This WACV paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# CCMR: High Resolution Optical Flow Estimation via <u>Coarse-to-Fine Context-Guided Motion Reasoning</u>

Azin Jahedi<sup>1</sup> Maximilian Luz<sup>2</sup> Marc Rivinius<sup>3</sup> Andrés Bruhn<sup>1</sup> <sup>1</sup>VIS, University of Stuttgart <sup>2</sup>RLL, University of Freiburg <sup>3</sup>SEC, University of Stuttgart {azin.jahedi,andres.bruhn}@vis.uni-stuttgart.de

luz@cs.uni-freiburg.de, marc.rivinius@sec.uni-stuttgart.de

# Abstract

Attention-based motion aggregation concepts have recently shown their usefulness in optical flow estimation, in particular when it comes to handling occluded regions. However, due to their complexity, such concepts have been mainly restricted to coarse-resolution single-scale approaches that fail to provide the detailed outcome of highresolution multi-scale networks. In this paper, we hence propose CCMR: a high-resolution coarse-to-fine approach that leverages attention-based motion grouping concepts to multi-scale optical flow estimation. CCMR relies on a hierarchical two-step attention-based context-motion grouping strategy that first computes global multi-scale context features and then uses them to guide the actual motion grouping. As we iterate both steps over all coarse-to-fine scales, we adapt cross covariance image transformers to allow for an efficient realization while maintaining scale-dependent properties. Experiments and ablations demonstrate that our efforts of combining multi-scale and attention-based concepts pay off. By providing highly detailed flow fields with strong improvements in both occluded and non-occluded regions, our CCMR approach not only outperforms both the corresponding single-scale attention-based and multiscale attention-free baselines by up to 23.0% and 21.6%, respectively, it also achieves state-of-the-art results, ranking first on KITTI 2015 and second on MPI Sintel Clean and Final. Code and trained models are available at https://github.com/cv-stuttgart/CCMR.

# 1. Introduction

Estimating the optical flow in terms of the 2D displacement field between two consecutive frames of an image sequence is a fundamental task in computer vision. It has found wide-spread application, ranging from autonomous driving and robotics over tracking and surveillance to frame interpolation and video compression.

Despite decades of research, the accurate estimation of



Figure 1. Comparison of our method to the ground truth (first row), recent approaches from the literature (second row) and the baselines (third row). Our method offers more structural details in the foreground and background (see boxes).

optical flow still remains challenging. In this context, most contemporary approaches rely on convolutional neural networks (CNNs) [7, 13, 33, 38, 40]. However, they lack global context, making reasoning about occlusions or large displacements difficult. Being capable of modeling dynamic long-range interactions, transformers and other attentionbased methods go far beyond the fixed receptive field of classical CNNs. Hence, it is not surprising that by providing this otherwise missing global information, they showed great potential in various computer vision tasks and advanced the field of optical flow estimation with respect to the aforementioned shortcomings [10, 18, 35, 44, 45, 47, 50]. Particularly interesting in this context is GMA [18], which has shown improvements in occluded areas by attentionbased aggregation of motion features and features of the reference frame. Recently, this approach has also been used by several other methods [6, 19, 36] as an improved backbone.

A second somewhat orthogonal successful concept, that seeks to embed the estimation in a more global context, is the computation of optical flow in a multi-scale coarse-tofine setting. It dates back to classical methods [2, 4] and early CNN-based approaches [12, 33, 38] where it primarily helped to improve the estimation of large displacements. In current methods, however, multi-scale approaches are hardly employed. Instead, they mainly follow the architectural shift introduced by RAFT [40] by estimating the optical flow on a single coarse scale and then upsampling the final result [10, 18, 25, 35, 39, 47]. Recently, however, MS-RAFT(+) [15,16] has shown that such schemes are still relevant and can, compared to single-scale approaches, not only improve robustness regarding ambiguous matches, but also produce more detailed flow fields. Thereby, valuable initializations from coarser scales allow them to operate more effectively on higher resolutions.

Evidently, it seems desirable to combine the strength of both attention-based motion guidance and high-resolution multi-scale schemes by integrating both concepts in a single framework. However, naively applying the (token) attention mechanism to the high resolution finer scales is prohibitive due to its quadratic complexity in the number of tokens (w.r.t. time and memory), thus requiring alternatives. Moreover, the desire to obtain a global context makes any of the local-attention based approaches [32, 41] proposed specifically to address the complexity problem less suitable. Methods attempting to retain global reasoning via hierarchical structures (e.g., [8, 21, 24, 43]; generally local at a scale) do inherently not retain their input resolution. This makes them less desirable, as our approach requires scale-specific global representations in all the coarse-to-fine scales.

An attractive option in this context seem to be crosscovariance image transformers (XCiT) [1], which propose cross-covariance attention (XCA) as an alternative to token attention [42] by essentially transposing feature-channels and tokens. This leads to an attention matrix quadratic over the number of input channels instead of tokens, making its complexity linear regarding the input size. Even though XCA loses the direct spatial structure inherent in standard token-attention, it retains implicit global interactions through the attention matrix, which can be seen as an expression of global channel statistics. As a result, XCiT performs favorably against various transformer-based models, even outperforming some [24, 43, 48].

**Contributions.** We propose CCMR, an optical flow approach that combines the benefits of attention-based motion aggregation concepts and high-resolution multi-scale approaches. In this context, our contributions are fourfold: (i) First, we consistently integrate context-based motion grouping concepts into a high-resolution coarse-to-fine estimation framework. This allows us to obtained detailed flow fields that also offer a high accuracy in occluded regions (see Fig. 1). (ii) Secondly, in this context, we propose a two-step motion grouping strategy, where we first compute global

self-attentive context features which are then used to guide the motion features iteratively over all scales. Thereby, an XCiT-based context-guided motion reasoning makes the processing over all coarse-to-fine scales feasible, while refraining from explicit patchification enables us to maintain scale-specific content (in contrast to [1]). (iii) Moreover, revisiting the U-Net style feature extractor of our multi-scale backbone, we increase its expressiveness while, independently, we reduce the size of the entire optical flow model by 28% using a leaner context encoder. Although these architectural changes are small, they strongly contribute to the improvements in non-occluded regions. (iv) Finally, experiments and ablations demonstrate the strong performance of our approach and the benefits of the underlying concepts. It achieves state-of-the-art results ranking first on KITTI 2015 and second on MPI-Sintel, Clean and Final.

# 2. Related Work

**Multi-Scale Optical Flow Estimation.** Most early neuralnetwork based optical flow methods [9, 12, 33, 38] owe their success to multi-scale estimation using coarse-to-fine warping schemes. These schemes estimate residual motion at each scale by iteratively proceeding from coarse to fine levels, compensating for previously estimated motion by warping the second frame towards the first, via the current flow estimate. First explored in classical approaches [2, 4], and therein found to be beneficial for estimating large motion and avoiding local minima, they have proven to be similarly successful for neural networks [11, 38]. Subsequent works focused on improving iterative refinement via weight sharing [12] and replacement of warping with sampling [9], avoiding ambiguities through ghosting effects.

Introducing a paradigm shift, RAFT [40] suggested the single-scale estimation of flow via a recurrent refinement strategy, now found in many contemporary and state-of-the-art approaches [6, 10, 16, 18, 25, 45, 47]. In each step, matching-costs are sampled based on the current flow estimate from a cost volume and fed through a recurrent unit to compute a residual flow update. While this poses a break from the until then predominant coarse-to-fine methods towards single-scale estimation, newer approaches, including DIP [51] and GMFlow(+) [44, 45] found it again beneficial to include a refinement stage on a secondary, finer scale, somewhat resembling earlier coarse-to-fine models. Recently, AnyFlow [19] proposes to aggregate the recurrent update with finer scale image features estimating a more detailed flow, however not in a coarse-to-fine manner.

Recently, MS-RAFT(+) [15,16], the architectural base of our method, reshaped RAFT into a coarse-to-fine estimation scheme using weight-sharing across scales, while relying on both multi-scale features and explicit multi-scale supervision. Thereby, the use of a coarse-to-fine scheme allows the estimation of flow at a higher final resolution, improving accuracy and preserving finer details, while also retaining the capability for estimating large motion via coarser scales.

In addition to the benefits inherited from MS-RAFT(+), our approach utilizes motion grouping concepts based on global self-attentive context and more expressive encoders, explicitly addressing occlusion handling and matching quality, respectively. In contrast to GMFlow(+) [44, 45] and DIP [51] that rely on additional refinement steps on a second scale, we consider a consistent coarse-to-fine method using four scales, operating at a higher final resolution  $(\frac{1}{2} \times (h, w)$  vs.  $\frac{1}{4} \times (h, w)$ ).

**Occlusions and Ambiguities.** Several approaches aim to explicitly predict occlusion masks, either in a supervised or self-supervised fashion [12,14,17,29] or indirectly in an unsupervised way [49]. We, however, focus on enabling networks to deal with occlusions and ambiguities intrinsically, intending to create a more holistic motion reasoning process without the need for any special training procedures.

In this context, both GMA [18] and KPA-Flow [25] aggregate motion-features based on the matching costs and current flow estimate with context features from the reference image via a token cross-attention mechanism. While GMA aims to resolve ambiguities and occlusions by incorporating global context-motion information, KPA-Flow takes inspiration from the smoothness-terms of classical energy-based optical flow approaches and aggregates motions only locally. Recently, GMA motion aggregation modules were also directly adopted by other methods, such as AnyFlow [19], MatchFlow [6] and CRAFT [36], as an improved backbone compared to RAFT.

GMFlow(+) [44, 45], on the other hand, follows the motion aggregation ideas used by GMA [18] and KPA-Flow [25] by performing an attention-based refinement step using self-similarity of image-features to propagate flow (instead of motion-features) into occluded and ambiguous areas. SeparableFlow [47], instead, aims to resolve difficult matches by applying an aggregation strategy to the cost volume itself. Finally, SKFlow [39] seeks to improve over GMA and RAFT by enlarging the receptive field of the convolutional layers in the motion encoder and update blocks specifically via their introduced super kernels.

In this work, we leverage motion grouping concepts to handle occlusions in a high-resolution coarse-to-fine framework. In contrast to GMA, our method not only operates on multiple scales, but it also relies on a two-step motion reasoning strategy that uses global self-attentive context features to guide the motion features. Thereby, the use of our motion reasoning concepts across all coarse-to-fine scales up to the finest scale  $(\frac{h}{2}, \frac{w}{2})$  is enabled by exploiting XCiT which has linear complexity over the number of pixels. Moreover, unlike other approaches which use XCiT in other vision tasks [26, 30, 31, 46], we refrain from the explicit patchification of its inputs originally proposed by



Figure 2. Coarse-to-fine architecture. The flow is first estimated on the coarsest scale, upsampled and used as initialization on the next finer scale. Note that the third scale is omitted for compactness.

XCiT and apply it directly to features at different scales maintaining their scale-specific properties. Besides, we further extend XCiT application beyond its original work [1], employing it not only for self- but also for cross-attention.

## 3. Approach

#### 3.1. Multi-Scale Flow Estimation

Our CCMR approach estimates the optical flow in a coarse-to-fine manner via recurrent updates, using a shared Gated Recurrent Unit (GRU), based on MS-RAFT+ [15]. Before starting with the coarse-to-fine estimation, for each scale s, image features  $F_{s,1}, F_{s,2}$  of both frames  $I_1, I_2$  are computed for matching. Moreover, context features  $C_s$  and, based on that, global context features  $GC_s$ , and the current scale's initial hidden state  $H_s$  of the recurrent unit are computed from the reference frame  $I_1$ . Fig. 2 illustrates the coarse-to-fine estimation process. Starting from the coarsest scale at  $\frac{1}{16} \times (h, w)$ , where h and w denotes the height and width of  $I_1$ , the flow is computed based on the abovementioned features  $F_{1,1}, F_{1,2}, C_1, GC_1, H_1$ , as well as an initial flow which is zero. After  $T_1$  recurrent flow updates, the estimated flow is upsampled using a shared  $\times 2$  convex upsampler, where the flow serves as initialization for the matching process on the next finer scale. This procedure continues until the flow is computed on the finest scale at  $\frac{1}{2} \times (h, w)$  and upsampled to the original resolution. For supervision, we use the same multi-scale multi-iteration loss as in [16]. To this end, the upsampled flow in each scale and each iteration is bilinearly upsampled to the original resolution for computing its difference to the ground truth.

In the following we detail on the key components of our CCMR architecture: the improved multi-scale feature consolidation unit to extract  $F_{s,1}, F_{s,2}, C_s$  and  $H_s$ , the computation of the global self-attentive context  $GC_s$ , the use of the global context for attention-based motion reasoning,



Figure 3. Improved multi-scale feature consolidation (top) and global context computation (bottom). Outputs shown on the right are used as inputs in the matching module in Fig. 2.

and the efficient realization of both components via XCiTbased self- and cross-attention, enabling the integration of such concepts in our high-resolution coarse-to-fine scheme.

#### 3.2. Improved Mutli-Scale Feature Consolidation

MS-RAFT(+) [15, 16] proposed to extract multi-scale image and context features using a U-Net-style feature extractor, inspired by feature pyramid networks [22]. To this end, intermediate features are computed in a top-down manner and then, to obtain multi-scale features, the betterstructured finer features  $\hat{F}_{s,1}$ ,  $\hat{F}_{s,2}$  and  $\hat{C}_s$  are semantically boosted by consolidating them with the deeper coarserscale features  $F_{s-1,1}$ ,  $F_{s-1,2}$  and  $C_{s-1}$  for  $s \in \{2,3,4\}$ . Thereby, the consolidation is realized by stacking the upsampled coarser features and the intermediate finer features and aggregating them via a residual unit.

This choice of architecture, however, is problematic for the computation of both image features and context features, due to the final ReLU activation applied in the residual unit. In the case of image features, this activation limits the range and therefore the expressiveness of the features of both images that are used in the computation of matching costs. In the case of context features, using this consolidation module is also not advisable, since the initial hidden state is computed as a Tanh activated split of the multi-scale context features after the consolidation, which is already activated via ReLU and does not contain negative values that Tanh expects for computing the hidden-state of the GRU.

We therefore improved the multi-scale feature consolidation module by adding an activation-free convolution layer after the residual unit at each scale, a small but effective modification to regain expressiveness. At the same time, using a leaner residual unit for the consolidation in the context encoder allowed us to reduce the size of entire optical flow network by 28% without harming the performance (see suppl. material). Our improved multi-scale feature consolidation is illustrated in the upper part of Fig. 3. Further details can be found in the suppl. material.



Figure 4. Visualization of context features. Top to bottom: reference frame, visualization of context  $C_4$ ,  $GC_4$ , two channels of  $GC_4$ . Similar to [1], the  $L_2$  norm of each feature map is shown as heat map. Each feature map is normalized individually.

#### 3.3. Global Context Computation

So far, we discussed how the multi-scale image and context features as well as the initial hidden-state per-scale are computed. Considering the context features  $C_s$  shown in Fig. 3, we now detail on how we compute our global context features. In this case, the goal is to obtain more attentive features that are later on used for guiding the motion. To this end, we perform an aggregation of the context features  $C_s$  via its channel statistics using a self-attention XCiT layer [1] (see lower part of Fig. 3), which offers linear complexity w.r.t the number of tokens. This architectural choice allows a feasible aggregation of context over all the coarseto-fine scales during the estimation. Importantly, our adaption of XCiT differs from its original approach, where the XCiT layer is actually applied on a coarser representation of its input realized via explicit input patchification and then again upsampled to obtain the original resolution. We, in contrast, apply the XCiT layer directly to the features in all coarse-to-fine scales, benefiting from scale-specific content.

To compute the global context, first, a sinusoidal positional embedding [42] is added to the context features  $C_s$ , yielding  $C_{s_p} \in \mathbb{R}^{N \times d}$ , where N and d denote the number of pixels at scale s and number of feature channels, respectively. This step is then followed by a layer normalization  $\mathcal{N}(C_{s_p})$  [3]<sup>1</sup>. At this point, to perform self-attention, all key (K), query (Q) and value features (V) features are computed from  $\mathcal{N}(C_{s_p})$  via

$$K_{C_s} = \mathcal{N}(C_{s_p}) \cdot W_{s,k},$$

$$Q_{C_s} = \mathcal{N}(C_{s_p}) \cdot W_{s,q},$$

$$V_{C_s} = \mathcal{N}(C_{s_p}) \cdot W_{s,v},$$
(1)

where  $W_{s,k}, W_{s,q}, W_{s,v} \in \mathbb{R}^{d \times d}$  are learned weight matrices. Before applying a cross-covariance attention step,

<sup>&</sup>lt;sup>1</sup>For a more compact illustration we omitted both steps in Fig. 3.



Figure 5. Matching block of our CCMR approach for each scale. Motion features are guided based on global context iteratively and used in the flow update computation. The motion encoder and the update block are shared among scales.

the channels of  $K_{C_s}, Q_{C_s}, V_{C_s} \in \mathbb{R}^{N \times d}$  are reshaped into h heads resulting in  $K_{C_s}, Q_{C_s}, V_{C_s} \in \mathbb{R}^{N \times d/h \times h}$ ; cf. [1]. The cross-covariance attention then reads

$$\operatorname{XCA}(K_{C_s}, Q_{C_s}, V_{C_s}) = V_{C_s} \cdot \mathcal{S}(\hat{K}_{C_s}^\top \cdot \hat{Q}_{C_s} / \tau_s), \quad (2)$$

where S denotes the Softmax function,  $\tau_s$  is the learned temperature parameter and  $\hat{K}_{C_s}$  and  $\hat{Q}_{C_s}$  denote  $L_2$ normalized key and query features, respectively. Note that the transpose operation and matrix multiplication in Eq. (2) are along the first two dimensions. Afterwards, a local patch interaction (LPI) layer and then a feed forward network (FFN) is applied. While the cross-covariance attention enables global interactions among channels in each head, LPI and FFN units provide explicit spatial interaction between tokens locally and communication across all channels [1], respectively. The XCiT operations can be summarized as

$$C_{\text{int}_1} = C_{s_p} + \gamma_{1_s} \cdot \text{XCA}(K_{C_s}, Q_{C_s}, V_{C_s}),$$
 (3)

$$C_{\text{int}_2} = C_{\text{int}_1} + \gamma_2 \cdot \text{LPI}(\mathcal{N}(C_{\text{int}_2})), \tag{4}$$

$$GC_s = C_{\text{int}_2} + \gamma_{3_s} \cdot \text{FFN}(\mathcal{N}(C_{\text{int}_2})), \tag{5}$$

where  $C_{\text{int}_1}$  and  $C_{\text{int}_2}$  are intermediate outcomes,  $GC_s$  is the resulting global self-attentive context features, and  $\gamma_{k_s}$ are learned scaling factors. The LPI module consists of two depth-wise convolutions and the FFN contains two linear layers. For simplicity and compactness, the weighted sum operations of Eqs. (3) to (5) are not shown in Fig. 3.

Before we discuss how the global self-attentive context features are used to guide motion features, let us first take a look at the impact of the global aggregation process. To this end, we visualize exemplary context maps in Fig. 4. Row 2 and 3 show the channel-wise  $L_2$  norm of the unaggregated context features  $C_4^2$  and of our global context features  $GC_4$ . As one can see, in the case of  $C_4$  the initial context features show high activations only mainly at



Figure 6. Comparison of memory consumption for our approach and the direct extension of GMA to multiple scales. X-axis shows height or width of squared input images.

edges, while in  $GC_4$  the attention of the global context is spread over entire objects. Additionally, Row 4 and 5 depict the magnitude of two channels (70 and 96) of  $GC_4$ . While the former attends to the foreground, the latter has large activations in the background.

#### 3.4. Global-Context Based Motion Grouping

In this section, we finally elaborate on the flow estimation via global-context based motion grouping based on the image features  $F_{s,1}, F_{s,2}$ , the initial and global context features  $C_s$  and  $GC_s$  and the initial hidden state  $H_s$  which where explained in Sec. 3.2 and Sec. 3.3. First, based on the initial flow in the first iteration (or the updated flow in the next iterations), matching costs in a neighborhood are computed from image features  $(F_{s,1}, F_{s,2})$ . Note that unlike RAFT [40], the all-pairs cost volume is not pre-computed, due to the large memory consumption for finer scales. The computed costs together with the current flow estimate are then processed via a motion encoder which outputs motion features (see Fig. 5) that are eventually used in the GRU to compute the flow update.

When computing the recurrent flow updates, additionally including global aggregated motion features based on context features can help resolve ambiguities in occluded areas. This is intuitive, because the motion of occluded pixels from a partially non-occluded object can typically be inferred from the motion of its non-occluded pixels. In contrast to GMA [18], which realizes this additional step with the means of a single token cross-attention [42] using coarse non-aggregated context to aggregate motion features in a single-scale setting, we follow our efficient strategy based on global channel statistics from the computation of the global context, which is performed on all coarse-tofine scales. More precisely, we perform the motion grouping by applying a cross-attention XCiT layer, on the global self-attentive context  $GC_s$  and motion features MF. Thus, we compute key and query features from our global context features  $GC_s$  and value features from the motion features

<sup>&</sup>lt;sup>2</sup>For fairness, an independent model without global context computation for motion reasoning was trained to generate the visualized  $C_4$ -maps.



Figure 7. Visual comparison of the results of our CCMR+ approach to two recent SOTA methods for each benchmark and to the two baselines MS-RAFT+ and GMA, when the original publication results available, From top to bottom: Temple-1 (Sintel-final), PERTURBEDmarket-3 (Sintel-clean), Sequences #9 and #15 (KITTI). Details best viewed electronically (by zooming in the flow fields).

directly from each scale without explicit patchification. After applying XCA, LPI and FFN to the key, query and value features, the context-guided motion features (CMFin Fig. 5), the context features  $C_s$  and the initial motion features MF are stacked and passed through the recurrent unit to compute the flow update iteratively.

Note that applying token cross-attention [42] to perform motion aggregation in high-resolution coarse-to-fine schemes is memory-wise infeasible. Fig. 6 shows that a simple GMA-based coarse-to-fine variant with the finest scale at  $(\frac{h}{2}, \frac{w}{2})$  requires an enormous amount of VRAM. Even for comparably small images (400 × 400 pixels) it needs more VRAM for inference than CCMR requires for training. Note that inference for Sintel images with  $436 \times 1024 \approx 668^2$  pixels in this case requires around 80 GB, while our approach needs less than 5 GB. In the suppl. material, we provide additional details on this comparison.

## 4. Experiments

# 4.1. Training Setup

We implemented our CCMR method in PyTorch. Following MS-RAFT+, during training we used 4, 5, 5 and 6 GRU iterations from the coarsest to the finest scale, respectively. We trained the model for 120K iterations on Chairs [7], 120K iterations on Things [27], and 120K iterations on a mixed set containing Sintel [5], KITTI 2015 [28], HD1K [20], and Things for Sintel evaluation. We then finetuned our model for 50K iterations on an automotive mixed set consisting of KITTI and Viper [34] for evaluating on *KITTI*. We used a batch size of 9 in the first training phase, 6 for the remaining three phases. For supervision, we used the same multi-scale multi-iteration loss as [16] with increasing weights over scales and iterations. In this context, we used the  $L_2$  loss for the first three training phases and the robust sample-wise loss [16] for the fine-tuning in the last phase. Note that, besides the 4-scale model with finest matching scale at:  $(\frac{h}{2}, \frac{w}{2})$  based on to MS-RAFT+ [15], we also investigated a 3-scale version based on MS-RAFT [16] with finest scale at  $(\frac{h}{4}, \frac{w}{4})$ . Consequently, in this section we rename the 4-scale CCMR method explained so far to CCMR+, while referring to the smaller 3-scale version as CCMR. Our CCMR and CCMR+ model with 10.8M and 11.5M parameters are 20% and 28% smaller than their MS-RAFT and MS-RAFT+ backbones with 13.5M and 16M parameters, respectively. This reduction of model size is due to our smaller modified context consolidation unit, which yielded better generalization results on average. For detailed inference settings, memory requirements and runtime as well as for an ablation on the larger consolidation unit; see suppl. material.

#### 4.2. Results

Let us now compare our CCMR(+) method to the state of the art and the underlying baselines (architectural baseline MS-RAFT(+) and conceptual baseline GMA), both in terms of accuracy and visual appearance. To this end, we compare the optical flow accuracy on the *Sintel* [5] and *KITTI* [28] benchmarks<sup>3</sup> and additionally show pre-training generalization results. Finally, we also delineate the architectural impact from the impact of our fine-tuning for KITTI to outline the advantage of our architecture compared to the baselines.

**Comparison to the Baselines.** Our results in Tab. 1 show large gains in both matched and unmatched regions, up to 46.6% and 22.2%, respectively. With large improvements against GMA in matched regions and against MS-RAFT(+) in unmatched regions, it becomes clear that our coarse-to-fine context-guided motion grouping is worthwhile. Specifically, our method outperforms GMA and MS-RAFT(+) by up to 23.0% and 21.6%, respectively. Also, pre-training

<sup>&</sup>lt;sup>3</sup>One unified architecture for each variant used on both benchmarks.

	Sinte	el Clean (	est) Sintel Final (test)			KITTI (test)			
Method	All↓	Mat↓	Unmat↓	All↓	Mat↓	Unmat↓	Method	Fl-all↓	Fl-noc↓
RAFT [40] ECCV'20	1.61	0.62	9.65	2.86	1.41	14.68	RAFT [40] ECCV'20	5.10	3.07
RAFT-it [37] ECCV'22	1.55	0.61	9.24	2.90	1.41	15.03	GMA <sup>†</sup> [18] ICCV'21	4.93	2.90
DIP [51] CVPR'22	1.44	0.52	8.92	2.83	1.28	15.49	MS-RAFT [16] ICIP'22	4.88	2.80
RAFT-OCTC [17] CVPR'22	1.42	0.54	8.57	2.57	1.24	13.44	FlowFormer [10] ECCV'22	4.68	2.69
GMA [18] ICCV'21	1.39	0.58	7.96	2.47	1.24	12.50	MatchFlow_GMA [6] CVPR'23	4.63	2.77
MS-RAFT [16] ICIP'22	1.37	0.48	8.67	2.67	1.20	14.70	KPA-Flow [25] CVPR'22	4.60	2.82
KPA-Flow [25] CVPR'22	1.35	0.50	8.23	2.27	1.09	11.89	FlowFormer++ [35] CVPR'23	4.52	-
MS-RAFT+ [15] arXiv'22	1.23	0.40	8.02	2.68	1.28	14.12	GMFlow+ [45] TPAMI'23	4.49	2.40
AnyFlow [19] CVPR'23	1.23	-	-	2.44	-	-	AnyFlow [19] CVPR'23	4.41	2.69
FlowFormer [10] ECCV'22	1.14	-	-	2.18	-	-	RAFT-OCTC [17] CVPR'22	4.33	-
MatchFlow_GMA [6] CVPR'23	1.16	0.43	7.13	2.37	1.06	13.07	RAFT-it [37] ECCV <sup>*</sup> 22	4.31	-
FlowFormer++ [35] CVPR'23	1.07	0.39	6.64	1.94	0.88	10.63	DIP [51] CVPR'22	4.21	2.43
GMFlow+ [45] TPAMI'23	1.03	<u>0.34</u>	<u>6.68</u>	2.37	1.10	12.74	MS-RAFT+ [15] arXiv'23	4.15	<u>2.18</u>
CCMR (ours)	1.19	0.41	7.50	2.14	1.00	11.44	CCMR (ours)	4.04	2.44
vs. GMA	+14.4%	+29.3%	+5.8%	+13.3%	+19.4%	+8.5%	vs. GMA	+18.1%	+15.9%
vs. MS-RAFT	+13.1%	+14.6%	+13.5%	+19.9%	+16.7%	+22.2%	vs. MS-RAFT	+17.2%	+12.9%
CCMR+ (ours)	1.07	0.31	7.24	2.10	0.93	11.58	CCMR+ (ours)	3.86	2.07
vs. GMA	+23.0%	+46.6%	+9.1%	+15.0%	+25.0%	+7.4%	vs. GMA	+21.7%	+28.6%
vs. MS-RAFT+	+13.0%	+22.5%	+9.7%	+21.6%	+27.3%	+18.0%	vs. MS-RAFT+	+7.0%	+5.1%

Table 1. Comparison to the SOTA and the baselines on Sintel (test) and KITTI (test) using the AEPE (Sintel) and the Fl error (KITTI). Mat = matched. Unmat = unmatched. Best result in bold, second best result is underlined. Results taken from the corresponding papers.

<sup>†</sup> GMA uses different architectures on Sintel and KITTI. Using the Sintel architecture for KITTI yields Fl-all of 5.15.

Table 2. Pre-training results. Results on *Sintel (train)* and *KITTI (train)* after pre-training with *Chairs* and *Things*. '-': not available.

	Sintel	(train)	KITTI (train)	
Method	Clean↓	Final↓	Fl↓	
RAFT	1.43	2.71	17.4	
GMA	1.30	2.74	17.1	
DIP	1.30	2.82	13.73	
KPA-flow	1.28	2.68	15.9	
AnyFlow	1.17	2.58	13.01	
MS-RAFT	1.13	2.60	-	
MatchFlow_GMA	1.03	2.45	15.6	
Flowformer++	0.90	2.30	14.13	
GMFlow+	<u>0.91</u>	-	-	
CCMR (ours)	1.07	2.40	13.30	
CCMR+ (ours)	0.98	<u>2.36</u>	12.84	

generalization results show consistent improvements of our method over the baselines (see suppl. material).

**Comparison to the State of the Art.** Comparing the performance of our CCMR(+) models to the best reported results in the published literature, Tab. 1 shows that our approach performs favorably. While our 3-scale CCMR model achieves very good results on both benchmarks, Our 4-scale CCMR+ models set a new state of the art on KITTI (3.86) and ranks second on Sintel Clean (1.07) and Final (2.10). In this context, our method shows good results in unmatched regions while it achieves an unprecedented accuracy for Sintel Clean (0.31) and KITTI (2.07) in matched

areas—even outperforming the leading *scene flow* method on KITTI (2.07 vs. 2.18, see [23] for results of CamLi-RAFT). Furthermore, it is noteworthy that, in contrast to other methods from the literature our approach achieves excellent results on *both* Sintel *and* KITTI. For instance, GM-Flow+ and FlowFormer++ achieve excellent results on Sintel clean (FlowFormer++ also on final), but are not among the SOTA on KITTI. RAFT-it and DIP in turn, that perform great on KITTI, offer less competitive results on Sintel. Finally, we report our pre-training generalization results on Sintel (train) and KITTI (train) in Tab. 2. Also in this case, our models yield very good results.

**Visual Comparison.** The previous findings are also confirmed by a visual comparison of the results in Fig. 7. Apart from our CCMR+ model, we also consider two recent SOTA methods for each benchmark as well as the GMA and MS-RAFT+ baseline, whenever available. The corresponding flow fields show that our model preserves more details both in the foreground and the background (cf. boxes).

**Architectural Impact.** As outlined in Sec. 4.1, our models were fined-tuned for KITTI. While this is common practice in the literature (see, e.g., [10, 17, 37, 39, 40]), our multiscale baselines MS-RAFT and MS-RAFT+ do not consider this additional fine-tuning step. Hence, to delineate the actual architectural improvements from the impact of the fine-tuning, we fined-tuned both our models (CCMR, CCMR+) and the multi-scale baselines (MS-RAFT, MS-RAFT+) on

Table 3. Comparison of our CCMR(+) models with MS-RAFT(+) and GMA baselines on a *KITTI validation split* after *fine-tuning*. The same training protocol has been used in all cases such that observed differences are only due to the improved architecture.

	GMA	MS-RAFT	CCMR	MS-RAFT+	CCMR+
KITTI-val	5.01	5.21	4.75	<u>4.71</u>	4.57
#Scales	1	3	3	4	4
Finest Res.	$\frac{1}{8}(h,w)$	$\frac{1}{4}(h,w)$	$\frac{1}{4}(h,w)$	$\frac{1}{2}(h,w)$	$\frac{1}{2}(h,w)$

KITTI using the same training scheme for a fair comparison. To this end, we excluded a part of the KITTI training data in the preceding mixed training and the subsequent fine-tuning and used this missing data afterwards as *validation split* (see suppl. material). The corresponding results in Tab. 3 show clear improvements from CCMR over MS-RAFT and CCMR+ over MS-RAFT+ as well as improvements when going from three to four scales. For completeness, we also included the result of our conceptual baseline, GMA, with the same training scheme.

### 4.3. Ablations

After we have evaluated our models, we now ablate their underlying architecture. To this end, we compare the results of our method with different architectural configurations on *Sintel (train)* [5] after pre-training the variants with *Chairs* [7] and *Things* [27]. Due to the large computational effort, we performed the ablations on the 3-scale CCMR model. Finally, we also present results for our 4-scale CCMR+ model for the selected setting.

Architectural Components. In our first ablation, we study the impact of the improved feature consolidation unit and of the global-context based motion reasoning. The results in Tab. 4 show that both concepts improve the estimation. While the context-based motion grouping has more impact on unmatched regions, the improved encoders have more influence in matched areas.

Motion Grouping Based on Global Context. In the next ablation, we disentangle the impact of using global context features and motion grouping for components with weights shared across scales or individual weights per scale; see Tab. 5. For the first two rows, no motion grouping was performed, and the global context was used directly in the recurrent unit for flow updates. In row three and four, no global context was computed and the motion grouping was done based on the un-aggregated context features. In the remaining rows, we investigate the use of both concepts using shared or individual components across the coarse-tofine scales, as well as going from the 3-scale CCMR model to the 4-scale CCMR+ model. The results show that both global context and motion grouping contribute to the improvements, with the best overall performance obtained by

Table 4. Ablation on the *global-context based motion grouping* (GC-MG) and the *improved feature consolidation unit* (IFCU) in all, matched and unmatched regions. Results on *Sintel (train)* after pre-training with *Chairs* and *Things*. Gray: CCMR configuration.

		Sinte	Sintel Clean (train)			Sintel Final (train)			
GC-MG	IFCU	All↓	Mat↓	Unmat↓	-	All↓	Mat↓	Unmat↓	
-	-	1.18	0.47	10.20		2.58	1.46	16.81	
$\checkmark$	-	1.10	<u>0.46</u>	<u>9.23</u>		<u>2.47</u>	<u>1.42</u>	<u>15.89</u>	
$\checkmark$	$\checkmark$	1.07	0.43	9.12		2.40	1.35	15.76	

Table 5. Ablation on *individual/shared* modules for global context computation and context-based motion grouping across scales. Results on *Sintel (train)* after pre-training with *Chairs* and *Things*. Gray: CCMR and CCMR+ configuration.

Global Context Computation	Motion Grouping	#Scales	$\frac{\text{Sintel (train)}}{\text{Clean All}\downarrow}$	$\frac{\text{Sintel (train)}}{\text{Final All}\downarrow}$
shared	-	3	1.11	2.50
individual	-	3	<u>1.06</u>	2.53
-	shared	3	1.08	2.50
-	individual	3	1.06	2.48
shared	shared	3	1.08	2.45
shared	individual	3	1.03	<u>2.45</u>
individual	shared	3	1.12	2.60
individual	individual	3	1.07	2.40
individual	individual	3	1.07	2.40
individual	individual	4	0.98	2.36

the model with individual units for each scale. Moreover, they demonstrate that going from three to four scales once again improves the results.

# 5. Conclusion

We presented CCMR, a high-accuracy optical flow approach that benefits from attention-based motion grouping concepts as well as from strong matching capabilities of high-resolution multi-scale neural networks. It relies on computing global context, which serves as guidance to perform motion reasoning across all scales via global channel statistics, and exploits an improved feature consolidation unit, that extracts more expressive features for matching. Experiments show strong improvements in both occluded and non-occluded regions, resulting from the global-context based motion reasoning and the improved feature consolidation, respectively. In this way, our approach produces highly detailed flow fields in the foreground and background, while achieving state-of-the-art results on KITTI 2015 (Rank 1) and MPI Sintel (Rank 2 on Clean and Final).

Acknowledgements. The authors thank the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 251654672 – TRR 161 (B04) and the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Azin Jahedi.

# References

- A. Ali, H. Touvron, M. Caron, P. Bojanowski, M. Douze, A. Joulin, I. Laptev, et al. XCiT: cross-covariance image transformers. In *NeuRIPS*, pages 20014–20027, 2021. 2, 3, 4, 5
- [2] P. Anandan. A computational framework and an algorithm for the measurement of visual motion. *IJCV*, 2(3):283–310, 1989.
- [3] J. L. Ba, J. R. Kiros, and G. E. Hinton. Layer normalization. arXiv:1607.06450, 2016. 4
- [4] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *ECCV*, pages 25–36, 2004. 2
- [5] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, pages 611–625, 2012. 6, 8
- [6] Q. Dong, C. Cao, and Y. Fu. Rethinking optical flow from geometric matching consistent perspective. In *CVPR*, pages 1337–1347, 2023. 1, 2, 3, 7
- [7] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. In *ICCV*, pages 2758–2766, 2015. 1, 6, 8
- [8] H. Fan, B. Xiong, K. Mangalam, Y. Li, Z. Yan, J. Malik, and C. Feichtenhofer. Multiscale vision transformers. In *ICCV*, pages 6824–6835, 2021. 2
- [9] M. Hofinger, S. R. Bulò, L. Porzi, A. Knapitsch, T. Pock, and P. Kontschieder. Improving optical flow on a pyramid level. In *ECCV*, pages 770–786, 2020. 2
- [10] Z. Huang, X. Shi, C. Zhang, Q. Wang, K. C. Cheung, H. Qin, J. Dai, and H. Li. FlowFormer: a transformer architecture for optical flow. In *ECCV*, pages 668– 685, 2022. 1, 2, 7
- [11] T.-W. Hui, X. Tang, and C. C. Loy. LiteFlowNet: A lightweight convolutional neural network for optical flow estimation. In *CVPR*, pages 8981–8989, 2018. 2
- [12] J. Hur and S. Roth. Iterative residual refinement for joint optical flow and occlusion estimation. In *CVPR*, pages 5747–5756, 2019. 2, 3
- [13] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, pages 1647–1655, 2017. 1
- [14] E. Ilg, T. Saikia, M. Keuper, and T. Brox. Occlusions, Motion and Depth Boundaries with a Generic Network for Disparity, Optical Flow or Scene Flow Estimation. In *ECCV*, pages 626–643, 2018. 3

- [15] A. Jahedi, M. Luz, L. Mehl, M. Rivinius, and A. Bruhn. High resolution multi-scale RAFT (ECCV Robust Vision Challenge 2022). arXiv:2210.16900, 2022. 2, 3, 4, 6, 7
- [16] A. Jahedi, L. Mehl, M. Rivinius, and A. Bruhn. Multi-Scale RAFT: Combining hierarchical concepts for learning-based optical flow estimation. In *ICIP*, pages 1236–1240, 2022. 2, 3, 4, 6, 7
- [17] J. Jeong, J. Menjay Lin, F. Porikli, and N. Kwak. Imposing consistency for optical flow estimation. In *CVPR*, pages 3171–3181, 2022. 3, 7
- [18] S. Jiang, D. Campbell, Y. Lu, H. Li, and R. Hartley. Learning to estimate hidden motions with global motion aggregation. In *ICCV*, pages 9772–9781, 2021. 1, 2, 3, 5, 7
- [19] H. Jung, Z. Hui, L. Luo, H. Yang, F. Liu, S. Yoo, R. Ranjan, and D. Demandolx. AnyFlow: arbitrary scale optical flow with implicit neural representation. In *CVPR*, pages 5455–5465, 2023. 1, 2, 3, 7
- [20] D. Kondermann, R. Nair, K. Honauer, K. Krispin, J. Andrulis, et al. The HCI benchmark suite: Stereo and flow ground truth with uncertainties for urban autonomous driving. In *CVPR Workshops*, pages 19–28, 2016. 6
- [21] Y. Li, C.-Y. Wu, H. Fan, K. Mangalam, B. Xiong, J. Malik, and C. Feichtenhofer. MViTv2: Improved multiscale vision transformers for classification and detection. In *CVPR*, pages 4804–4814, 2022. 2
- [22] T. Lin, P. Dollar, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature Pyramid Networks for Object Detection. In *CVPR*, pages 936–944, 2017. 4
- [23] H. Liu, T. Lu, Y. Xu, J. Liu, W. Li, and L. Chen. CamLiFlow: bidirectional camera-LiDAR fusion for joint optical flow and scene flow estimation. In *CVPR*, pages 5791–5801, 2022. 7
- [24] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. 2
- [25] A. Luo, F. Yang, X. Li, and S. Liu. Learning optical flow with kernel patch attention. In *CVPR*, pages 8896–8905, 2022. 2, 3, 7
- [26] M. Maaz, A. Shaker, H. Cholakkal, S. Khan, S. W. Zamir, R. M. Anwer, et al. Edgenext: efficiently amalgamated CNN-transformer architecture for mobile vision applications. In *ECCV Workshops*, pages 3–20, 2023. 3
- [27] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A large dataset to train convolutional networks for disparity, optical flow, and

scene flow estimation. In *CVPR*, pages 4040–4048, 2016. 6, 8

- [28] M. Menze and A. Geiger. Object scene flow for autonomous vehicles. In *CVPR*, pages 3061–3070, 2015.
- [29] M. Neoral, J. Šochman, and J. Matas. Continual Occlusion and Optical Flow Estimation. In ACCV, pages 159–174, 2019. 3
- [30] P. Panteleris and A. Argyros. Pe-former: Pose estimation transformer. In *ICPRAI*, pages 3–14, 2022. 3
- [31] L. Phan, H. T. H. Nguyen, H. Warrier, and Y. Gupta. Patch embedding as local features: Unifying deep local and global features via vision transformer for image retrieval. In ACCV, pages 2527–2544, 2022. 3
- [32] P. Ramachandran, N. Parmar, A. Vaswani, I. Bello, et al. Stand-alone self-attention in vision models. In *NeurIPS*, pages 68–80, 2019. 2
- [33] A. Ranjan and M. J. Black. Optical flow estimation using a spatial pyramid network. In *CVPR*, pages 4161– 4170, 2017. 1, 2
- [34] S. R. Richter, Z. Hayder, and V. Koltun. Playing for benchmarks. In *ICCV*, pages 2232–2241, 2017. 6
- [35] X. Shi, Z. Huang, D. Li, M. Zhang, K. C. Cheung, S. See, H. Qin, J. Dai, and H. Li. FlowFormer++: masked cost volume autoencoding for pretraining optical flow estimation. In *CVPR*, pages 1599–1610, 2023. 1, 2, 7
- [36] X. Sui, S. Li, X. Geng, Y. Wu, X. Xu, Y. Liu, R. S. M. Goh, and H.an Zhu. CRAFT: cross-attentional flow transformer for robust optical flow. In *CVPR*, pages 17581–17590, 2022. 1, 3
- [37] D. Sun, C. Herrmann, F. A. Reda, M. Rubinstein, D. J. Fleet, and W. T. Freeman. Disentangling architecture and training for optical flow. In *ECCV*, pages 165– 182, 2022. 7
- [38] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *CVPR*, pages 8934–8943, 2018. 1, 2
- [39] S. Sun, Y. Chen, Y. Zhu, G. Gou, and G. Li. Learning optical flow with super kernels. In *NeurIPS*, pages 11313–11326, 2022. 2, 3, 7
- [40] Z. Teed and J. Deng. RAFT: Recurrent all-pairs field transforms for optical flow. In *ECCV*, pages 402–419, 2020. 1, 2, 5, 7
- [41] A. Vaswani, P. Ramachandran, A. Srinivas, N. Parmar, B. Hechtman, and J. Shlens. Scaling local selfattention for parameter efficient visual backbones. In *CVPR*, pages 12894–12904, 2021. 2

- [42] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017. 2, 4, 5, 6
- [43] W. Wang, E. Xie, X. Li, D.-P. Fan, K. Song, D. Liang, T. Lu, P. Luo, and L. Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, pages 568–578, 2021.
  2
- [44] H. Xu, J. Zhang, J. Cai, H. Rezatofighi, and D. Tao.
   GMFlow: Learning optical flow via global matching. In *CVPR*, pages 8111–8120, 2022. 1, 2, 3
- [45] H. Xu, J. Zhang, J. Cai, H. Rezatofighi, F. Yu, D. Tao, and A. Geiger. Unifying flow, stereo and depth estimation. *TPAMI*, 2023. 1, 2, 3, 7
- [46] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, and Ming-H. Yang. Restormer: Efficient Transformer for High-Resolution Image Restoration. In *CVPR*, pages 5728–5739, 2022. 3
- [47] F. Zhang, O. J. Woodford, V. A. Prisacariu, and P. H. S. Torr. Separable Flow: Learning motion cost volumes for optical flow estimation. In *CVPR*, pages 10807– 10817, 2021. 1, 2, 3
- [48] P. Zhang, X. Dai, J. Yang, B. Xiao, L. Yuan, L. Zhang, and J. Gao. Multi-scale vision longformer: A new vision transformer for high-resolution image encoding. In *ICCV*, pages 2978–2988, 2021. 2
- [49] S. Zhao, Y. Sheng, Y. Dong, E. Chang, and Y. Xu. MaskFlownet: Asymmetric feature matching with learnable occlusion mask. In *CVPR*, pages 6278– 6287, 2020. 3
- [50] S. Zhao, L. Zhao, Z. Zhang, E. Zhou, and D. N. Metaxas. Global matching with overlapping attention for optical flow estimation. In *CVPR*, pages 17571– 17580, 2022. 1
- [51] Z. Zheng, N. Nie, Z. Ling, P. Xiong, J. Liu, H. Wang, and J. Li. DIP: Deep inverse patchmatch for highresolution optical flow. In *CVPR*, pages 8925–8934, 2022. 2, 3, 7