

Video Frame Interpolation with Densely Queried Bilateral Correlation

Chang Zhou, Jie Liu*, Jie Tang and Gangshan Wu

State Key Laboratory for Novel Software Technology, Nanjing University, China

zhouchang@smail.nju.edu.cn, {liujie, tangjie, gswu}@nju.edu.cn

Abstract

Video Frame Interpolation (VFI) aims to synthesize non-existent intermediate frames between existent frames. Flow-based VFI algorithms estimate intermediate motion fields to warp the existent frames. Real-world motions’ complexity and the reference frame’s absence make motion estimation challenging. Many state-of-the-art approaches explicitly model the correlations between two neighboring frames for more accurate motion estimation. In common approaches, the receptive field of correlation modeling at higher resolution depends on the motion fields estimated beforehand. Such receptive field dependency makes common motion estimation approaches poor at coping with small and fast-moving objects. To better model correlations and to produce more accurate motion fields, we propose the Densely Queried Bilateral Correlation (DQBC) that gets rid of the receptive field dependency problem and thus is more friendly to small and fast-moving objects. The motion fields generated with the help of DQBC are further refined and up-sampled with context features. After the motion fields are fixed, a CNN-based SynthNet synthesizes the final interpolated frame. Experiments show that our approach enjoys higher accuracy and less inference time than the state-of-the-art. Source code is available at <https://github.com/kinoud/DQBC>.

1 Introduction

Video Frame Interpolation (VFI) aims to synthesize non-existent intermediate frames between existent frames. VFI is used for a wide range of applications such as slow-motion generation [Jiang *et al.*, 2018], video compression [Wu *et al.*, 2018], video restoration [Haris *et al.*, 2020; Kim *et al.*, 2020; Xiang *et al.*, 2020], novel view synthesis [Zhou *et al.*, 2016; Kalantari *et al.*, 2016], animation production [Siyao *et al.*, 2021; Chen and Zwicker, 2022; Shen *et al.*, 2022] and video prediction [Wu *et al.*, 2022]. Nowadays, the flow-based paradigm dominates the methodology of VFI owing to its superior performance. However, the

*Corresponding author.

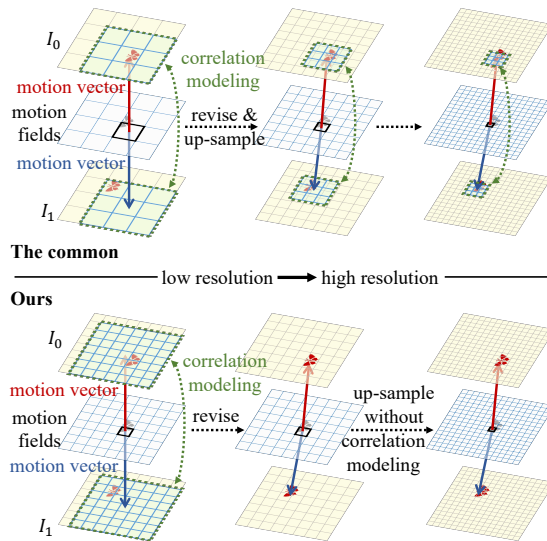


Figure 1: The common correlation modeling paradigm vs. ours. The correlation modeling is conducted along with the motion estimation process. The common approach iteratively models correlations from low resolution to high resolution. Our approach models correlations at a single resolution.

absence of the reference frame makes it challenging to estimate the intermediate optical flow fields or motion fields used to warp the neighboring frames backward. [Bao *et al.*, 2019; Jiang *et al.*, 2018; Liu *et al.*, 2019; Niklaus and Liu, 2018; Xue *et al.*, 2019] approximate the intermediate motion fields by linearly combining the optical flow fields between the two input frames. [Reda *et al.*, 2022; Huang *et al.*, 2022; Kong *et al.*, 2022] directly estimate the motion fields using Convolutional Neural Networks (CNN) in a coarse-to-fine manner. [Park *et al.*, 2020; Park *et al.*, 2021; Lu *et al.*, 2022] explicitly model correlations between two input frames, which appears as a strong prior in the VFI task and helps to produce more accurate motion fields.

Correlation modeling is often conducted locally in a coarse-to-fine manner coupled with the motion fields estimation as in [Park *et al.*, 2020; Park *et al.*, 2021; Lu *et al.*, 2022]. The top half (“The common”) of Figure 1 illustrates this idea. At each resolution level, for every location in the

motion fields, two local windows on I_0 and I_1 are picked to model correlations between them. In correlation modeling, similarities of patches in two local windows (Each local window contains 3×3 patches in the top half of Figure 1.) are computed. These similarity values play an important role in motion fields revisions performed by neural networks. After being revised, the motion fields are up-sampled, and correlation modeling is conducted again at a higher resolution. This process is repeated across several resolution levels.

However, such a paradigm suffers from the “receptive field dependency” problem, which will be explained in detail next. As shown in the top half of Figure 1, correlation modeling is conducted within two local windows, and the locations of the two local windows are determined by motion vectors. Considering the local windows are relatively small at higher resolution, if the motion vectors estimated at low resolution are inaccurate, the local windows at higher resolution may not cover the butterfly at all. In that case, the correlation modeling at higher resolution does not work as expected because the local windows are not placed in the intended areas. In other words, the receptive field of correlation modeling depends on the motion fields estimated beforehand. So, in the common approach shown in the top half of Figure 1, we expect that the motion fields estimated at lower resolution are reasonably accurate to prevent misleading the correlation modeling at higher resolution. However, it is not fair to expect the motion fields estimated at low resolution to be “reasonably accurate”, especially for small objects. The reasons are as follows: At low resolution, patches used for computing similarities are relatively large, and small objects may only occupy a small fraction of the patch. Thus the enormous background may dominate the similarity computing regarding patches that contain small objects. Besides, suppose a large patch on one frame includes two or more small objects, and these small objects are separated across different patches on the other frame. In that case, the correlation modeling and motion estimation at low resolution may face ambiguity.

Low-resolution correlation modeling is expected to cope with fast-moving objects but not small objects, and high-resolution correlation modeling is expected to cope with small objects but not fast-moving objects. Superficially, they are complementary. But considering the aforementioned receptive field dependency problem, if low-resolution correlation modeling fails to capture small and fast-moving objects (which may easily happen), high-resolution correlation modeling may also miss that object. The receptive field dependency problem makes the common approach poor at coping with small and fast-moving objects.

In this work, we propose the Densely Queried Bilateral Correlation (DQBC) to better model correlations between two frames and produce more accurate motion fields. We conduct correlation modeling differently compared with the common approach. The bottom half of Figure 1 illustrates our paradigm. Instead of modeling correlations iteratively from very low resolution, we model correlations all at once at a single and relatively high resolution. Within the local windows, DQBC is calculated asymmetrically: When computing similarities between patches, query patches keep their resolution and key patches are down-sampled several times

for larger receptive fields. DQBC gets rid of the receptive field dependency problem and helps generate more accurate motion fields, especially for small and fast-moving objects. To refine and up-sample the motion fields to full size, we further devise a Motion Refinement Module (MRM). MRM refines and up-samples the input motion fields with the help of context features. Compared with [Lu *et al.*, 2022] that uses heavy transformers to synthesize the intermediate frame after motion estimation, we devise a relatively lightweight CNN called SynthNet to synthesize the final result.

Our contributions are summarized as follows:

- We propose the Densely Queried Bilateral Correlation (DQBC) to better model correlations for the VFI task. DQBC gets rid of the receptive field dependency problem and thus is more friendly to small and fast-moving objects.
- With the DQBC, our VFI framework achieves state-of-the-art performance on various VFI benchmarks. In our VFI framework, we devise a Motion Refinement Module (MRM) to refine and up-sample the motion fields and a CNN-based SynthNet to synthesize the interpolated frame.

2 Related Work

VFI is a long-standing task, and various approaches have been proposed.

Kernel-based approaches such as [Niklaus *et al.*, 2017a; Niklaus *et al.*, 2017b; Cheng and Chen, 2020; Cheng and Chen, 2021; Lee *et al.*, 2020] estimate spatially-adaptive convolutional kernels for each pixel to be synthesized. [Meyer *et al.*, 2015; Meyer *et al.*, 2018] represent motions in the phase shifts of individual pixels. [Choi *et al.*, 2020; Choi *et al.*, 2021] distribute spatial information into the channel dimension and perform channel attention for pixel-level frame synthesis.

Besides these, the flow-based approaches enjoy a long history and have been prevailing. In recent flow-based VFI works, [Niklaus and Liu, 2018] proposes to warp the context features using the motion fields for better interpolation quality. [Jiang *et al.*, 2018] proposes a network for variable-length multi-frame video interpolation. [Bao *et al.*, 2019] exploits the additional depth information to better handle occlusions in the sampling process. [Peleg *et al.*, 2019; Sim *et al.*, 2021] explore high-resolution video interpolation. [Tulyakov *et al.*, 2021] explores the event-based approach. [Park *et al.*, 2020] proposes to model correlations between two frames explicitly via cost volume and is improved by [Park *et al.*, 2021] who further models the asymmetric motions. [Huang *et al.*, 2022; Kong *et al.*, 2022] aim at efficient VFI algorithms and yet have shown compelling performance. [Lu *et al.*, 2022] proposes a transformer-based synthesis network to blend two warped frames better and hallucinate details.

Our work follows the flow-based paradigm and gives focuses on correlation modeling. Among the approaches mentioned above, [Park *et al.*, 2020; Park *et al.*, 2021; Lu *et al.*, 2022] explicitly model the correlations between two frames. [Park *et al.*, 2020] computes similarities between any two centrosymmetric patch pairs between two local windows.




Figure 2: Overview of our VFI pipeline. Firstly, the Densely Queried Bilateral Correlation ($DQBC$) is extracted. The Motion Generation Module (MGM) generates preliminary motion fields with the help of $DQBC$. The Motion Refinement Module (MRM) refines and up-samples the motion fields to full size and produces a preliminary occlusion map O . The SynthNet estimates a residual occlusion map ΔO and a residual image R for final synthesizing.




Figure 3: Extraction process of $DQBC$. (a): Tokens in F_1 are taken as queries and tokens in local windows on several down-sampled F_0 are taken as keys to compute similarities. The unilateral correlation embeddings are then enhanced by the enhancement block and spatially aligned by the feature distributing operation. $DQBC$ is the concatenation of two bilateral correlation embeddings. (b): A toy example illustrating the intuition of spatial alignment. The feature map of unilateral correlation embeddings is spatially aligned with I_0 but not with $I_{\frac{1}{2}}$. To be aligned with $I_{\frac{1}{2}}$, the entire feature map is shifted by $\frac{1}{2}v_c$.

[Park *et al.*, 2021] further models the correlation between the intermediate anchor frame produced by [Park *et al.*, 2020] and two input frames. [Lu *et al.*, 2022] adopts the IFNet [Huang *et al.*, 2022] without explicit correlation modeling as the base motion estimator and further refines the motion fields with explicit correlation modeling. It computes similarities between the center patch of one local window and all patches of the other. Despite the difference in the modeling details, they all model correlations locally in a coarse-to-fine manner, as illustrated in the top half of Figure 1.

Our work takes two frames as inputs. There are also VFI works that take four input frames and are beyond the scope of this paper. We refer interested readers to [Xu *et al.*, 2019; Liu *et al.*, 2020; Chi *et al.*, 2020; Shi *et al.*, 2022; Danier *et al.*, 2022; Yu *et al.*, 2022] for more details.

3 The Proposed Method

3.1 Overview

Given 2 adjacent frames I_0 and I_1 , we synthesize the interpolated frame $\hat{I}_t(t = 0.5)$ as:

$$\hat{I}_t = O \cdot \overleftarrow{w}(I_0, M_{t \rightarrow 0}) + (1 - O) \cdot \overleftarrow{w}(I_1, M_{t \rightarrow 1}) + R. \quad (1)$$

In Equation (1), $M_{t \rightarrow \tau} \in \mathbb{R}^{H \times W \times 2}$ is the estimated optical flow field from \hat{I}_t to I_τ where $\tau \in \{0, 1\}$. They are referred to as *motion fields*. $O \in [0, 1]^{H \times W}$ is the occlusion map and $R \in \mathbb{R}^{H \times W \times 3}$ is the image residual. The $\overleftarrow{w}(x, y)$

operation is to spatially sample the input tensor x with the vector field y .

As shown in Figure 2, our VFI pipeline can be roughly divided into three stages: (1) *Correlation Modeling*, (2) *Motion Estimation*, and (3) *Frame Synthesis*. In stage (1), the Densely Queried Bilateral Correlation ($DQBC$) is extracted from two input frames. In stage (2), $DQBC$ helps to generate preliminary motion fields that are further refined and up-sampled by the Motion Refinement Module (MRM). MRM also produces an initial occlusion map. In stage (3), the SynthNet predicts the image residual and the occlusion map residual. Then the interpolated frame is synthesized following Equation (1).

3.2 Densely Queried Bilateral Correlation

The proposed $DQBC$ explicitly encodes the comprehensive correlation information needed for motion estimation.

As shown in Figure 3(a), $DQBC$ extraction consists of three parts: (1) *Unilateral Correlation Gathering*, (2) *Correlation Enhancement* (via “ $\text{Enh}(\cdot)$ ”), (3) *Spatial Alignment* (via “ $\text{Dist}(\cdot)$ ”), which can be summarized as follows:

$$\begin{aligned} BiCorr_{t \rightarrow 0} &= \text{Dist}_{0 \rightarrow t}(\text{Enh}(UniCorr_{0 \rightarrow 1})), \\ BiCorr_{t \rightarrow 1} &= \text{Dist}_{1 \rightarrow t}(\text{Enh}(UniCorr_{1 \rightarrow 0})), \\ DQBC &= [BiCorr_{t \rightarrow 0}, BiCorr_{t \rightarrow 1}], \end{aligned} \quad (2)$$

where $[\cdot, \cdot]$ denotes concatenation along channels.

Unilateral Correlation Gathering

As shown in Figure 3(a), we take one frame as the query frame and the other as the key frame to gather unilateral correlations. Firstly, a feature extractor $f(\cdot)$ extracts high-dimensional visual features F_0, F_1 from I_0, I_1 . F_1 is kept as a query set and won't ever be down-sampled. Thus the queries are spatially dense. F_0 is down-sampled $(L - 1)$ times by factor $1/2$ to form the key sets of different resolution. For every location (x_q, y_q) on F_1 and each resolution level $l \in \{0, 1, \dots, L - 1\}$, we calculate the embedding vector $UniCorr_{1 \rightarrow 0}^l(x_q, y_q)$ whose i^{th} element is:

$$UniCorr_{1 \rightarrow 0}^l(x_q, y_q)_i = \langle F_1(x_q, y_q), F_0^l(\frac{x_q + x_i^i}{2^l}, \frac{y_q + y_i^i}{2^l}) \rangle, \quad (3)$$

where $\langle \cdot, \cdot \rangle$ is the inner product operation and F_0^l is spatially down-sampled F_0 by factor $1/2^l$ and (x_i^i, y_i^i) is the i^{th} location in the local window LW_l defined as:

$$LW_l = \{(2^l i, 2^l j) | i, j \in \{-r_l, -r_l + 1, \dots, r_l\}\}, \quad (4)$$

where r_l is the radius of this local window.

Then the embedding vectors from all resolution levels at (x_q, y_q) are concatenated together to form the unilateral correlation embedding for (x_q, y_q) :

$$UniCorr_{1 \rightarrow 0}(x_q, y_q) = \text{Concatenate}_{l \in \{0, 1, \dots, L-1\}} \{UniCorr_{1 \rightarrow 0}^l(x_q, y_q)\}. \quad (5)$$

Note that $UniCorr_{0 \rightarrow 1}$ can be acquired in a similar way.

Correlation Enhancement

In Figure 3(a), the correlation enhancement block $Enh(\cdot)$ performs denoising and enhancement on unilateral correlation embeddings. Due to occlusions and 3D rotations, some objects may only appear on one of the two input frames. Thus the vanilla unilateral correlation embeddings may contain noise. The $Enh(\cdot)$ consists of two convolutional layers and a skip connection.

Spatial Alignment

The enhanced unilateral correlation embeddings are not spatially aligned with the intermediate frame. In Figure 3(a), the feature distributing operation $Dist(\cdot)$ performs spatial alignment by shifting feature maps.

We first define the displacement vector for each feature map of $UniCorr$. According to Equation (3), the i^{th} feature map of $UniCorr_{1 \rightarrow 0}^l$ encodes the similarities between query locations on I_1 and key locations on I_0 . We define the displacement vector as the displacement from the query location to the key location. Thus according to Equation (3), the displacement vector \mathbf{v}_i of the i^{th} feature map of $UniCorr_{1 \rightarrow 0}^l$ is just (x_i^i, y_i^i) . Equation (5) implies that every feature map of $UniCorr_{1 \rightarrow 0}$ also has its displacement vector.

The $Dist(\cdot)$ function in Equation (2) shifts each feature map of the input embeddings by a fraction of its displacement vector. $Dist(\cdot)$ shifts each feature map of its input by $t\mathbf{v}$ where $\frac{0 \rightarrow t}{1 \rightarrow t}$ shifts each feature map of its input by $(1 - t)\mathbf{v}$ where \mathbf{v} is the displacement vector of the feature map.

Figure 3(b) illustrates the intuition of $Dist(\cdot)$. A butterfly appears at $P_0, P_{\frac{1}{2}}, P_1$ on $I_0, I_{\frac{1}{2}}, I_1$ respectively. Following Equation (3) and Equation (5), we first gather the unilateral correlation embeddings from I_0 to I_1 , denoted as $Corr_{0 \rightarrow 1}$. By definition, the value at (x, y) on the c^{th} feature map of $Corr_{0 \rightarrow 1}$ (denoted as $Corr_{0 \rightarrow 1, c}$) is the similarity between (x, y) on I_0 and $(x, y) + \mathbf{v}_c$ on I_1 where \mathbf{v}_c is the displacement vector of this feature map. If $P_0 + \mathbf{v}_c$ is reasonably close to P_1 , $Corr_{0 \rightarrow 1, c}$ will have a relatively high activation value at P_0 . In $Dist_{0 \rightarrow \frac{1}{2}}(\cdot)$, we displace this high activation value at P_0 to be near $P_{\frac{1}{2}}$ by shifting the entire feature map by $\frac{1}{2}\mathbf{v}_c$.

Now we have described every part of the $DQBC$ extraction process summarized in Equation (2). The $DQBC$ is $1/8$ in spatial resolution of I_0 and I_1 since the feature extractor $f(\cdot)$ contains three convolutional layers with stride 2.

3.3 Motion Estimation

As shown in the middle part of Figure 2, in the motion estimation stage, the Motion Generation Module (MGM) generates preliminary motion fields facilitated by $DQBC$ and then the Motion Refinement Module (MRM) refines and up-samples the motion fields to full size.

Motion Generation

With the extracted $DQBC$, MGM generates the preliminary motion fields as:

$$[\tilde{M}_{t \rightarrow 0}, \tilde{M}_{t \rightarrow 1}] = \mathbf{g}([\mathbf{c}([I_0, I_1]), \mathbf{m}(DQBC)]). \quad (6)$$

In Equation (6), $[\cdot, \cdot]$ denotes concatenation along channels. $\mathbf{c}(\cdot)$ is a CNN that extracts context information from two input frames. $\mathbf{m}(\cdot)$ is a Multi-Layer Perceptron (MLP) with one hidden layer to perform dimensionality reduction on $DQBC$. $\mathbf{g}(\cdot)$ is a CNN to generate the motion fields.

The preliminary motion fields are $1/8$ in spatial resolution of I_0 and I_1 , just the same as $DQBC$.

Motion Refinement

MRM up-samples the motion fields to full size and compensates more details with the warped context features.

Context features are extracted by ContextNet which contains three down-sampling convolutional blocks. It takes I_τ as input and produces the context feature C_τ^i ($i \in \{0, 1, 2\}$) in $1/2^{3-i}$ spatial resolution of I_τ .

MRM consists of three UpBlocks each of which up-samples the input motion fields by factor 2. The first UpBlock takes the motion fields produced by Equation (6) and context features C_0^0, C_1^0 as inputs and produces the up-sampled motion fields $\tilde{M}_{t \rightarrow 0}^1, \tilde{M}_{t \rightarrow 1}^1$ and a hidden feature H_1 as:

$$\tilde{M}_{t \rightarrow 0}^1, \tilde{M}_{t \rightarrow 1}^1, H_1 = \text{UpBlock}_1(\tilde{M}_{t \rightarrow 0}, \tilde{M}_{t \rightarrow 1}, C_0^0, C_1^0). \quad (7)$$

The second and third UpBlock are similar to the first UpBlock but also take the hidden feature produced from the previous UpBlock as input:

$$\tilde{M}_{t \rightarrow 0}^{i+1}, \tilde{M}_{t \rightarrow 1}^{i+1}, H_{i+1} = \text{UpBlock}_{i+1}(\tilde{M}_{t \rightarrow 0}^i, \tilde{M}_{t \rightarrow 1}^i, C_0^i, C_1^i, H_i), \quad i \in \{1, 2\}. \quad (8)$$




Figure 4: Structure of the UpBlock in the MRM. An Upblock refines and up-samples the input motion fields by a factor of 2.

Figure 4 shows the structure of the $(i + 1)^{\text{th}}$ UpBlock ($i \in \{0, 1, 2\}$). H_i is surrounded by parentheses since there is no H_0 for the first UpBlock. \overleftarrow{C}_τ^i denotes the warped context feature $\overleftarrow{w}(C_\tau^i, \tilde{M}_{t \rightarrow \tau}^i)$. A convolutional head outputs residual motion fields ΔM^i and the up-sampling parameter w . The input motion fields are added with ΔM^i and then up-sampled with the parameter w . We up-sample the motion fields by taking the convex combination of neighboring 3×3 pixels on the input motion fields. Softmax is performed on w to get the convex combination weights that sum up to 1.

The occlusion map O is estimated from the last hidden feature H_3 using a convolutional head.

3.4 Frame Synthesis

We synthesize the interpolated frame following Equation (1). As shown in Figure 2, a SynthNet is devised to estimate a residual occlusion map ΔO and a residual image R for better interpolation quality. The residual image helps to compensate for details and to eliminate artifacts.

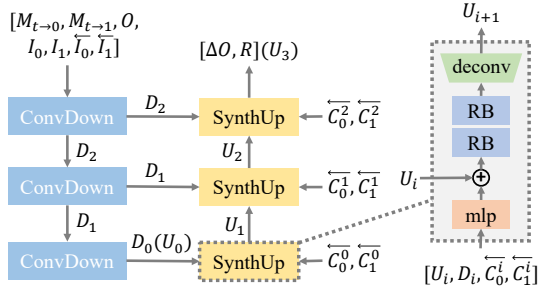


Figure 5: Structure of SynthNet. The Residual Block (RB) consists of 2 convolutional layers and a skip connection.

Our SynthNet is a modified U-Net as shown in Figure 5. The down-sampling path consists of three ConvDown blocks each of which has two convolutional layers and the first one has stride 2. The up-sampling path consists of three SynthUp blocks that take the warped context features shared by the MRM as auxiliary inputs and fuse them with hidden features.

3.5 Training

Loss Functions

Our model is trained end-to-end with three loss terms: reconstruction loss L_{rec} , teacher reconstruction loss L_{rec}^{tea} and distillation loss $L_{distill}$. The overall loss L is:

$$L = L_{rec} + \lambda_1 L_{rec}^{tea} + \lambda_2 L_{distill}. \quad (9)$$

L_{rec} is the L_1 loss between the estimated intermediate frame \hat{I}_t and the ground truth I_t^{gt} :

$$L_{rec} = \|\hat{I}_t - I_t^{gt}\|_1. \quad (10)$$

Following [Huang *et al.*, 2022], we adopt a privileged teacher to better supervise the learning of motion estimation. Specifically, during the training process an auxiliary IFBlock [Huang *et al.*, 2022] which is accessible to I_t^{gt} takes our estimated motion fields $M_{t \rightarrow 0}, M_{t \rightarrow 1}$ as preliminary motion fields and produces the revised ones $M_{t \rightarrow 0}^{tea}, M_{t \rightarrow 1}^{tea}$. The teacher reconstruction loss is defined as:

$$L_{rec}^{tea} = \|O \cdot \overleftarrow{w}(I_0, M_{t \rightarrow 0}^{tea}) + (1 - O) \cdot \overleftarrow{w}(I_1, M_{t \rightarrow 1}^{tea}) - I_t^{gt}\|_1, \quad (11)$$

where O is the occlusion map produced by our model.

The distillation loss is defined as the weighted sum of L_2 losses between $M_{t \rightarrow 0}^{tea}, M_{t \rightarrow 1}^{tea}$ and all the in-process motion fields that our model produces during the motion generation and refinement process.

Datasets

We use the training set of the Vimeo90K dataset proposed by [Xue *et al.*, 2019] as our training set. The training dataset contains 51,312 triplets with a resolution of 448×256 .

4 Experiments

4.1 Implementation Details

Network Architecture

The down-sampling levels of key sets when extracting $DQBC$ and the radii (r_l in Equation (4)) of local windows of each resolution level typically determine the preferred motion range that the network copes with. We use three levels of resolution. The radii for each level are 6, 5, and 4, respectively.

Training Details

Our network is optimized by AdamW [Loshchilov and Hutter, 2017] with weight decay 10^{-4} on 256×256 patches. The patches are randomly cropped from the training set and are randomly flipped, rotated and reversed. We use a total batch size of 64 distributed over 4 Tesla V100 GPUs for 510 epochs. The training process lasts for about three days. The memory consumption for each GPU is about 18 GB. The learning rate gradually decays from 2×10^{-4} to 2×10^{-6} using cosine annealing during the training process.

4.2 Evaluation Datasets

Vimeo90K. The evaluation set of Vimeo90K dataset [Xue *et al.*, 2019] contains 3,782 triplets with resolution 448×256 .

UCF101. The UCF101 dataset was originally proposed by [Soomro *et al.*, 2012] and contains human action videos of resolution 256×256 . [Liu *et al.*, 2017] selected 379 triplets for VFI evaluation.

Middlebury. Proposed by [Baker *et al.*, 2011], the Middlebury benchmark provides two sets: *Other* and *Evaluation*. We adopt the *Other* set which has ground truth intermediate frames and contains 12 triplets for VFI evaluation.

Method	Vimeo90K	UCF101	Middle-bury	SNU-FILM				Params (M)	MACs (G)	Time (ms)
				Easy	Medium	Hard	Extreme			
ToFlow	33.73/0.9682	34.58/0.9667	2.15	39.08/0.9890	34.39/0.9740	28.44/0.9180	23.39/0.8310	-	-	-
SepConv	33.79/0.9702	34.78/0.9669	2.27	39.41/0.9900	34.97/0.9762	29.36/0.9253	24.31/0.8448	-	-	-
CyclicGen	32.09/0.9490	35.11/0.9684	-	37.72/0.9840	32.47/0.9554	26.95/0.8871	22.70/0.8083	-	-	-
DAIN	34.71/0.9756	34.99/0.9683	2.04	39.73/0.9902	35.46/0.9780	30.17/0.9335	25.09/0.8584	-	-	-
CAIN	34.65/0.9730	34.91/0.9690	2.28	39.89/0.9900	35.61/0.9776	29.90/0.9292	24.78/0.8507	42.8	43.5	62
AdaCoF	34.47/0.9730	34.90/0.9680	2.24	39.80/0.9900	35.05/0.9753	29.46/0.9244	24.31/0.8439	21.8	-	15
BMBC	35.01/0.9764	35.15/0.9689	2.04	39.90/0.9902	35.31/0.9774	29.33/0.9270	23.92/0.8432	11.0	174.2	2168
RIFE	35.61/0.9779	35.29/0.9690	1.96	40.02/0.9904	35.72/0.9786	30.07/0.9326	24.82/0.8529	10.1	11.7	31
ABME	36.18/0.9805	35.38/0.9698	2.01	39.59/0.9901	35.77/0.9789	30.58/0.9363	25.42/0.8639	17.5	-	235
VFIformer	<u>36.50/0.9816</u>	<u>35.43/0.9700</u>	<u>1.82</u>	<u>40.13/0.9907</u>	<u>36.09/0.9799</u>	<u>30.67/0.9378</u>	<u>25.43/0.8643</u>	24.2	356.7	756
Ours	36.37/0.9812	35.35/0.9696	1.86	<u>40.15/0.9907</u>	<u>36.10/0.9796</u>	<u>30.78/0.9371</u>	25.41/0.8628	18.3	57.5	54
Ours-Aug	36.57/0.9817	35.44/0.9700	1.78	40.31/0.9909	36.25/0.9799	30.94/0.9378	25.61/0.8648	18.3	229.9	206

Table 1: Quantitative comparisons with state-of-the-art VFI methods. The average interpolation error IE (the lower, the better) is reported on the Middlebury benchmark and PSNR/SSIM (the higher, the better) are reported on the others. The **best** and second-best results for accuracy are emphasized with different styles.

SNU-FILM. Proposed with [Choi *et al.*, 2020], the SNU-FILM benchmark contains four splits of different difficulty for VFI: *Easy*, *Medium*, *Hard* and *Extreme*. The difficulty is categorized by motion amplitude. Each split contains 310 triplets of resolution up to 1280×720 .

4.3 Comparisons with the State-of-the-Art

We quantitatively compared our model with some other competitive state-of-the-art VFI models, including ToFlow [Xue *et al.*, 2019], SepConv [Niklaus *et al.*, 2017b], CyclicGen [Liu *et al.*, 2019], DAIN [Bao *et al.*, 2019], CAIN [Choi *et al.*, 2020], AdaCoF [Lee *et al.*, 2020], BMBC [Park *et al.*, 2020], RIFE [Huang *et al.*, 2022], ABME [Park *et al.*, 2021] and VFIformer [Lu *et al.*, 2022]. Results are presented in Table 1.

Besides the accuracy metrics (IE/PSNR/SSIM), we also provide Table 1 with numbers of model parameters, MACs and inference time for some methods. MACs are calculated on 256×256 patches using *thop*¹. The inference time is recorded as the average inference time per sample on the Vimeo90K evaluation set and is tested on the same GTX-1080-Ti GPU.

Considering that our base model (“**Ours**” in Table 1) runs much faster than VFIformer both theoretically (indicated by MACs) and practically (indicated by inference time), we use test augmentation to boost the performance at the price of more computation. We combine temporal reversal and 90° spatial rotation to augment the test frame pair to 4 groups and directly average the 4 outcomes to get the final interpolation result. The augmented version is denoted as “**Ours-Aug**” in Table 1.

Compared with ABME, our base model achieves **+0.19 dB** improvement and runs in less than **1/4** inference time on the Vimeo90K evaluation set. Compared with VFIformer, our augmented model achieves **+0.07 dB** improvement and runs in less than **1/3** inference time on the Vimeo90K evaluation set.

As illustrated by Figure 6, we visualized the motion fields estimated by VFIformer, ABME and our base model. It can

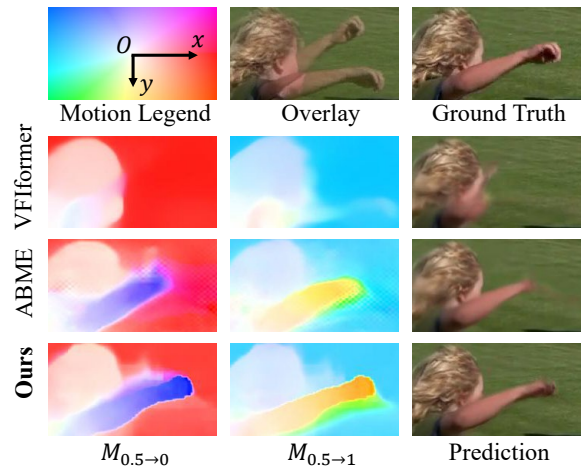


Figure 6: Visualizations of motion fields that are estimated by VFIformer, ABME and DQBC (Ours). Colors represent the directions of motion vectors in motion fields.

be seen that our model is more capable of coping with small and fast-moving objects (the waving arm and fist in this case). The waving arm and fist are totally missed in the motion estimation of VFIformer, leading to a poor interpolation result. ABME succeeds in capturing the motion of the waving arm but misses the fist which is smaller and moves faster.

We also present some visual comparisons of the interpolated results as shown in Figure 7. Both the quantitative comparisons and interpolation visualizations show that our approach is pretty competitive.

4.4 Ablation Study

For the ablation experiments below, we report the results on Vimeo90K evaluation set at 300 training epochs.

Ablations for DQBC

To verify the effectiveness of DQBC, we compared our model with the baseline model in which we adopt the bilateral motion estimation module of [Park *et al.*, 2021] as the motion

¹<https://pypi.org/project/thop/>




Figure 7: Visualizations of interpolated results from some competitive approaches and ours.

generation module. The motion refinement module and SynthNet stay the same. Basically, the baseline model adopts the paradigm illustrated by the top half of Figure 1. Table 2 shows that DQBC outperforms the baseline by a notable margin. Besides, it can be seen that the correlation enhancement block and the feature distributing operation also contribute to performance promotion. These results verified the effectiveness of the proposed DQBC.

Model	PSNR	SSIM
Baseline	36.16	0.9804
Ours	36.31	0.9810
Ours w/o enh	36.25	0.9809
Ours w/o dist	36.28	0.9809

Table 2: Ablation experiments for DQBC. “enh” represents the correlation enhancement block and “dist” represents the feature distributing operation.

Model	PSNR	SSIM
Bilinear	36.25	0.9808
Convex w/o refine	36.11	0.9802
Convex (Ours)	36.31	0.9810

Table 3: Contrastive experiments for alternative designs for Motion Refinement Module (MRM).

Ablations for Motion Refinement Module

The Motion Refinement Module (MRM) refines motion fields by producing residuals and up-samples the motion fields by estimating convex combination weights. We compared our approach with the naive bilinear way, where we first bilinearly up-sample the coarse motion fields and then estimate residuals to refine the up-sampled motion fields. We also explored the effectiveness of the refinement function of MRM by removing the residual estimation from it. Table 3 shows that our approach is better than the naive bilinear approach. The “convex” approach only works well along with the residual estimation, indicating that it is necessary for MRM to have the ability to revise the motion fields rather than merely up-sample them.

5 Conclusion

We have proposed the Densely Queried Bilateral Correlation (DQBC) for better modeling correlations between two input frames in the VFI scenario. DQBC gets rid of the receptive field dependency problem in common approaches and thus is more friendly to small and fast-moving objects. A Motion Refinement Module (MRM) has also been devised to up-sample the motion fields to full size and refine them with context information. We have verified the effectiveness of DQBC and MRM via ablation studies. Both quantitative experiments and interpolation visualizations show that our approach has reached and outperformed the state-of-the-art level on various VFI benchmarks.

Acknowledgements

This work is supported by the Fundamental Research Funds for the Central Universities, No.40020214380099.

References

- [Baker *et al.*, 2011] Simon Baker, Daniel Scharstein, JP Lewis, Stefan Roth, Michael J Black, and Richard Szeliski. A database and evaluation methodology for optical flow. *International journal of computer vision*, 92(1):1–31, 2011.
- [Bao *et al.*, 2019] Wenbo Bao, Wei-Sheng Lai, Chao Ma, Xiaoyun Zhang, Zhiyong Gao, and Ming-Hsuan Yang. Depth-aware video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3703–3712, 2019.
- [Chen and Zwicker, 2022] Shuhong Chen and Matthias Zwicker. Improving the perceptual quality of 2d animation interpolation. In *European Conference on Computer Vision*, pages 271–287. Springer, 2022.
- [Cheng and Chen, 2020] Xianhang Cheng and Zhenzhong Chen. Video frame interpolation via deformable separable convolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10607–10614, 2020.
- [Cheng and Chen, 2021] Xianhang Cheng and Zhenzhong Chen. Multiple video frame interpolation via enhanced deformable separable convolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [Chi *et al.*, 2020] Zhixiang Chi, Rasoul Mohammadi Nasiri, Zheng Liu, Juwei Lu, Jin Tang, and Konstantinos N Plataniotis. All at once: Temporally adaptive multi-frame interpolation with advanced motion modeling. In *European Conference on Computer Vision*, pages 107–123. Springer, 2020.
- [Choi *et al.*, 2020] Myungsub Choi, Heewon Kim, Bohyung Han, Ning Xu, and Kyoung Mu Lee. Channel attention is all you need for video frame interpolation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10663–10671, 2020.
- [Choi *et al.*, 2021] Myungsub Choi, Suyoung Lee, Heewon Kim, and Kyoung Mu Lee. Motion-aware dynamic architecture for efficient frame interpolation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13839–13848, 2021.
- [Danier *et al.*, 2022] Duolikun Danier, Fan Zhang, and David Bull. St-mfnet: A spatio-temporal multi-flow network for frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3521–3531, 2022.
- [Haris *et al.*, 2020] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. Space-time-aware multi-resolution video enhancement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2859–2868, 2020.
- [Huang *et al.*, 2022] Zhewei Huang, Tianyuan Zhang, Wen Heng, Boxin Shi, and Shuchang Zhou. Real-time intermediate flow estimation for video frame interpolation. In *European Conference on Computer Vision*, pages 624–642. Springer, 2022.
- [Jiang *et al.*, 2018] Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik Learned-Miller, and Jan Kautz. Super slomo: High quality estimation of multiple intermediate frames for video interpolation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9000–9008, 2018.
- [Kalantari *et al.*, 2016] Nima Khademi Kalantari, Ting-Chun Wang, and Ravi Ramamoorthi. Learning-based view synthesis for light field cameras. *ACM Transactions on Graphics (TOG)*, 35(6):1–10, 2016.
- [Kim *et al.*, 2020] Soo Ye Kim, Jihyong Oh, and Munchurl Kim. Fsr: Deep joint frame interpolation and super-resolution with a multi-scale temporal loss. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11278–11286, 2020.
- [Kong *et al.*, 2022] Lingtong Kong, Boyuan Jiang, Donghao Luo, Wenqing Chu, Xiaoming Huang, Ying Tai, Chengjie Wang, and Jie Yang. Ifrnet: Intermediate feature refine network for efficient frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1969–1978, 2022.
- [Lee *et al.*, 2020] Hyeongmin Lee, Taoh Kim, Tae-young Chung, Daehyun Pak, Yuseok Ban, and Sangyoun Lee. Adacof: Adaptive collaboration of flows for video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5316–5325, 2020.
- [Liu *et al.*, 2017] Ziwei Liu, Raymond A Yeh, Xiaoou Tang, Yiming Liu, and Aseem Agarwala. Video frame synthesis using deep voxel flow. In *Proceedings of the IEEE international conference on computer vision*, pages 4463–4471, 2017.
- [Liu *et al.*, 2019] Yu-Lun Liu, Yi-Tung Liao, Yen-Yu Lin, and Yung-Yu Chuang. Deep video frame interpolation using cyclic frame generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8794–8802, 2019.
- [Liu *et al.*, 2020] Yihao Liu, Liangbin Xie, Li Siyao, Wenxiu Sun, Yu Qiao, and Chao Dong. Enhanced quadratic video interpolation. In *European Conference on Computer Vision*, pages 41–56. Springer, 2020.
- [Loshchilov and Hutter, 2017] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [Lu *et al.*, 2022] Liying Lu, Ruizheng Wu, Huajia Lin, Jiangbo Lu, and Jiaya Jia. Video frame interpolation with transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3532–3542, 2022.

- [Meyer *et al.*, 2015] Simone Meyer, Oliver Wang, Henning Zimmer, Max Grosse, and Alexander Sorkine-Hornung. Phase-based frame interpolation for video. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1410–1418, 2015.
- [Meyer *et al.*, 2018] Simone Meyer, Abdelaziz Djelouah, Brian McWilliams, Alexander Sorkine-Hornung, Markus Gross, and Christopher Schroers. Phasenet for video frame interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 498–507, 2018.
- [Niklaus and Liu, 2018] Simon Niklaus and Feng Liu. Context-aware synthesis for video frame interpolation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1701–1710, 2018.
- [Niklaus *et al.*, 2017a] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive convolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 670–679, 2017.
- [Niklaus *et al.*, 2017b] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive separable convolution. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 261–270, 2017.
- [Park *et al.*, 2020] Junheum Park, Keunsoo Ko, Chul Lee, and Chang-Su Kim. Bmbc: Bilateral motion estimation with bilateral cost volume for video interpolation. In *European Conference on Computer Vision*, pages 109–125. Springer, 2020.
- [Park *et al.*, 2021] Junheum Park, Chul Lee, and Chang-Su Kim. Asymmetric bilateral motion estimation for video frame interpolation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14539–14548, 2021.
- [Peleg *et al.*, 2019] Tomer Peleg, Pablo Szekely, Doron Sabo, and Omry Sendik. Im-net for high resolution video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2398–2407, 2019.
- [Reda *et al.*, 2022] Fitsum Reda, Janne Kontkanen, Eric Tabellion, Deqing Sun, Caroline Pantofaru, and Brian Curless. Film: Frame interpolation for large motion. In *European Conference on Computer Vision*, pages 250–266. Springer, 2022.
- [Shen *et al.*, 2022] Wang Shen, Cheng Ming, Wenbo Bao, Guangtao Zhai, Li Chenn, and Zhiyong Gao. Enhanced deep animation video interpolation. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 31–35. IEEE, 2022.
- [Shi *et al.*, 2022] Zhihao Shi, Xiangyu Xu, Xiaohong Liu, Jun Chen, and Ming-Hsuan Yang. Video frame interpolation transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17482–17491, 2022.
- [Sim *et al.*, 2021] Hyeonjun Sim, Jihyong Oh, and Munchurl Kim. Xvfi: Extreme video frame interpolation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14489–14498, 2021.
- [Siyao *et al.*, 2021] Li Siyao, Shiyu Zhao, Weijiang Yu, Wenxiu Sun, Dimitris Metaxas, Chen Change Loy, and Ziwei Liu. Deep animation video interpolation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6587–6595, 2021.
- [Soomro *et al.*, 2012] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- [Tulyakov *et al.*, 2021] Stepan Tulyakov, Daniel Gehrig, Stamatiou Georgoulis, Julius Erbach, Mathias Gehrig, Yuanyou Li, and Davide Scaramuzza. Time lens: Event-based video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16155–16164, 2021.
- [Wu *et al.*, 2018] Chao-Yuan Wu, Nayan Singhal, and Philipp Krahenbuhl. Video compression through image interpolation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 416–431, 2018.
- [Wu *et al.*, 2022] Yue Wu, Qiang Wen, and Qifeng Chen. Optimizing video prediction via video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17814–17823, 2022.
- [Xiang *et al.*, 2020] Xiaoyu Xiang, Yapeng Tian, Yulun Zhang, Yun Fu, Jan P Allebach, and Chenliang Xu. Zooming slow-mo: Fast and accurate one-stage space-time video super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3370–3379, 2020.
- [Xu *et al.*, 2019] Xiangyu Xu, Li Siyao, Wenxiu Sun, Qian Yin, and Ming-Hsuan Yang. Quadratic video interpolation. *Advances in Neural Information Processing Systems*, 32, 2019.
- [Xue *et al.*, 2019] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision*, 127(8):1106–1125, 2019.
- [Yu *et al.*, 2022] Zhiyang Yu, Yu Zhang, Xujie Xiang, Dongqing Zou, Xijun Chen, and Jimmy S Ren. Deep bayesian video frame interpolation. In *European Conference on Computer Vision*, pages 144–160. Springer, 2022.
- [Zhou *et al.*, 2016] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A Efros. View synthesis by appearance flow. In *European conference on computer vision*, pages 286–301. Springer, 2016.