

An Efficient Quality Metric for Video Frame Interpolation Based on Motion-Field Divergence

Conall Daly, Darren Ramsook, and Anil Kokaram

Sigmedia Group, School of Engineering

Trinity College Dublin

Ireland

{dalyc21, ramsookd, anil.kokaram}@tcd.ie

Abstract—Video frame interpolation is a fundamental tool for temporal video enhancement, but existing quality metrics struggle to evaluate the perceptual impact of interpolation artefacts effectively. Metrics like PSNR, SSIM and LPIPS ignore temporal coherence. State-of-the-art quality metrics tailored towards video frame interpolation, like FloLPIPS, have been developed but suffer from computational inefficiency that limits their practical application. We present PSNR_{DIV}, a novel full-reference quality metric that enhances PSNR through motion divergence weighting, a technique adapted from archival film restoration where it was developed to detect temporal inconsistencies. Our approach highlights singularities in motion fields which is then used to weight image errors. Evaluation on the BVI-VFI dataset (180 sequences across multiple frame rates, resolutions and interpolation methods) shows PSNR_{DIV} achieves statistically significant improvements: +0.09 Pearson Linear Correlation Coefficient over FloLPIPS, while being 2.5× faster and using 4× less memory. Performance remains consistent across all content categories and are robust to the motion estimator used. The efficiency and accuracy of PSNR_{DIV} enables fast quality evaluation and practical use as a loss function for training neural networks for video frame interpolation tasks. An implementation of our metric is available at www.github.com/conalld/psnr-div.

Index Terms—Video frame interpolation quality, temporal consistency metrics.

I. INTRODUCTION

Video frame interpolation (VFI) is a widely studied problem in video processing, aimed at generating intermediate frames between existing ones in a sequence. By increasing the effective frame rate, VFI enhances visual quality and enables applications such as slow-motion effects [1], efficient rendering of animated content [2], and broadcast frame-rate conversion [3]. In recent years, significant advancements have been made in VFI techniques, driven by sophisticated motion estimation algorithms [4] and deep neural networks (DNN). DNN based VFI algorithms are capable of accurately predicting intricate, non-linear motion patterns. They also handle occlusions more effectively than classical methods. Techniques used for DNN-based VFI have included motion-based [5], kernel-based [6], [7], and transformer networks [8]. Despite these significant ad-

vancements in interpolation quality, evaluating and comparing VFI methods remains challenging.

Traditional metrics such as PSNR and SSIM are commonly applied to video quality assessment by averaging frame-level scores. However, this approach fails to account for temporal information, leading to poor correlation with perceptual judgments in video frame interpolation (VFI). While perceptual metrics like LPIPS [9] improve upon traditional methods by leveraging deep learning to compare structural and textural differences, they still operate on a per-frame basis. As a result, they overlook motion consistency which is a critical aspect of perceived video quality [10]. This shortcoming is particularly evident in VFI, where temporal artifacts significantly impact subjective evaluations [11], as well as in diffusion-based video generation, where maintaining frame coherence remains challenging [12].

To address this limitation, recent adaptations have incorporated temporal awareness into quality assessment. For instance, FloLPIPS [13] enhances LPIPS by weighting motion errors to highlight motion-related distortions. While this improves detection of temporal artifacts, its high computational cost (reported as 5.5× slower than standard LPIPS) makes it impractical for training DNNs.

To address these limitations we propose a novel motion-focused quality metric for VFI evaluation that is simpler and correlates more closely to subjective scores than FloLPIPS. Our key contributions are as follows.

- 1) We propose a metric which takes a novel approach to weighting image errors based on blotch detection work in the field of archival footage restoration.
- 2) This metric only requires the motion field for one of the compared video sequences resulting in a reduction in computational load.
- 3) We demonstrate that our metric can outperform FloLPIPS with better correlation to subjective scores.

The rest of the paper is structured as follows: Section II reviews background of quality metrics for VFI. Section III details our proposed VFI quality metric. Section IV details the BVI-VFI subjective dataset we use to evaluate our metric and the other metrics we compare against. Section V discusses the results of our evaluation and looks at the robustness of our metric to motion fields generated by different estimation algorithms.

This publication has emanated from research jointly funded by Taighde Éireann – Research Ireland, and Overcast HQ under Grant number EP-SPG/2023/1515. Thanks also to Neil Birkbeck and Balu Adsumilli from the YouTube Video Algorithms team for their helpful comments and supporting this work.

II. BACKGROUND

This section reviews previous works on quality measurement for VFI and how temporal consistency measures can be leveraged as a video quality measure for VFI.

A. Video Frame Interpolation Quality Metrics

Video quality metrics are categorized as no-reference (NR), reduced-reference (RR), or full-reference (FR). Our work focuses on FR metrics, which compare interpolated sequences against a ground-truth reference. Metric performance is typically evaluated by correlation with human subjective scores, using statistical measures such as Pearson’s Linear Correlation Coefficient (PLCC), Spearman’s Rank Correlation Coefficient (SRCC), Kendall’s Rank Correlation Coefficient (KRCC) and Root Mean Squared Error (RMSE), as standardized by VQEG [14].

Recently, video quality metrics increasingly leverage spatio-temporal features [15], [16]. ST-GREED [17] analyses spatial/temporal band-pass coefficients via Generalized Gaussian Distributions (GGD), measuring entropy differences between reference and interpolated videos to detect distortions from frame rate changes and compression. In contrast, Wu *et al.* [18] directly use estimated motion, tracking keypoint trajectories and extracting motion field patches around them. By comparing histograms of these patches between reference and interpolated sequences, they quantify temporal consistency deviations.

However, general-purpose video metrics often correlate poorly with subjective scores for VFI-specific artefacts [19], prompting the development of bespoke solutions. Hou *et al.* [20] employ a Swin Transformer-based architecture trained directly on subjective scores from their own custom dataset (VFIPS). It achieves good correlation in terms of PLCC (0.87) and SRCC (0.79). However, it is computationally intensive, requiring 2.5 GB of GPU memory to process a 256×256 12-frame sequence.

Danier *et al.* propose FloLPIPS, a video quality metric that extends LPIPS by incorporating temporal distortion sensitivity. LPIPS computes perceptual differences between reference ($\mathbf{I}$) and interpolated ($\hat{\mathbf{I}}$) frames using deep feature maps $\phi_l(\hat{\mathbf{I}})$, $\phi_l(\mathbf{I})$ from layer l of a pre-trained convolutional neural network. For each spatial position $\mathbf{x}$, LPIPS calculates $\|w_l \odot (\phi_l(\hat{\mathbf{I}}) - \phi_l(\mathbf{I}))\|_2^2$ across channels, where w_l is a weighting to emphasise or de-emphasise the features of layer l . It then averages these differences spatially. FloLPIPS enhances LPIPS by introducing motion error weighting. For consecutive frames $\hat{\mathbf{I}}_{n-1}$, $\hat{\mathbf{I}}_n$ and $\mathbf{I}_{n-1}$, $\mathbf{I}_n$, it computes motion fields $\hat{\mathbf{F}}$ and $\mathbf{F}$ using a pre-trained estimator. The end-point-error (EPE) between these motion fields is calculated.

$$\Delta\mathbf{F}(\mathbf{x}) = \|\hat{\mathbf{F}}(\mathbf{x}) - \mathbf{F}(\mathbf{x})\|_2 \quad (1)$$

This EPE map is normalised spatially to produce distortion-sensitive weights.

$$\mathbf{w}(\mathbf{x}) = \frac{\Delta\mathbf{F}(\mathbf{x})}{\sum_{\mathbf{x}} \Delta\mathbf{F}(\mathbf{x})} \quad (2)$$

This weights LPIPS’s uniform spatial averaging, emphasizing regions with motion inconsistencies. The final FloLPIPS score aggregates the weighted feature differences across layers and frames.

$$\text{FloLPIPS} = \sum_l \sum_{\mathbf{x}} \mathbf{w}(\mathbf{x}) \|w_l \odot (\phi_l(\hat{\mathbf{I}}) - \phi_l(\mathbf{I}))\|_2^2 \quad (3)$$

FloLPIPS achieves scores of 0.71 in terms of PLCC and 0.68 in terms of SRCC on the BVI-VFI dataset. As mentioned earlier, FloLPIPS is quite computationally intensive. FloLPIPS runtime is reported as 332 ms for a 1080p frame whereas LPIPS is 59 ms, making FloLPIPS 5.5× slower. FloLPIPS power comes from using EPE as a measure of temporal consistency. As such, we propose to use a simpler more computationally efficient measure of temporal consistency.

B. Divergence as a Measure of Temporal Consistency

Temporal consistency is critical for perceived video quality [21], distortions causing motion discontinuities are of particular interest for VFI. Our metric adapts techniques from archival film restoration, where pathological motion (PM) [22] characterised by irregular, non-smooth motion like fast deformations or intermittent movements pose a key challenge. PM introduces intensity discontinuities resembling physical film artefacts (scratches/blotches), leading to false alarms for traditional displaced frame difference (DFD) detectors [23]. These methods often misclassify rapid motion or complex textures as artefacts, because their motion field patterns mimic true damage.

Corrigan *et al.* [24] addressed this by introducing motion field divergence to spatially identify PM. Divergence quantifies local motion irregularities: high absolute values indicate abrupt vector changes, violating smooth motion assumptions. By integrating divergence over a five-frame window, their method distinguishes PM from true artefacts, even when temporal profiles overlap, reducing false alarms while preserving detection in stable regions. This spatial approach complements DFD methods, robustly flagging PM.

We repurpose divergence as a computationally efficient metric for temporal consistency. Unlike EPE, which requires two motion fields for comparison, divergence can highlight motion inconsistencies within a single motion field. This approach is particularly suited to VFI, where interpolation artefacts often manifest as non-smooth motion fields.

III. OUR PROPOSED METRIC

A flow diagram for our proposed VFI quality metric is shown in Figure 1. We define two video sequences. An interpolated sequence $\hat{\mathbf{G}}$ which has been interpolated from a low frame rate to a higher frame rate, and a reference sequence $\mathbf{G}$ which is the ground truth target high frame rate video. For a current time n , we have an interpolated frame, $\hat{\mathbf{I}}_n \in \hat{\mathbf{G}}$, and an interpolated frame $\hat{\mathbf{I}}_{n+1} \in \hat{\mathbf{G}}$ at the successive future time $n+1$. The motion field between these two frames $\hat{\mathbf{F}}_n = [\hat{\mathbf{u}}(\mathbf{x}), \hat{\mathbf{v}}(\mathbf{x})]$ is computed. The divergence of this motion field can be calculated from its two components as follows.




Fig. 1. Flow diagram for our proposed metric. Inputs and outputs are annotated above flow arrows. Operations are denoted by blue boxes. Our metric uses the divergence of the motion field for weighting the mean squared error between pixel values of the interpolated and reference sequences.

$$\nabla \cdot \hat{\mathbf{F}}_n(\mathbf{x}) = \frac{\partial \hat{u}(\mathbf{x})}{\partial x} + \frac{\partial \hat{v}(\mathbf{x})}{\partial y} \quad (4)$$

In practice the absolute value of $\nabla \cdot \hat{\mathbf{F}}_n$ is taken as we simply want to detect motion field singularities. This divergence magnitude is normalised between $[0, 1]$ by dividing by the maximum value for our motion field at time n .

$$\mathbf{d}(\mathbf{x}) = \frac{|\nabla \cdot \hat{\mathbf{F}}_n(\mathbf{x})|}{\max(|\nabla \cdot \hat{\mathbf{F}}_n(\mathbf{x})|)} \quad (5)$$

Using this normalised value $\mathbf{d}(\mathbf{x})$ allows us to more consistently assess importance of local motion field singularities. Essentially regions where $\mathbf{d}(\mathbf{x}) \approx 0$ can be taken to be temporally smooth. This can be used as a good proxy to indicate regions of a frame where humans may notice motion errors. We use a simple binary mask, $\mathbf{w}(\mathbf{x})$, based on a pre-defined threshold of $\mathbf{d}(\mathbf{x})$, T (we use $T = 0.01$, based on our parameter optimisation in Section V-B), to segment our motion field into regions of high divergence and low divergence.

$$\mathbf{w}(\mathbf{x}) = \begin{cases} 0 & \mathbf{d}(\mathbf{x}) \leq T \\ 1 & \text{Otherwise} \end{cases} \quad (6)$$

We then take the squared error between the interpolated frame $\hat{\mathbf{I}}_n$ and the corresponding reference frame $\mathbf{I}_n \in \mathbf{G}$. This binary mask along with the squared error is used to calculate a weighted mean squared error (MSE), referred to as divergence masked MSE (MSE_w). The normalising factor for the MSE is based on the sum of the number of pixels where $\mathbf{d}(\mathbf{x}) \geq T$, denoted as $Z = \sum_{\mathbf{x}} \mathbf{w}(\mathbf{x})$. This results in only calculating image errors for regions of temporal inconsistency.

$$\text{MSE}_w = \frac{1}{Z} \sum_{\mathbf{x}} \mathbf{w}(\mathbf{x}) (\mathbf{I}_n(\mathbf{x}) - \hat{\mathbf{I}}_n(\mathbf{x}))^2 \quad (7)$$

This modified MSE can simply be used in the conventional expression for PSNR, resulting in our VFI quality metric PSNR_{DIV} .

$$\text{PSNR}_{\text{DIV}} = 20 \cdot \log_{10} \left(\frac{255}{\sqrt{\text{MSE}_w}} \right) \quad (8)$$

Concisely, our metric is a modified PSNR that uses MSE weighted by motion divergence, which is a measure of temporal inconsistency.

IV. EXPERIMENTAL SETUP

In this section, we describe the dataset and statistical measures used to evaluate and compare the performance of our proposed metric.

A. Evaluation Dataset

We evaluate the proposed method using the BVI-VFI database [25], the only publicly available subjective quality dataset containing uncompressed video sequences, in a 8-bit YUV420 format, with distortions exclusively induced by video frame interpolation (VFI). Compared to the two other most commonly used VFI quality datasets, VFIPS [20] and VFIIQA [15], it is the only one that reports Differential Mean Opinion Scores (DMOS) making it the most suitable for FR metric assessment. The full BVI-VFI dataset spans three spatial resolutions (540p, 1080p, 2160p) and three frame rates (30, 60, 120fps), offering diverse content to study perceptual artefacts. For this work, we focus on a subset comprising 540p and 1080p resolutions at 30 and 60fps, as these are more common resolution and framerate conversion scenarios in addition to being more challenging for quality metrics.

The subset includes 24 source sequences, each processed by five representative VFI methods at two different frame rates (in total this gives 240 scored sequences): basic frame averaging/repeating (to simulate motion blur and judder),

motion-based deep learning approaches (DVF [5] and QVI [26]), and the state-of-the-art kernel-based ST-MFNet [7], which combines deformable convolutions with texture synthesis based post-processing. These methods collectively span classical and modern VFI algorithms, enabling comprehensive analysis of artefact types. For each interpolated sequence BVI-VFI provides DMOS derived from a large-scale subjective study involving 189 participants, which quantify perceptual quality degradation relative to a reference sequence. DMOS is obtained by firstly computing the differential opinion score d_{ij} between a subject i 's score $\hat{s}_{ij}$ for an interpolated video j and the score for its corresponding reference video s_{ij} assigned by the same participant.

$$d_{ij} = s_{ij} - \hat{s}_{ij} \quad (9)$$

The DMOS for each interpolated video is obtained by averaging the differential opinion scores for each of the N_j participants. Each video in the BVI-VFI dataset has been scored by at least 20 participants ($N_j \geq 20$).

$$\text{DMOS}_j = \frac{1}{N_j} \sum_{i=1}^{N_j} d_{ij} \quad (10)$$

BVI-VFI provides these conventional DMOS values along with ones that have been post-processed to P.910 specifications [27], which accounts for biases in the scoring of participants and re-weights the differential opinion scores based on this. We use these de-biased DMOS values to evaluate our metric.

B. Compared Quality Metrics

We benchmark our metric against widely adopted quality metrics and specialized VFI quality assessors. Notably few FR metrics have been developed since FloLPIPS. Recently there has been a resurgence in VFI quality metric development [15], [16]. As of writing, weights are yet to be released for these models, however we compare against the same metrics evaluated in those papers which all happen to be pre-2022. PSNR and SSIM are included as foundational baselines, measuring pixel-wise fidelity and structural similarity, respectively. For PSNR and PSNR_{DIV}, we compute scores solely on the luminance channel (Y). Gradient Magnitude Similarity Deviation (GMSD) [28] is included, leveraging spatial gradient deviations to capture edge and texture degradation. We include LPIPS, a deep perceptual metric that compares activated network features between reference and interpolated frames. Although these image metrics lack temporal modelling, we include them to maintain consistency with prior VFI literature. We use FloLPIPS as a representative VFI video quality metric that incorporates temporal information. FloLPIPS is used with AlexNet [29] as its feature extractor (we also use LPIPS with AlexNet) and PWC-Net as its motion estimator [30].

As an ablation, we adapt FloLPIPS's motion error weighting scheme (2) to MSE, creating a baseline (PSNR_{EPE}) where motion field EPE weights pixel-wise errors. Which is normalised by the sum of the differences between the motion fields. We

take this as the weighting function for MSE and refer to it as PSNR_{EPE}.

C. Evaluation Metrics

To evaluate the alignment between objective quality metrics and subjective judgments, we employ the four statistical measures previously mentioned in Section II-A which are PLCC, SRCC, KRCC, and RMSE.

We follow the VQEG recommendations for fitting quality metric values to perceptual scores [14]. A logistic function $Y(x)$ with four-parameter (β_{1-4}) is fitted to map metric values to DMOS scores using initial parameter estimates from the VQEG report.

$$Y(x) = \beta_2 + \frac{\beta_1 - \beta_2}{1 + \exp\left(-\frac{x - \beta_3}{|\beta_4|}\right)} \quad (11)$$

The optimization process combines exhaustive search for coarse tuning with the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm in SciPy¹ for refinement of the logistic fit parameters.

V. RESULTS AND DISCUSSION

In this section we quantitatively evaluate the performance of our metric by comparing it with six commonly used image/video quality metrics. We also examine the effect of changing the motion estimator, and changing the divergence threshold parameter on metric performance. Finally, we compare the computational complexity of our metric compared to FloLPIPS.

A. Evaluation

The evaluation of PSNR_{DIV} alongside other compared quality models on the BVI-VFI dataset are presented in Table I. For PSNR_{DIV} and PSNR_{EPE}, we compute motion fields using OpenCV's implementation² of the F rneback motion estimation algorithm [31]. It can be observed that PSNR_{DIV} outperforms the other 6 quality metrics according to all four performance measurements overall.

Compared to FloLPIPS, PSNR_{DIV} achieves a large gain in terms of PLCC and RMSE (+0.09 and -1.38 respectively). Small gains are also seen in terms of SRCC (+0.05). When broken down by framerate and resolution our metric outperforms FloLPIPS in all categories. GMSD occasionally outperforms PSNR_{DIV} in SRCC (e.g., for high-motion sequences), likely due to its sensitivity to edge distortions and ability to capture structural degradation.

PSNR_{DIV} is across the board better than PSNR_{EPE}. This demonstrates the effectiveness of our motion divergence based weighting scheme. It is worth noting that the the poor performance of LPIPS suggests that the motion weighting step plays a larger role than the perceptual part of LPIPS showing the power of using motion as a quality feature.

¹<https://docs.scipy.org/doc/scipy/reference/optimize.minimize-lbfgsb.html>

²https://docs.opencv.org/3.4/d4/dee/tutorial_optical_flow.html

TABLE I

THE PERFORMANCE OF EACH QUALITY ASSESSMENT MODEL IN PREDICTING DMOS ON THE BVI-VFI DATASET. FOR EACH STATISTICAL METRIC, THE BEST AND SECOND BEST RESULTS ARE **BOLDED** AND UNDERLINED RESPECTIVELY.

Metric	540p				1080p				Overall			
	15fps→30fps		30fps→60fps		15fps→30fps		30fps→60fps		PLCC		SRCC	
	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	KRCC	SRCC	RMSE
PSNR	0.41	0.27	0.35	0.45	0.40	0.45	0.61	0.66	0.42	0.32	0.46	16.70
SSIM	0.40	0.35	0.48	0.60	0.42	0.44	0.62	0.71	0.48	0.36	0.53	16.06
GMSD	0.48	0.41	<u>0.56</u>	0.67	0.56	0.58	0.68	0.73	<u>0.59</u>	<u>0.44</u>	0.63	<u>14.78</u>
LPIPS	0.38	0.29	0.44	0.52	0.53	0.45	0.66	0.66	0.51	0.36	0.52	15.83
FloLPIPS	<u>0.49</u>	<u>0.42</u>	0.53	0.55	<u>0.60</u>	0.55	<u>0.78</u>	<u>0.74</u>	0.58	0.40	<u>0.58</u>	14.98
PSNR _{EPE}	0.36	0.31	0.47	0.50	0.44	0.34	0.44	0.34	0.38	0.24	0.35	16.96
PSNR _{DIV}	0.52	0.45	0.65	<u>0.63</u>	0.61	<u>0.57</u>	0.80	0.76	0.67	0.45	0.63	13.60

TABLE II

BOOTSTRAPPED P-VALUES USING T-TEST COMPARING PSNR_{DIV} AND FLOLPIPS PERFORMANCE. $p < 0.05$ IS DEEMED STATISTICALLY SIGNIFICANT.

Statistical Measure	Subset (%)		
	30	50	70
PLCC	0.00	0.00	0.00
KRCC	0.78	0.11	0.01
SRCC	0.00	0.00	0.00
RMSE	0.00	0.00	0.00

To determine whether the performance difference between PSNR_{DIV} and FloLPIPS is statistically significant, we employ a bootstrapping method. We repeatedly sample (with replacement) subsets of sequences from the BVI-VFI dataset, including their PSNR_{DIV}, FloLPIPS, and corresponding DMOS scores. For each subset, we fit the DMOS logistic curve (as described earlier) for both metrics and compute four statistical measures (PLCC, SRCC, etc.). This is done for 2000 bootstrap iterations, generating distributions for each measure, from which we derive confidence intervals. Using a relative t-test, we calculate p-values for different sampling percentages (Table II), with $p < 0.05$ indicating statistical significance. Most measures show that PSNR_{DIV} outperforms FloLPIPS significantly. However, KRCC improvements require larger sample sizes (70%) to reach significance, suggesting that ranking consistency (KRCC) is more sensitive to dataset diversity than absolute error measures (PLCC/RMSE).

A visualisation of our weighting scheme is shown in Figure 2. As can be seen in (c-d) our divergence based metric accurately highlights areas for which temporal inconsistency is introduced by DVF. Correspondingly fewer regions of temporal inconsistency are seen in (e-f) for the more advanced ST-MFNet interpolator, but it does still highlight regions of high motion blur in the background and occlusion (e.g. edge of bicycle wheel) where interpolation errors are most likely to occur. Notably, while ST-MFNet fails to interpolate the football as seen in (f) it does appear to create a temporally consistent interpolation compared to the warping errors introduced by DVF.




Fig. 2. Frames from two reference sequences in the BVI-VFI dataset are shown in (a-b). The corresponding interpolated frame generated by DVF (c-d) and ST-MFNet (e-f) are overlaid with values of our divergence measure $d(x)$. Regions where $d(x) \geq 0.01$ are shown in yellow and regions where $d(x) = 0$ are shown in blue. Crops are shown inset outlined in red. This demonstrates the effectiveness of divergence as a metric for detecting regions of temporal inconsistency.

B. Effect of Motion Estimator and Threshold Value

To examine the effect of the choice of motion divergence threshold we performed a search across a range of values and measured the previous statistical parameters for each. In addition to make sure our metric is robust to choice of motion estimator we measure our metrics performance for three different motion estimators. We use the Färneback algorithm, being one of the most popular classical dense motion estimators, PWC-Net as a fair comparison to FloLPIPS, and RAFT [4] as a more recent learning-based motion estimator. A sweep of the motion divergence threshold parameter (T) was performed

TABLE III

THE PERFORMANCE OF OUR METRIC ON THE BVI-VFI DATASET. STATISTICAL METRICS FOR EACH MOTION ESTIMATOR ARE MEASURED USING THE MOTION DIVERGENCE THRESHOLD THAT MAXIMISES ITS CORRELATION TO PERCEPTUAL SCORES.

Motion Estimator	Threshold	540p				1080p				Overall			
		15fps→30fps		30fps→60fps		15fps→30fps		30fps→60fps		PLCC	KRCC	SRCC	RMSE
		PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC				
Farneback	0.01	0.52	0.45	0.65	0.63	0.61	0.57	0.80	0.76	0.67	0.45	0.63	13.60
PWC-Net	0.01	0.52	0.47	0.57	0.65	0.65	0.60	0.83	0.83	0.67	0.47	0.65	13.63
RAFT	0.001	0.52	0.48	0.56	0.64	0.70	0.62	0.83	0.85	0.67	0.47	0.66	13.70




Fig. 3. Plots of variation in (a) PLCC and (b) SRCC with respect to motion divergence threshold (T) for different motion estimators. Values of the threshold for which correlation is maximised are marked by a yellow square.

across values spaced by an order of magnitude. The variation of the two most salient measures, PLCC and SRCC, for this parameter sweep are shown in Figure 3 (a-b) respectively. From these we can get an estimate of a good value of threshold T to maximise correlation.

Detailed results for the optimal parameters are shown in Table III. It can be seen that for the Farneback and PWC-Net motion estimators the same threshold value results in the best correlation for both of these methods with a PLCC of 0.67 and SRCC values in the 0.63-0.65 range. RAFT’s occlusion handling reduces spurious divergence in motion fields, allowing a lower threshold to capture true artefacts without over-penalizing valid interpolations. Looking at the overall results it appears that it is better to use a deep learning based motion estimator. Notably our metric when using the same motion

estimator as FloLPIPS is still able to maintain good correlation overall and across individual framerate resolution subsets.

C. Computational Complexity

All metrics were evaluated on a 12th Gen Intel Core i7-12700K CPU (32GB RAM) and an NVIDIA RTX A4000 GPU. For fairness, both PSNR_{DIV} and FloLPIPS used PWC-Net for motion estimation. We use the implementation of FloLPIPS provided by Danier *et. al*³. We measure memory usage using the maximum memory allocation measurement provided by PyTorch⁴. As Table IV shows, PSNR_{DIV} reduces memory usage by 4× (3.7 GB vs. 15.2 GB) and runtime by 2.5× (124 ms vs. 305 ms per frame). With further optimization of the implementation, PSNR_{DIV} could approach real-time evaluation (33 ms/frame), enabling integration as a loss function for VFI DNN training.

TABLE IV
MEMORY AND RUNTIME REQUIREMENTS FOR QUALITY EVALUATION OF ONE FRAME FROM A 1080P SEQUENCE.

Metric	Max. Memory Usage (GB)	Runtime (ms)
FloLPIPS	15.2	305
PSNR _{DIV}	3.9	124

VI. CONCLUSIONS

In this paper, we introduced PSNR_{DIV}, a full reference video quality assessment metric tailored towards video frame interpolation quality assessment. Our metric builds on a method used to detect motion smoothness in archival footage restoration. We use this to weight MSE and calculate a modified PSNR that is based on areas with high temporal inconsistency. We evaluated our metric on the BVI-VFI subjective dataset. We show that this metric is competitive with a popular deep learning based VFI metric and can correlate better with human perception of interpolated content quality. Our metric is also 2.5× faster than FloLPIPS while using 4× less memory. Our metric is likely to achieve even greater efficiency with an optimised implementation. This makes it quite suitable for use as a loss function.

³<https://github.com/danier97/flopiips>

⁴https://docs.pytorch.org/docs/stable/generated/torch.cuda.max_memory_allocated.html

REFERENCES

- [1] A. Paliwal and N. Khademi Kalantari, "Deep slow motion video reconstruction with hybrid imaging system," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 42, no. 7, pp. 1557–1569, 2020.
- [2] K. M. Briedis, A. Djelouah, M. Meyer, I. McGonigal, M. Gross, and C. Schroers, "Neural frame interpolation for rendered content," *ACM Transactions on Graphics (TOG)*, vol. 40, no. 6, pp. 1–13, 2021.
- [3] Y. Guo, L. Chen, Z. Gao, and X. Zhang, "Frame rate up-conversion method for video processing applications," *IEEE Transactions on Broadcasting (TOB)*, vol. 60, no. 4, pp. 659–669, 2014.
- [4] Z. Teed and J. Deng, "RAFT: Recurrent all-pairs field transforms for optical flow," in *16th European Conference on Computer Vision (ECCV)*. Springer, 2020, pp. 402–419.
- [5] Z. Liu, R. A. Yeh, X. Tang, Y. Liu, and A. Agarwala, "Video frame synthesis using deep voxel flow," in *International Conference on Computer Vision (ICCV)*. IEEE, 2017, pp. 4463–4471.
- [6] S. Niklaus, L. Mai, and F. Liu, "Video frame interpolation via adaptive separable convolution," in *International Conference on Computer Vision (ICCV)*. IEEE, 2017, pp. 261–270.
- [7] D. Danier, F. Zhang, and D. Bull, "ST-MFNet: A spatio-temporal multi-flow network for frame interpolation," in *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE/CVF, 2022, pp. 3521–3531.
- [8] Z. Shi, X. Xu, X. Liu, J. Chen, and M.-H. Yang, "Video frame interpolation transformer," in *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE/CVF, 2022, pp. 17 482–17 491.
- [9] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE/CVF, 2018, pp. 586–595.
- [10] A. Mackin, K. C. Noland, and D. R. Bull, "High frame rates and the visibility of motion artifacts," *SMPTE Motion Imaging Journal*, vol. 126, no. 5, pp. 41–51, 2017.
- [11] D. Danier, F. Zhang, and D. Bull, "A subjective quality study for video frame interpolation," in *International Conference on Image Processing (ICIP)*. IEEE, 2022, pp. 1361–1365.
- [12] S. Ge, A. Mahapatra, G. Parmar, J.-Y. Zhu, and J.-B. Huang, "On the content bias in fr chet video distance," in *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE/CVF, 2024, pp. 7277–7288.
- [13] D. Danier, F. Zhang, and D. Bull, "FloLPIPS: A bespoke video quality metric for frame interpolation," in *Picture Coding Symposium (PCS)*. IEEE, 2022, pp. 283–287.
- [14] "Final report from the video quality experts group on the validation of objective quality metrics for video quality assessment," Video Quality Experts Group, 2000.
- [15] J. Han, X. Min, J. Jia, Y. Gao, X. Liu, and G. Zhai, "Full-reference and no-reference quality assessment for video frame interpolation," *IEEE Transactions on Circuits and Systems for Video Technology*, 2025.
- [16] G. M.  kmez, Y. Zhang, C. Schroers, and T. O. Aydin, "Clip-fusion: A spatio-temporal quality metric for frame interpolation," in *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2025, pp. 7450–7459.
- [17] P. C. Madhusudana, N. Birkbeck, Y. Wang, B. Adsumilli, and A. C. Bovik, "ST-GREED: Space-time generalized entropic differences for frame rate dependent video quality prediction," *IEEE Transactions on Image Processing (TIP)*, vol. 30, pp. 7446–7457, 2021.
- [18] J. Wu, Y. Liu, W. Dong, G. Shi, and W. Lin, "Quality assessment for video with degradation along salient trajectories," *IEEE Transactions on Multimedia (TOM)*, vol. 21, no. 11, pp. 2738–2749, 2019.
- [19] H. Men, V. Hosu, H. Lin, A. Bruhn, and D. Saupe, "Visual quality assessment for interpolated slow-motion videos based on a novel database," in *12th International Conference on Quality of Multimedia Experience (QoMEX)*. IEEE, 2020, pp. 1–6.
- [20] Q. Hou, A. Ghildyal, and F. Liu, "A perceptual quality metric for video frame interpolation," in *European Conference on Computer Vision (ECCV)*. Springer, 2022, pp. 234–253.
- [21] P. V. Pahalawatta and A. M. Tourapis, "Motion estimated temporal consistency metrics for objective video quality assessment," in *International Workshop on Quality of Multimedia Experience (QoMEX)*, 2009, pp. 174–179.
- [22] A. C. Kokaram, *Motion Picture Restoration: Digital Algorithms for Artefact Suppression in Degraded Motion Picture Film and Video*, 1st ed. Berlin, Heidelberg: Springer-Verlag, 1998.
- [23] A. Rares, M. Reinders, and J. Biemond, "Statistical analysis of pathological motion areas," in *Seminar Digital Restoration of Film and Video Archives*. IEE, 2001.
- [24] D. Corrigan, N. Harte, and A. Kokaram, "Pathological motion detection for robust missing data treatment in degraded archived media," in *International Conference on Image Processing (ICIP)*. IEEE, 2006, pp. 621–624.
- [25] D. Danier, F. Zhang, and D. R. Bull, "BVI-VFI: A video quality database for video frame interpolation," *IEEE Transactions on Image Processing (TIP)*, vol. 32, pp. 6004–6019, 2023.
- [26] X. Xu, L. Siyao, W. Sun, Q. Yin, and M.-H. Yang, "Quadratic video interpolation," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 32, 2019.
- [27] "Recommendation ITU-T P. 910. subjective video quality assessment methods for multimedia applications," International Telecommunication Union, 2000.
- [28] W. Xue, L. Zhang, X. Mou, and A. C. Bovik, "Gradient magnitude similarity deviation: A highly efficient perceptual image quality index," *IEEE Transactions on Image Processing*, vol. 23, no. 2, pp. 684–695, 2013.
- [29] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 25, 2012.
- [30] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz, "PWC-Net: Cnns for optical flow using pyramid, warping, and cost volume," in *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE/CVF, 2018, pp. 8934–8943.
- [31] G. Farneb ck, "Two-frame motion estimation based on polynomial expansion," in *13th Scandinavian Conference on Image Analysis (SCIA)*. Springer, 2003, pp. 363–370.