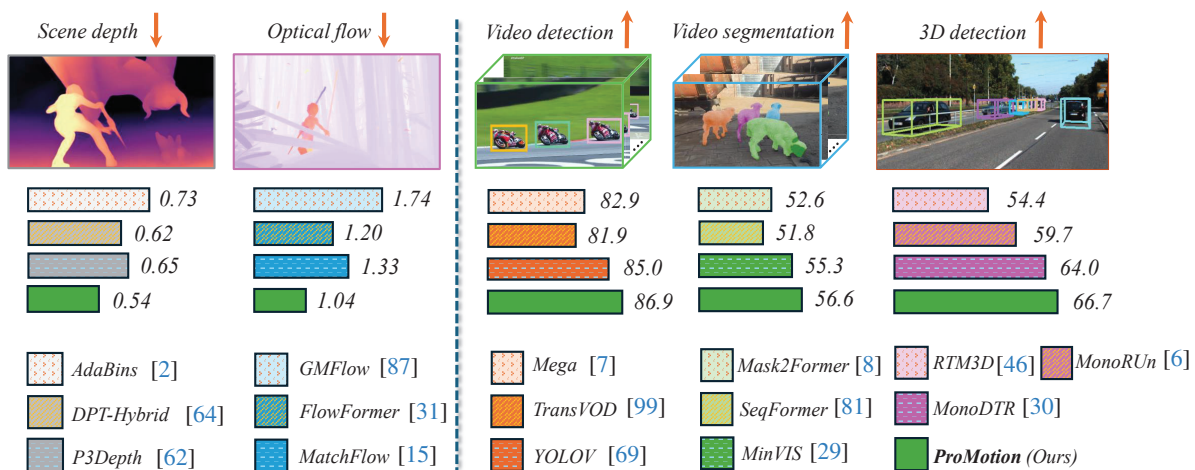


Figure 1. **PROMOTION** harmoniously handle core motion tasks (*i.e.*, optical flow and scene depth estimation), within an elegant prototype-based framework. This adaptability further enhances its utility in enabling seamless knowledge transfer to various downstream applications. For depth [2, 62, 64], we averagely reduce 0.13 *Abs Rel*; For flow [15, 31, 87], we averagely reduce 0.38 *AEPE*; For downstream tasks, we averagely achieve 3.6% *mAP* boost in video object detection [7, 69, 99], 3.4% *AP* boost in video object segmentation [8, 29, 81], and 7.3% *AP*_{3D} boost in 3D object detection [6, 30, 46].

Abstract

In this work, we introduce PROMOTION, a unified prototypical transformer-based framework engineered to model fundamental motion tasks. PROMOTION offers a range of compelling attributes that set it apart from current task-specific paradigms. ① We adopt a prototypical perspective, establishing a unified paradigm that harmonizes disparate motion learning approaches. This novel paradigm streamlines the architectural design, enabling the simultaneous assimilation of diverse motion information. ② We capitalize

on a dual mechanism involving the feature denoiser and the prototypical learner to decipher the intricacies of motion. This approach effectively circumvents the pitfalls of ambiguity in pixel-wise feature matching, significantly bolstering the robustness of motion representation. ③ We demonstrate a profound degree of transferability across distinct motion patterns. This inherent versatility reverberates robustly across a comprehensive spectrum of both 2D and 3D downstream tasks. Empirical results demonstrate that PROMOTION outperforms various well-known specialized architectures, achieving 0.54 and 0.054 *Abs Rel* error on the Sintel and KITTI depth datasets, 1.04 and 2.01 average

*Corresponding author.

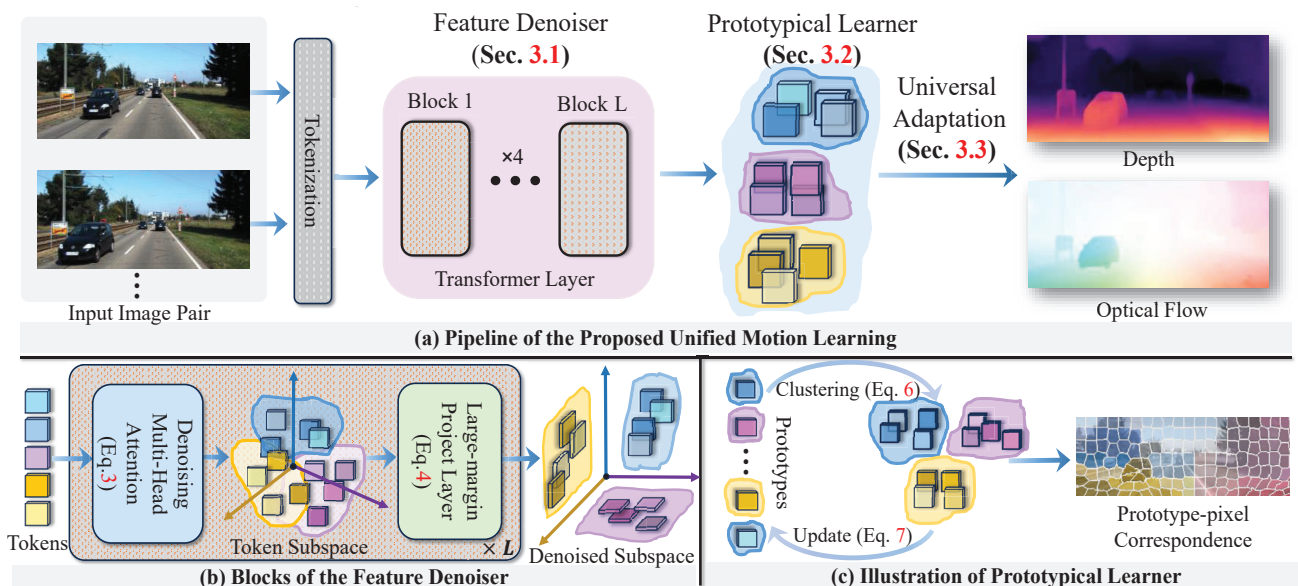


Figure 2. (a) Overall pipeline of PROMOTION. (b) Each Transformer block in *feature denoiser* maps the input tokens into different feature subspaces (Eq. 3) and then projects them to the orthogonal direction (Eq. 4), thereby mitigating the uncertainties in motion for robustness. (c) The *prototypical learner* clusters the subspace into prototypes (Eq. 6) and performs iterations to update them (Eq. 7). The learned prototype can capture different motion patterns, enabling representation learning for various dynamic characteristics.

endpoint error on the clean and final pass of Sintel flow benchmark, and 4.30 F1-all error on the KITTI flow benchmark. For its efficacy, we hope our work can catalyze a paradigm shift in universal models in computer vision.

1. Introduction

“The eventual goal of science is to provide a single theory that describes the whole universe.”

– Stephen Hawking [25]

The quote coined by the famous physicist Stephen Hawking reflects that, the nature of science is to quest a *unified* approach to describe and understand various phenomena in the universe. So, *what about motion*?

Motion learning, such as *optical flow* and *scene depth*, in videos is one of the most fundamental problems in computer vision, and has a wide spectrum of key applications in many downstream tasks, *e.g.*, 2D video segmentation and tracking and 3D detection. Current works, somehow unexpected, are often focused on developing a specialized architecture for a particular aspect in motion learning, *e.g.*, either optical flow [15, 31, 68] or scene depth [2, 47], while neglecting the underlying connections between different motion learners and their unification for motion representation, which potentially could be transformative for motion learning. In addition, existing task-specific motion learning mechanism may lack the flexibility to generalize to broader downstream vision tasks, *e.g.*, being applicable to either 2D or 3D tasks,

not both. Therefore, the following question, as in the above quote but for motion learning, naturally arises: ① *Is there a manner to unify different motion learners?*

Besides, for motion learning tasks, the notorious photometric inconsistency (*e.g.*, caused by the shadows and illumination variations) widely exists, which introduces unexpected uncertainty during matching and hence compromise the underlying motion representation learning by inaccurate or even false pixel-wise feature matching, degrading performance. Thus, along with question ①, another question needs to be answered for unified motion learning: ② *How can we mitigate uncertainty in motion learning?* or in other words, *how to learn uncertainty reduced motion representation?*

In an effort to embrace these challenges, we propose **ProMotion**, a novel **Prototypical Motion** architecture based on Transformer for unified learning of the optical flow and depth from videos. With the learned knowledge, *ProMotion* can seamlessly adapt to downstream tasks with ease (Fig. 1). The key motivation behind PROMOTION is that, motion (*e.g.*, optical flow or scene depth) essentially is to describe movement of an object or a semantic region, which can be seen as a *prototypical learner* to a set of pixel exemplars. Viewed in this light, prototype, as a underlying connection between different motion tasks, is leveraged for unified motion learning, answering the question ①. Besides, by modeling motion using semantic prototypical features instead of pixel features, and optimizing them in the designed *feature denoiser*, we reduce the risk of noisy as well as outlier pixels for feature matching, which effectively alleviates the uncertainty issue, answering the question ②.

At the core of PROMOTION lies a novel prototypical Transformer architecture. Different from existing task-specific motion learning models [31, 47, 68] that directly extract feature embeddings from Transformer, PROMOTION gradually denoises the feature space to preserve the most dominant features and then optimizes the denoised features to subspaces in appearance, texture, and geometry along orthogonal direction, which enforces intra-class compactness. Afterwards, the optimized subspaces are grouped into a set of semantic prototypes with iterative clustering. In comparison with the feature embeddings from vanilla Transformer [78], the optimized subspace features from PROMOTION are denoised to some extent and thus are more suitable for prototype learning. Eventually, the prototypes are fed to universal adaptation head to produce motion information such as optical flow and depth. Fig. 2 illustrates the proposed PROMOTION for unified motion learning.

To the best of our knowledge, PROMOTION is the first unified motion learners with prototype learning. In comparison to current task-specific designs, PROMOTION enjoys a few attractive qualities: **① Unified motion learning.** PROMOTION, from the prototype perspective, provides a single paradigm (see Fig. 2 (a)) for different motion learners, which simplifies the architecture design of motion learning and allows unified learning of various motion information. **② Uncertainty-aware motion representation.** By learning motion from *feature denoiser* (see Fig. 2 (b)) and *prototypical learner* (see Fig. 2 (c)), PROMOTION avoids noisy and outlier pixel-wise feature matching, which mitigates the uncertainties in motion for robustness. **③ Generalization to more downstream tasks.** PROMOTION learns different motion information within one single framework, which could benefit more 2D and 3D downstream tasks.

To effectively assess our method, we present compelling experimental results on various datasets and different settings. We show experimentally in §4.1, with the task of optical flow estimation, *ProMotion* outperforms existing counterparts, *e.g.*, 0.16 and 0.29 lower Average End-Point Error (AEPE) compared to Flowformer [31] and MatchFlow [15] on the clean pass of Sintel. In §4.2, under the same paradigm, our depth estimation outperforms competitors by significant margins in the *Abs Rel* and *Sq Rel* metrics. For instance, we reduce the relative loss by 0.08 for *Abs Rel* and 0.27 for *Sq Rel* compared to recent DPT [64] on Sintel. Our algorithms are extensively tested, and the efficacy of the core components is demonstrated through ablative studies in §4.3. We believe this work can provide insights into this field.

2. Related Work

Motion Task. Motion tasks are characterized by the intricate processes of identifying, modeling, and predicting the motion patterns of objects and scenes, making it indispensable for a wide range of computer vision appli-

cations and subsequent tasks. Motion tasks play a pivotal role in detecting the motion of vehicles and pedestrians [11, 27, 38, 49, 86] for self-driving car navigation and perception, identifying abnormal activities in video surveillance, and recognizing actions [45, 76, 98]. In addition, motion tasks improve video compression and segmentation [21, 28, 56, 63] by predicting object and frame-level relationships. Among the myriad of motion-related tasks, two highly representative tasks stand out as particularly noteworthy: optical flow and scene depth estimation.

Optical Flow and Depth. Current methods for estimating optical flow and scene depth have created specialized architectures for each, ignoring their intrinsic relationships.

Optical flow, which involves finding 2D pixel displacements between images, is a fundamental problem in computer vision. FlowNet [33], as a pioneering work, first introduced the promise of a CNN-based architecture for direct flow regression. Building upon synthetic data for pre-training, it achieves coarse estimation performance. This sparked further explorations in architectures and training strategies, including iterative refinement and pyramid regression [65, 73, 75], incorporating photometric [50] and forward-backward consistency [37], and transformer architecture [31, 55, 68, 72]. These advances have outperformed earlier approaches by significantly reducing end-point errors.

Learning-based scene depth prediction typically takes a single image as input and utilizes generic network architectures such as ResNet [26] and ViT [17] to predict scene depth. DepthNet [19] pioneered the direct depth regression networks and established optimization through the Scale-Invariant Log Loss (SIlog). Significant progress has been driven by innovations in network architectures [42, 43, 60, 85], optimization schemes [9, 20, 42], and multi-view fusion [23, 66]. To improve the architectures, Laina et al. [42] developed a residual network, Xu et al. [85] presented attention-guided neural fields, and Lee et al. [43] used a multi-scale fusion scheme to improve contextual reasoning. For optimization, advanced loss functions, including reverse Huber loss [42] and ordinal regression [20] provide stronger supervision signals. Multi-view feature fusion incorporates more geometric cues from additional perspectives [66].

However, unlike the specialized architectures that focus on improvements in architecture, optimization, etc., our worldview cognition is more ambitious in unifying motion tasks under one paradigm, which is a rare exploration.

Unified Vision Models. As a recent trend in scientific research, unified theories have seen their emergence. There have been successful attempts to train different language modeling objectives within a single model [14, 39, 40, 54, 84]. However, the development of unified models for computer vision still lags far behind language models. In the vision regime, some of the early research efforts have focused on the development of either encoders [12, 16, 17, 52]

or decoders [8, 44, 48, 92]. The encoders focus on the effort to develop generic backbones that are trained on extensive data, so they can be adapted for different downstream tasks. For instance, Swin [52] and CSWin Transformer [16] serve as a general-purpose backbone for visual tasks. In contrast, studies on decoders [8, 44, 48, 92] are designed to address homogeneous specific tasks (e.g., object recognition [44, 100], instance segmentation [8, 44], semantic segmentation [8, 44]) by representing visual patterns from queries. Despite this, there has been little effort to create a unified model for motion tasks involving optical flow and scene depth estimation.

Prototype Learning. Prototype-based methods have gained substantial attention in the field of machine learning due to their alignment with human cognitive processes and intuitive appeal. Studies in psychology have demonstrated that humans tend to rely on prototypical examples as a basis for learning and problem-solving when confronted with novel challenges [61, 91]. In the realm of machine learning, prototype-based methodologies deviate from conventional techniques, such as support vector machines [24] and multilayer perceptrons [3]. They instead reason about observations by making comparisons with representative examples. Among the earliest prototype-based classifiers, the nearest neighbor algorithm [10] stands out and laid the foundation for subsequent methods like learning vector quantization (LVQ) [41], generalized LVQ [67] and deep nearest centroids [79]. Recent times have witnessed a surge of interest in deep learning techniques, combined with prototype learning, which have shown promising results across various domains within machine learning. These include few-shot learning [51, 71], zero-shot learning [35, 95], unsupervised learning [83, 88], and supervised learning [59, 82, 90]. Drawing from the valuable insights offered by these prior works, we embark on an endeavor to extend prototype learning into the motion tasks. In this context, we interpret the motion of objects or regions as an ensemble of prototypes. This innovative approach inherently encapsulates the dynamic features of motion, resulting in an advancement in the foundational learning of diverse motion attributes.

3. ProMotion

In this section, we present PROMOTION, a novel framework for unified motion tasks (see Fig. 2 (a)). The model has a serial of hierarchical transformer blocks in *feature denoiser* that enables compact feature learning, followed by *prototypical learner* to capture valuable motion concepts. We elaborate the pipeline in two stages — first *feature denoiser* (§3.1) and then *prototypical learner* (§3.2).

3.1. Feature Denoiser

Proposition 1. *Feature denoising enhances representation learning by capturing more informative subspaces.*

To suppress noises and mitigate uncertainty in motion representation learning, we propose a novel *feature denoiser* for enhancement. Specifically, considering a mapping function $\Phi(\cdot)$ and input patch tokens $X \in \mathbb{R}^{D \times N}$, the goal of feature denoiser is to learn a more compact and dominant feature representation $Z \in \mathbb{R}^{d \times N}$ via *subspace denoising* and *large-margin projection*:

$$\Phi : x \in \mathbb{R}^{D \times N} \rightarrow z \in \mathbb{R}^{d \times N}, \quad d < D \quad (1)$$

For subspace denoising, we abstract the raw representation x using the expected denoised tokens z with an additive Gaussian noise w as: $x = z + \sigma w$, where $w \sim N(0, 1)$. The objective is to minimize the distance between raw representation x and clean representation z , under noise level σ^l as:

$$E[z|\cdot] = \arg \min_f E_{z,w} [\|f(z + \sigma^l w) - z\|_2^2] \quad (2)$$

where $f(\cdot)$ is the transition function at the l^{th} transformer layer. If z^l is the feature representation under noise level σ^l , and the updated representation at the next stage is $z^{l'} = E[z|z^l]$. Using Tweedie’s formula [18], the optimization in Eq. 2 can be solved by computing posterior expectations as:

$$\begin{aligned} z^{l'} &= z^l + (\sigma^l)^2 \nabla_x \log p(z^l) \\ &\approx \sum_{k=1}^K (w_k * U_k * (U_k^T * z^l)) \end{aligned} \quad (3)$$

where w_k satisfies $\text{Softmax}(\frac{1}{2 \times (\sigma^l)^2} * \|U_k^T z^l\|_2^2)$. This reveals that the signal $z^{l'}$ can be naturally reconstructed as a weighted sum of multiple linear projections $U_k * (U_k^T * z^l)$, which enables reformulation of the transformer encoder to avoid non-linear attention blocks, much simpler yet effective.

To improve the robustness of feature embeddings to noise, we further apply a large-margin projection to enforce inter-class separability and intra-class clustering that maps features in orthogonal directions. For this purpose, the Iterative Shrinkage Thresholding Algorithm (ISTA) [1, 22, 80, 93] is used to infer the z^{l+1} as:

$$\begin{aligned} z^{l+1} &= \text{ISTA}(z^{l'} | \mathcal{O}^l) \\ &= \text{ReLU} \left(z^{l'} + \epsilon \mathcal{O}^{l*} (z^{l'} - \mathcal{O}^l z^{l'}) \right) \end{aligned} \quad (4)$$

where $\epsilon > 0$ is a step size and \mathcal{O}^l is the sparsifying orthogonal dictionary. With this approach, the projected subspace can be interpreted as ReLU activation function of multiple linear transformations. The design of subspace denoising and large-margin projection aligns with the trend of achieving representative representation through sparsifying representation, which enjoys the following appealing characteristics:

- *Underlying Data Structure:* Subspace denoising maps the input tokens to different feature subspaces and removes irrelevant variances, allowing the most representative data

structure to emerge clearly. The large-margin projection further projects feature subspaces in a low-dimensional orthogonal direction to increase inter-class separation. As a result, *Feature Denoiser* reflects the most representative or characteristic relationships.

- *Uncertainty Mitigation*: Aligning closely with the underlying data structure, subspace denoising improves the signal-to-noise ratio and allows for more repeatable matching correspondences, while large-margin projection eliminates redundant dimensions, allowing for more compact features for matching. *Feature Denoiser* mitigates uncertainty after subspace denoising and large-margin projection.

3.2. Prototypical Learner

Proposition 2. *Dynamic patterns can be captured by a set of prototypes, therefore unifying different motion tasks.*

After subspace denoising in the *feature denoiser*, we propose to employ prototypical learning on the denoised and compact representation for unified motion learning. The essence of prototypical learning is that, for every token embedding \mathcal{X} , find a mapping function f_c to generate K prototypes with prototype centers $C \in \mathbb{R}^{K \times D}$ that is initialized from a D -dimensional image grid. For each image x_i , the expected set of prototype segments are:

$$f_c(\mathcal{X}) = (\mathbf{r}_i)_{i=1}^N, \quad \mathbf{r}_i = (r_{i,k})_{k=1}^K \quad (5)$$

The initial prototypes are iteratively updated by Sinkhorn-Knopp based clustering as an approximate solution of optimal transport problem:

$$\max_{\mathbf{Q} \in \mathcal{Q}} \text{Tr}(\mathbf{Q}^\top \mathbf{V}), \quad \text{s.t. } \mathbf{Q} \mathbf{1}_N = \frac{1}{K} \mathbf{1}_K, \mathbf{Q}^\top \mathbf{1}_K = \frac{1}{N} \mathbf{1}_N \quad (6)$$

where $\mathbf{Q} \in \mathbb{R}^{K \times N}$ is the soft cluster assignment matrix and is constrained in the transportation polytope under the equipartition constraint [13], \mathbf{V} is the similarity matrix where $\mathbf{V} = \mathbf{C} \mathcal{X}^\top \in \mathbb{R}^{K \times N}$. Current prototypes are updated according to \mathbf{Q}_{t-1} for generating new prototypes centered at C_t :

$$C_t \leftarrow \mathbf{Q}_{t-1} \mathcal{X} \in \mathbb{R}^{K \times D} \quad (7)$$

By modeling the intrinsic similarities of prototypes via prototypical learning, the method refines the feature embeddings and improves the overall understanding of the underlying context and structure. All feature representations can be seamlessly used for various target tasks via adaptation heads (see §3.3). Prototypical learning has the following features:

- *Unified Paradigm*: Prototypical learning promotes the unified paradigm to encapsulate the underlying structure and context of data that can be leveraged across a variety of data and visual tasks. The learned prototypes capture the most representative and informative prototype-wise structures that are invariant across diverse data distributions and tasks, leading to great generalizability and adaptability.

- *Transferability*: The prototypes generated capture the underlying data structures and the essence of a category, which can be used directly by adding task-specific heads for downstream tasks. For complex data or tasks, the learned prototypes can also be reused by adding new prototypes or re-weighting the existing ones. This elegant design facilitates representation learning for knowledge transferability in accordance with the unified paradigm.

3.3. Implementation Details

The implementation details and overall framework of PRO-MOTION are shown in Fig. 2.

- *Encoder*: The goal of the encoding process is to generate token prototypes from the given image I . The whole pipeline starts with tokenization to convert the I into token embeddings. Subsequently, the tokens are passed to *feature denoiser* (subspace denoising (Eq. 3) and large-margin projection (Eq. 4)) to learn denoised and compact subspaces. Once the denoised subspaces are learned, a set of prototype features is inferred from the subspaces (Eq. 5), and are then updated via iterative clustering (Eq. 7).
- *Universal Adaptation*: For optical flow and scene depth, the flow and depth heads upsample the features of the token prototypes by predicting a convex mask to retrieve the flow map and depth. The initial predictions (\hat{f}_0, \hat{d}_0) are updated by continuously estimating the residuals $(\delta \hat{f}, \delta \hat{d})$. The last estimate is used as the final prediction as $\hat{f}_n, \hat{d}_n = \Delta \hat{f}, \Delta \hat{d} + \hat{f}_{n-1}, \hat{d}_{n-1}$. When transferring to downstream tasks, the pre-trained motion features are frozen, and only the task-specific adaptation head is fine-tuned, leveraging the core designs of the *feature denoiser* (§3.1) and *prototypical learner* (§3.2).
- *Loss Functions*: We use the SILog loss [20] and a weighted L_1 loss [75] for guiding the depth and flow learning.

4. Experiments

4.1. Experiments on Optical Flow

Training. We follow standard optical flow training steps [31, 75] to train our model on FlyingChair (C) for 80K iterations and FlyingThings3D (T) for another 80K iterations (denoted as ‘C+T’) with batch size of 16. Following this, we fine-tune the model on a larger combination set of FlyingThings3D (T), Sintel (S), KITTI (K) and HD1K (H) (‘C+T+S+K+H’). The method is implemented on 8 NVIDIA A100 GPUs with one-cycle learning strategy [70] and AdamW optimizer [53] with a maximum learning rate of $3e-4$.

Testing. This model is evaluated on Sintel and KITTI benchmarks. Generalization and ablation are on the validation split of Sintel and KITTI following [31, 75].

Metrics. Following the evaluation metrics from the Sintel and KITTI benchmarks, we use the average end-point error (AEPE) to measure the average pixel-wise flow error on

Training	Method	Sintel (val)		KITTI-15 (val)		Sintel (test)		KITTI-15 (test)
		Clean	Final	Fl-epe	Fl-all	Clean	Final	Fl-all
C+T	PWC-Net _[CVPR2018] [73]	2.55	3.93	10.35	33.7	-	-	-
	LiteFlowNet _[TPAMI2020] [32]	2.24	3.78	8.97	25.9	-	-	-
	RAFT _[ECCV2020] [75]	1.43	2.71	5.04	17.4	-	-	-
	Separable Flow _[ICCV2021] [94]	1.30	2.59	4.60	15.9	-	-	-
	GMA _[ICCV2021] [36]	1.30	2.74	4.69	17.1	-	-	-
	AGFlow _[AAAI2022] [58]	1.31	2.69	4.82	17.0	-	-	-
	KPA-Flow _[CVPR2022] [57]	1.28	2.68	4.46	15.9	-	-	-
	DIP _[CVPR2022] [97]	1.30	2.82	4.29	13.7	-	-	-
	GMFlowNet _[CVPR2022] [96]	1.14	2.71	4.24	15.4	-	-	-
	GMFlow _[CVPR2022] [87]	1.08	2.48	7.77	23.40	-	-	-
	CRAFT _[CVPR2022] [72]	1.27	2.79	4.88	17.5	-	-	-
	FlowFormer _[ECCV2022] [31]	<u>1.01</u>	<u>2.40</u>	<u>4.09</u> [†]	14.7 [†]	-	-	-
	SKFlow _[NIPS2022] [74]	1.22	2.46	4.27	15.5	-	-	-
MatchFlow _[CVPR2023] [15]	1.14	2.61	4.19 [†]	13.6 [†]	-	-	-	
Ours	0.91	2.32	3.95	14.1	-	-	-	
C+T+S+K(+H)	PWC-Net _[CVPR2018] [73]	-	-	-	-	4.39	5.04	9.60
	RAFT _[ECCV2020] [75]	(0.76)	(1.22)	(0.63)	(1.5)	1.61	2.86	5.10
	Separable Flow _[ICCV2021] [94]	(0.69)	(1.10)	(0.69)	(1.60)	1.50	2.67	4.64
	GMA _[ICCV2021] [36]	(0.62)	(1.06)	(0.57)	(1.2)	1.39	2.47	5.15
	AGFlow _[AAAI2022] [58]	(0.65)	(1.07)	(0.58)	(1.2)	1.43	2.47	4.89
	KPA-Flow _[CVPR2022] [57]	(0.60)	(1.02)	(0.52)	(1.1)	1.35	2.36	4.60
	DIP _[CVPR2022] [97]	-	-	-	-	1.44	2.83	4.21
	GMFlowNet _[CVPR2022] [96]	(0.59)	(0.91)	(0.64)	(1.51)	1.39	2.65	4.79
	GMFlow _[CVPR2022] [87]	-	-	-	-	1.74	2.90	9.32
	CRAFT _[CVPR2022] [72]	(0.60)	(1.06)	(0.58)	(1.34)	1.45	2.42	4.79
	FlowFormer _[ECCV2022] [31]	<u>(0.48)</u>	<u>(0.74)</u>	<u>(0.53)</u>	<u>(1.11)</u>	<u>1.20</u>	<u>2.12</u>	4.68 [†]
	SKFlow _[NIPS2022] [74]	(0.52)	(0.78)	<u>(0.51)</u>	(0.94)	1.28	2.23	4.84
	MatchFlow _[CVPR2023] [15]	(0.51)	(0.81)	(0.59)	(1.3)	1.33	2.64	4.72
Ours	(0.40)	0.66	0.48	0.98	1.04	2.01	4.30	

Table 1. **Quantitative comparison on Sintel.** ‘C+T’: Succeeding training on FlyingChairs (C) and FlyingThings3D (T), the models are evaluated for generalization on Sintel (S) and KITTI (K) validation sets. ‘C+T+S+K(+H)’: Training on a combination of C, T, S, K, and HD1K (+H) is evaluated. The first and second place results are **bolded** and underlined, respectively. [†] indicates the use of tile technique [34].

Sintel, and *Fl-epe* (*AEPE* on outlier pixels) and *Fl-all* (*AEPE* over all pixels) on KITTI.

Quantitative Results. Following [75], we first evaluate the generalization performance after training on C+T. As shown in Table 1, PROMOTION achieves a 36.4% and 20.2% reduction in *AEPE* on Sintel clean pass, surpassing both the strong baseline [75] and the recent SOTA method [15]. On KITTI, there is a 21.6% and 5.7% decrease of *Fl-epe* in comparison to [75] and [15]. These outcomes indicate the good generalization of our approach.

The method is then evaluated in a more general manner to report leaderboard results for the Sintel and KITTI test sets with training on C+T+S+K(+H). On the Sintel, PROMOTION obtains *AEPE* of 1.04 and 2.01 on Clean and Final passes respectively, which significantly lowers baseline model [75] error by 35.4% and 29.7%. Similarly, for the KITTI, PROMOTION reduces [75] by 15.7%, from 5.10 to 4.30. The results demonstrate overall superior performance compared to other recent published works.

Qualitative Examples. We also provide qualitative results on Sintel and KITTI in Fig. 3. Under severe illumination and shadow variations and object occlusion, the estimation from

PROMOTION exhibits more complete objects and robust estimation (second and third rows), which is attributed to the design of *feature denoiser* to mitigate uncertainties. Within highly similar textures and patterns (first and fourth rows), the estimation of PROMOTION shows better capability in capturing object-level shapes and boundaries, which can be explained by the design of *prototypical learner* to capture more object-level motion patterns.

Dataset	Method	Abs Rel ↓	Sq Rel ↓	RMSE ↓
Sintel	TransDepth _[ICCV2021] [89]	0.78	1.07	0.83
	AdaBins _[CVPR2021] [2]	0.73	0.74	0.57
	DPT-Hybrid _[ICCV2021] [64]	<u>0.62</u>	<u>0.65</u>	0.58
	P3Depth _[CVPR2022] [62]	0.65	0.67	<u>0.39</u>
	Ours	0.54	0.38	0.36
KITTI	TransDepth _[ICCV2021] [89]	0.064	0.252	2.755
	AdaBins _[CVPR2021] [2]	<u>0.058</u>	0.198	2.360
	DPT-Hybrid _[ICCV2021] [64]	0.059	<u>0.190</u>	<u>2.315</u>
	P3Depth _[CVPR2022] [62]	0.060	0.206	2.519
	Ours	0.054	0.182	2.298

Table 2. **Depth performance on Sintel clean and KITTI test sets.** The best results are in **bold** and the second-best results are underlined. Our method outperforms other recent methods on most of the common error metrics.




Figure 3. **Qualitative comparison of optical flow on Sintel and KITTI v1 set.** Notable areas are marked with red circles. Compared to [31] and [15], our approach shows better ability to reduce matching uncertainties due to similar patterns, illumination changes, shadows, etc.

4.2. Experiments on Scene Depth

Training. We train scene depth under the same architecture as for flow, except for the dimensions in the headers. We follow the standard Eigen split [19] for training and then fine-tune the model to the smaller MPI Sintel [4] consisting of 14 sequences with clear edges and varying levels of motions.

Testing. The model is tested on the Sintel depth benchmark (1,104 images) and KITTI depth split (697 images).

Metrics. Following standard depth estimation metrics, we use absolute mean relative error (*Abs Rel*), squared mean relative error (*Sq Rel*), and root mean square error (*RMSE*).

Quantitative Results. As shown in Table 2, although the leading algorithms approach a saturation point to some extent (*Abs Rel* and *Sq Rel* are close among approaches), PROMOTION showed improvements with reductions of 26.0% and 16.9% on Sintel and 6.9% and 10.0% on KITTI in *Abs Rel* compared to the strong baseline AdaBins [2] and recent SOTA [62] for monocular depth estimation. This suggests that our method is suitable for a range of depth scenarios, including synthetic Sintel and real-world KITTI.

Qualitative Examples. As depicted in Fig. 4, the compared methods AdaBins [2] and DPT-Hybrid [64] tend to generate inconsistent noisy depth due to the texture-less regions, shadows, and lighting changes (first, second, and third rows), while PROMOTION shows more consistent and smoother depth estimates without being affected by the uncertainties, which is in line with the goal of our design *feature denoiser*. Moreover, the proposed method is able to preserve more object characteristics (*e.g.*, shape, surface, and boundaries) for both foreground objects and the out-of-focus image background (first and fourth rows) compared to AdaBins and DPT-Hybrid. This may be explained by the design of *prototypical learner* for better object-level representation learning.

4.3. Ablative Study

We conduct a series of ablation studies on PROMOTION with evaluating on the validation set of Sintel, as in Table 3.

Key Components. We first assess the effect of each core component in PROMOTION. As seen in Table 3a, solely baseline architecture without any proposed designs performs much worse than the model with feature denoiser. With further prototypical learner, we can get lower errors on all datasets, from 0.51 on clean pass to 0.45 on flow. The full settings achieves the best for both flow and depth, which validates the benefit of PROMOTION.

Attention Types. We also try vanilla self-attention, multi-head self-attention [17], and spatial reduction attention [77] in Table 3b. Considering that vanilla self-attention is a simple case of multi-head attention with only a single head, and spatial reduction attention focuses only on local context and ignores global dependencies, multi-head self-attention achieves the lowest error among the compared three, but still worse than the subspace self-attention in PROMOTION.

Number of Prototypes. Table 3c reports the performance of our approach with respect to different numbers of prototypes. The number of prototypes 10 achieves a baseline performance. As the number gradually increases to 100, we observe a clear performance gain from 0.46 to 0.40 in error for flow and from 0.65 to 0.59 for depth. However, continuing to increase beyond 100 results in marginal changes.

Decoder Iteration Number. To gain insights into the adaptation head, we ablate the effect of iteration number T in Table 3d. We find that the performance gradually improves from 0.58 and 0.84 on clean pass for flow and depth to 0.40 and 0.59 when increasing T from 1 to 12, but remains almost unchanged after running more iterations. However, with more iteration steps, the training and inference time will be stretched. We therefore set $T = 12$ by default for a




Figure 4. **Qualitative comparison of scene depth on the Sintel and KITTI val set.** Notable areas are marked with red circles. Compared to [2] and [64], our approach produces more consistent and smooth depths with complete object shapes and clear boundaries. Sparse ground truths in KITTI are interpolated for better visualization.

Algorithm Component	Flow (val)		Depth (val)		Attention Types	Flow (val)		Depth (val)	
	Clean	Final	Clean	Final		Clean	Final	Clean	Final
BASELINE	0.54	0.77	0.72	0.71	Vanilla Self-Attention	0.45	0.69	0.63	0.61
+ Feature Denoiser	0.51	0.73	0.67	0.66	Multi-head Self Attention	0.42	0.67	0.61	0.58
+ Prototypical Learner	0.45	0.70	0.63	0.62	Spatial Reduction Attention	0.43	0.68	0.63	0.60
+ Unified Learning (Our ProMotion)	0.40	0.66	0.59	0.57	Subspace Self-Attention (Ours)	0.40	0.66	0.59	0.57

Variant Number of Prototype	Flow (val)		Depth (val)		Decoder Iteration Number	Flow (val)		Depth (val)	
	Clean	Final	Clean	Final		Clean	Final	Clean	Final
10	0.46	0.71	0.65	0.62	1	0.58	0.89	0.84	0.77
50	0.42	0.69	0.61	0.59	6	0.42	0.71	0.62	0.61
100	0.40	0.66	0.59	0.57	12	0.40	0.66	0.59	0.57
200	0.41	0.68	0.61	0.58	24	0.39	0.66	0.60	0.56

Variant Prototype Updating Strategy	Flow (val)		Depth (val)		Head Dimension	Flow (val)		Depth (val)	
	Clean	Final	Clean	Final		Clean	Final	Clean	Final
Cosine Similarity	0.43	0.70	0.62	0.61	1	0.47	0.76	0.68	0.65
K-Means	0.41	0.68	0.62	0.60	4	0.43	0.69	0.62	0.61
Gaussian Mixture Models	0.41	0.67	0.61	0.58	8	0.40	0.66	0.59	0.57
Sinkhorn-Knopp	0.40	0.66	0.59	0.57	16	0.40	0.65	0.59	0.55

Table 3. **A set of ablative studies** on Sintel [5] flow and depth datasets (see §4.3). The adopted designs are marked in red.

better trade-off between accuracy and cost.

Prototype Updating Strategy. We further probe the influence of prototype updating methods, by comparing it with vanilla cosine similarity, K-Means, and Gaussian Mixture Models. As shown in Table 3e, the adoption is effective, improving cosine similarity and K-Means by 3.3% and 0.9%.

Head Dimensions. Finally, we compare the different head dimensions for the attention heads in Table 3f. We find that error significantly reduces from 0.47 to 0.40 in *AEPE* for flow and from 0.68 to 0.59 in *Abs Rel* for depth when increasing the dimension from 1 to 8, but there is no obvious decrease as the dimension grows to 16. For a fair comparison with other methods [15, 31], we choose it as 8.

5. Conclusion

In this work, we centered on a prototypical motion paradigm and a feature denoising optimization, which advocates a unified motion framework named PROMOTION. This work aims to mitigate the uncertainties by modeling motion as prototypical features and optimizing them in feature space. Empirical results suggest that PROMOTION achieves superior performance on two most representative motion tasks: optical flow and scene depth. Our work may potentially benefit the broader domain of dense motion estimation.

Acknowledgement

This work is supported by the National Science Foundation under Award No. 2242243.

References

- [1] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009. 4
- [2] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *CVPR*, 2021. 1, 2, 6, 7, 8
- [3] Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*. Springer, 2006. 4
- [4] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, 2012. 7
- [5] Daniel J Butler, Jonas Wulff, Garrett B Stanley, and Michael J Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, 2012. 8
- [6] Hansheng Chen, Yuyao Huang, Wei Tian, Zhong Gao, and Lu Xiong. Monorun: Monocular 3d object detection by reconstruction and uncertainty propagation. In *CVPR*, 2021. 1
- [7] Yihong Chen, Yue Cao, Han Hu, and Liwei Wang. Memory enhanced global-local aggregation for video object detection. In *CVPR*, 2020. 1
- [8] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *CVPR*, 2022. 1, 4
- [9] Zhiyuan Cheng, James Liang, Hongjun Choi, Guan hong Tao, Zhiwen Cao, Dongfang Liu, and Xiangyu Zhang. Physical attack on monocular depth estimation with optimal adversarial patches. In *ECCV*, 2022. 3
- [10] Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE transactions on information theory*, 13(1):21–27, 1967. 4
- [11] Yiming Cui, Cheng Han, and Dongfang Liu. Collaborative multi-task learning for multi-object tracking and segmentation. *Journal on Autonomous Transportation Systems*, 2023. 3
- [12] Mostafa Dehghani, Josip Djolonga, Basil Mustafa, Piotr Padlewski, Jonathan Heek, Justin Gilmer, Andreas Peter Steiner, Mathilde Caron, Robert Geirhos, Ibrahim Alabdulmohsin, et al. Scaling vision transformers to 22 billion parameters. In *ICML*, 2023. 3
- [13] Cuturi M Sinkhorn Distances. Lightspeed computation of optimal transport. *NIPS*, 2013. 5
- [14] Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. Unified language model pre-training for natural language understanding and generation. *NIPS*, 2019. 3
- [15] Qiaole Dong, Chenjie Cao, and Yanwei Fu. Rethinking optical flow from geometric matching consistent perspective. In *CVPR*, 2023. 1, 2, 3, 6, 7, 8
- [16] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. In *CVPR*, 2022. 3, 4
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 3, 7
- [18] Bradley Efron. Tweedie’s formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614, 2011. 4
- [19] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *NIPS*, 2014. 3, 7
- [20] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *CVPR*, 2018. 3, 5
- [21] Han Gao, Jinzhong Cui, Mao Ye, Shuai Li, Yu Zhao, and Xiatian Zhu. Structure-preserving motion estimation for learned video compression. In *ACMMM*, 2022. 3
- [22] Karol Gregor and Yann LeCun. Learning fast approximations of sparse coding. In *ICML*, 2010. 4
- [23] Vitor Guizilini, Rareş Ambruş, Dian Chen, Sergey Zakharov, and Adrien Gaidon. Multi-frame self-supervised depth with transformers. In *CVPR*, 2022. 3
- [24] Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. *The elements of statistical learning: data mining, inference, and prediction*. Springer, 2009. 4
- [25] Stephen Hawking. A brief history of time (new york: Bantam). *Hawking suggests that because God made the universe” evolve in a very regular way according to certain laws... it therefore seems equally reasonable to suppose that there are also laws governing the initial state*, page 11, 1988. 2
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 3
- [27] Hou-Ning Hu, Qi-Zhi Cai, Dequan Wang, Ji Lin, Min Sun, Philipp Krahenbuhl, Trevor Darrell, and Fisher Yu. Joint monocular 3d vehicle detection and tracking. In *ICCV*, 2019. 3
- [28] Zhihao Hu, Guo Lu, and Dong Xu. Fvc: A new framework towards deep video compression in feature space. In *CVPR*, 2021. 3
- [29] De-An Huang, Zhiding Yu, and Anima Anandkumar. Minvis: A minimal video instance segmentation framework without video-based training. *NIPS*, 2022. 1
- [30] Kuan-Chih Huang, Tsung-Han Wu, Hung-Ting Su, and Winston H Hsu. Monodr: Monocular 3d object detection with depth-aware transformer. In *CVPR*, 2022. 1
- [31] Zhaoyang Huang, Xiaoyu Shi, Chao Zhang, Qiang Wang, Ka Chun Cheung, Hongwei Qin, Jifeng Dai, and Hongsheng Li. Flowformer: A transformer architecture for optical flow. In *ECCV*, 2022. 1, 2, 3, 5, 6, 7, 8
- [32] Tak-Wai Hui, Xiaou Tang, and Chen Change Loy. Lite-flownet: A lightweight convolutional neural network for optical flow estimation. In *CVPR*, 2018. 6
- [33] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, 2017. 3

- [34] Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, David Ding, Skanda Kop-pula, Daniel Zoran, Andrew Brock, Evan Shelhamer, et al. Perceiver io: A general architecture for structured inputs & outputs. *arXiv preprint arXiv:2107.14795*, 2021. 6
- [35] Saumya Jetley, Bernardino Romera-Paredes, Sadeep Jayasumana, and Philip Torr. Prototypical priors: From improving classification to zero-shot learning. *arXiv preprint arXiv:1512.01192*, 2015. 4
- [36] Shihao Jiang, Dylan Campbell, Yao Lu, Hongdong Li, and Richard Hartley. Learning to estimate hidden motions with global motion aggregation. In *ICCV*, 2021. 6
- [37] Rico Jonschkowski, Austin Stone, Jonathan T Barron, Ariel Gordon, Kurt Konolige, and Anelia Angelova. What matters in unsupervised optical flow. In *ECCV*, 2020. 3
- [38] Anouar Ben Khalifa, Ihsen Alouani, Mohamed Ali Mahjoub, and Najoua Essoukri Ben Amara. Pedestrian detection using a moving camera: A novel framework for foreground detection. *Cognitive Systems Research*, 60:77–96, 2020. 3
- [39] Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh Hajishirzi. Unifiedqa: Crossing format boundaries with a single qa system. *EMNLP - findings*, 2020. 3
- [40] Daniel Khashabi, Yeganeh Kordi, and Hannaneh Hajishirzi. Unifiedqa-v2: Stronger generalization via broader cross-format training. *arXiv preprint arXiv:2202.12359*, 2022. 3
- [41] Teuvo Kohonen. The self-organizing map. *Proceedings of the IEEE*, 78(9):1464–1480, 1990. 4
- [42] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *3DV*, 2016. 3
- [43] Jae-Han Lee, Minhyeok Heo, Kyung-Rae Kim, and Chang-Su Kim. Single-image depth estimation based on fourier domain analysis. In *CVPR*, 2018. 3
- [44] Feng Li, Hao Zhang, Huaizhe Xu, Shilong Liu, Lei Zhang, Lionel M Ni, and Heung-Yeung Shum. Mask dino: Towards a unified transformer-based framework for object detection and segmentation. In *CVPR*, 2023. 4
- [45] Jing Li, Qingwang Huang, Yingjun Du, Xiantong Zhen, Shengyong Chen, and Ling Shao. Variational abnormal behavior detection with motion consistency. *TIP*, 2021. 3
- [46] Peixuan Li, Huaici Zhao, Pengfei Liu, and Feidao Cao. Rtm3d: Real-time monocular 3d detection from object keypoints for autonomous driving. In *ECCV*, 2020. 1
- [47] Zhenyu Li, Zehui Chen, Xianming Liu, and Junjun Jiang. Depthformer: Exploiting long-range correlation and local information for accurate monocular depth estimation. *Machine Intelligence Research*, pages 1–18, 2023. 2, 3
- [48] Weicong Liang, Yuhui Yuan, Henghui Ding, Xiao Luo, Weihong Lin, Ding Jia, Zheng Zhang, Chao Zhang, and Han Hu. Expediting large-scale vision transformer for dense prediction without fine-tuning. *NIPS*, 2022. 4
- [49] Dongfang Liu, Yiming Cui, Yingjie Chen, Jiyong Zhang, and Bin Fan. Video object detection for autonomous driving: Motion-aid feature calibration. *Neurocomputing*, 409:1–11, 2020. 3
- [50] Pengpeng Liu, Michael Lyu, Irwin King, and Jia Xu. Selfflow: Self-supervised learning of optical flow. In *CVPR*, 2019. 3
- [51] Yuanwei Liu, Nian Liu, Xiwen Yao, and Junwei Han. Intermediate prototype mining transformer for few-shot semantic segmentation. *NIPS*, 2022. 4
- [52] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 3, 4
- [53] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 5
- [54] Nicholas Lourie, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Unicorn on rainbow: A universal commonsense reasoning model on a new multitask benchmark. In *AAAI*, 2021. 3
- [55] Yawen Lu, Qifan Wang, Siqi Ma, Tong Geng, Yingjie Victor Chen, Huaijin Chen, and Dongfang Liu. Transflow: Transformer as flow learner. In *CVPR*, 2023. 3
- [56] Yawen Lu, Jie Zhang, Su Sun, Qianyu Guo, Zhiwen Cao, Songlin Fei, Baijian Yang, and Yingjie Chen. Label-efficient video object segmentation with motion clues. *IEEE TCSVT*, 2023. 3
- [57] Ao Luo, Fan Yang, Xin Li, and Shuaicheng Liu. Learning optical flow with kernel patch attention. In *CVPR*, 2022. 6
- [58] Ao Luo, Fan Yang, Kunming Luo, Xin Li, Haoqiang Fan, and Shuaicheng Liu. Learning optical flow with adaptive graph reasoning. In *AAAI*, 2022. 6
- [59] Pascal Mettes, Elise Van der Pol, and Cees Snoek. Hyper-spherical prototype networks. *NIPS*, 2019. 4
- [60] S Mahdi H Miangoleh, Sebastian Dille, Long Mai, Sylvain Paris, and Yagiz Aksoy. Boosting monocular depth estimation models to high-resolution via content-adaptive multi-resolution merging. In *CVPR*, 2021. 3
- [61] Allen Newell, Herbert Alexander Simon, et al. *Human problem solving*. Prentice-hall Englewood Cliffs, NJ, 1972. 4
- [62] Vaishakh Patil, Christos Sakaridis, Alex Liniger, and Luc Van Gool. P3depth: Monocular depth estimation with a piecewise planarity prior. In *CVPR*, 2022. 1, 6, 7
- [63] Zheyun Qin, Xiankai Lu, Xiushan Nie, Dongfang Liu, Yilong Yin, and Wenguan Wang. Coarse-to-fine video instance segmentation with factorized conditional appearance flows. *IEEE/CAA Journal of Automatica Sinica*, 10(5):1192–1208, 2023. 3
- [64] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *ICCV*, 2021. 1, 3, 6, 7, 8
- [65] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *CVPR*, 2017. 3
- [66] Alexander Rich, Noah Stier, Pradeep Sen, and Tobias Höllerer. 3dvnnet: Multi-view depth prediction and volumetric refinement. In *3DV*, 2021. 3
- [67] Atsushi Sato and Keiji Yamada. Generalized learning vector quantization. *NIPS*, 1995. 4
- [68] Xiaoyu Shi, Zhaoyang Huang, Dasong Li, Manyuan Zhang, Ka Chun Cheung, Simon See, Hongwei Qin, Jifeng Dai, and Hongsheng Li. Flowformer++: Masked cost volume

- autoencoding for pretraining optical flow estimation. In *CVPR*, 2023. 2, 3
- [69] Yuheng Shi, Naiyan Wang, and Xiaojie Guo. Yolov: making still image object detectors great at video object detection. In *AAAI*, 2023. 1
- [70] Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial intelligence and machine learning for multi-domain operations applications*, 2019. 5
- [71] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *NIPS*, 2017. 4
- [72] Xiuchao Sui, Shaohua Li, Xue Geng, Yan Wu, Xinxing Xu, Yong Liu, Rick Goh, and Hongyuan Zhu. Craft: Cross-attentional flow transformer for robust optical flow. In *CVPR*, 2022. 3, 6
- [73] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *CVPR*, 2018. 3, 6
- [74] Shangkun Sun, Yuanqi Chen, Yu Zhu, Guodong Guo, and Ge Li. Skflow: Learning optical flow with super kernels. *arXiv preprint arXiv:2205.14623*, 2022. 6
- [75] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*, 2020. 3, 5, 6
- [76] Radu Tudor Ionescu, Sorina Smeureanu, Bogdan Alexe, and Marius Popescu. Unmasking the abnormal events in video. In *ICCV*, 2017. 3
- [77] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, 2021. 7
- [78] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pvt v2: Improved baselines with pyramid vision transformer. *Computational Visual Media*, 8(3):415–424, 2022. 3
- [79] Wenguan Wang, Cheng Han, Tianfei Zhou, and Dongfang Liu. Visual recognition with deep nearest centroids. In *ICLR*, 2023. 4
- [80] John Wright and Yi Ma. *High-dimensional data analysis with low-dimensional models: Principles, computation, and applications*. Cambridge University Press, 2022. 4
- [81] Junfeng Wu, Yi Jiang, Song Bai, Wenqing Zhang, and Xiang Bai. Seqformer: Sequential transformer for video instance segmentation. In *ECCV*, 2022. 1
- [82] Zhirong Wu, Alexei A Efros, and Stella X Yu. Improving generalization via scalable neighborhood component analysis. In *ECCV*, 2018. 4
- [83] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 4
- [84] Tianbao Xie, Chen Henry Wu, Peng Shi, Ruiqi Zhong, Torsten Scholak, Michihiro Yasunaga, Chien-Sheng Wu, Ming Zhong, Pengcheng Yin, Sida I Wang, et al. Unifiedskg: Unifying and multi-tasking structured knowledge grounding with text-to-text language models. *arXiv preprint arXiv:2201.05966*, 2022. 3
- [85] Dan Xu, Wei Wang, Hao Tang, Hong Liu, Nicu Sebe, and Elisa Ricci. Structured attention guided convolutional neural fields for monocular depth estimation. In *CVPR*, 2018. 3
- [86] He Xu, Mingtao Guo, Nadia Nedjah, Jindan Zhang, and Peng Li. Vehicle and pedestrian detection algorithm based on lightweight yolov3-promote and semi-precision acceleration. *IEEE Transactions on Intelligent Transportation Systems*, 23(10):19760–19771, 2022. 3
- [87] Haoifei Xu, Jing Zhang, Jianfei Cai, Hamid Rezatofighi, and Dacheng Tao. Gmflow: Learning optical flow via global matching. In *CVPR*, 2022. 1, 6
- [88] Wenjia Xu, Yongqin Xian, Jiuniu Wang, Bernt Schiele, and Zeynep Akata. Attribute prototype network for zero-shot learning. *NIPS*, 2020. 4
- [89] Guanglei Yang, Hao Tang, Mingli Ding, Nicu Sebe, and Elisa Ricci. Transformer-based attention networks for continuous pixel-wise prediction. In *ICCV*, 2021. 6
- [90] Hong-Ming Yang, Xu-Yao Zhang, Fei Yin, and Cheng-Lin Liu. Robust classification with convolutional prototype learning. In *CVPR*, 2018. 4
- [91] Yi Yang, Yueting Zhuang, and Yunhe Pan. Multiple knowledge representation for big data artificial intelligence: framework, applications, and case studies. *Frontiers of Information Technology & Electronic Engineering*, 22(12):1551–1558, 2021. 4
- [92] Qihang Yu, Huiyu Wang, Dahun Kim, Siyuan Qiao, Maxwell Collins, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Cmt-deeplab: Clustering mask transformers for panoptic segmentation. In *CVPR*, 2022. 4
- [93] Yaodong Yu, Sam Buchanan, Druv Pai, Tianzhe Chu, Ziyang Wu, Shengbang Tong, Benjamin D Haefele, and Yi Ma. White-box transformers via sparse rate reduction. In *NIPS*, 2023. 4
- [94] Feihu Zhang, Oliver J Woodford, Victor Adrian Prisacariu, and Philip HS Torr. Separable flow: Learning motion cost volumes for optical flow estimation. In *ICCV*, 2021. 6
- [95] Hui Zhang and Henghui Ding. Prototypical matching and open set rejection for zero-shot semantic segmentation. In *ICCV*, 2021. 4
- [96] Shiyu Zhao, Long Zhao, Zhixing Zhang, Enyu Zhou, and Dimitris Metaxas. Global matching with overlapping attention for optical flow estimation. In *CVPR*, 2022. 6
- [97] Zihua Zheng, Ni Nie, Zhi Ling, Pengfei Xiong, Jiangyu Liu, Hao Wang, and Jiankun Li. Dip: Deep inverse patchmatch for high-resolution optical flow. In *CVPR*, 2022. 6
- [98] Joey Tianyi Zhou, Jiawei Du, Hongyuan Zhu, Xi Peng, Yong Liu, and Rick Siow Mong Goh. Anomalydet: An anomaly detection network for video surveillance. *IEEE Transactions on Information Forensics and Security*, 14(10):2537–2550, 2019. 3
- [99] Qianyu Zhou, Xiangtai Li, Lu He, Yibo Yang, Guangliang Cheng, Yunhai Tong, Lizhuang Ma, and Dacheng Tao. Transvod: end-to-end video object detection with spatial-temporal transformers. *TPAMI*, 2022. 1
- [100] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Simple multi-dataset detection. In *CVPR*, 2022. 4