

BluRef: Unsupervised Image Deblurring with Dense-Matching References

Bang-Dang Pham^{1†} Anh Tran² Cuong Pham^{2,3} Minh Hoai^{2,4}

¹University of Wisconsin-Madison ²Qualcomm AI Research[‡] ³Posts & Telecommunications Inst. of Tech. ⁴Adelaide University

<https://qualcomm-ai-research.github.io/BluRef/>

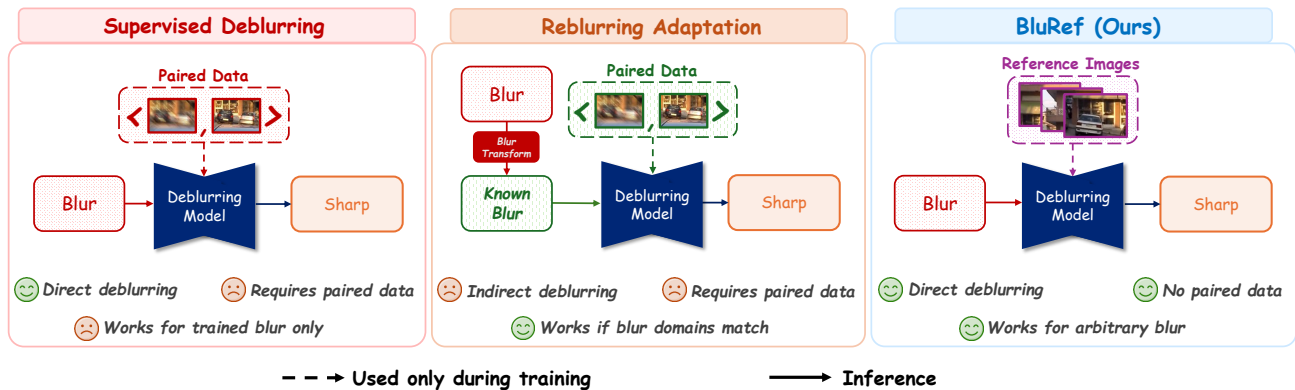


Figure 1. Comparison of three approaches to image deblurring. Supervised deblurring requires costly paired data, limiting scalability. Reblurring-based methods rely on indirect mappings and suitable intermediate domains, which can be hard to identify. In contrast, the proposed BluRef directly learns from unpaired blurry and reference images within the target domain, offering an efficient and scalable solution without requiring paired supervision or complex multi-stage pipelines.

Abstract

This paper introduces a novel unsupervised approach for image deblurring that utilizes a simple process for training data collection, thereby enhancing the applicability and effectiveness of deblurring methods. Our technique does not require meticulously paired data of blurred and corresponding sharp images; instead, it uses unpaired blurred and sharp images of similar scenes to generate pseudo-ground truth data by leveraging a dense matching model to identify correspondences between a blurry image and reference sharp images. Thanks to the simplicity of the training data collection process, our approach does not rely on existing paired training data or pre-trained networks, making it more adaptable to various scenarios and suitable for networks of different sizes, including those designed for low-resource devices. We demonstrate that this novel approach achieves state-of-the-art performance, marking a significant advancement in the field of image deblurring.

1. Introduction

Motion blur is a common issue in images and videos, reducing their visual quality and potentially impairing the performance of subsequent computer vision applications. Therefore, a reliable image deblurring technique is crucial in many scenarios. Currently, the leading approach is data-

driven, involving the training of a deblurring function using paired data comprising blurry and corresponding sharp images. However, acquiring such paired training data requires a complex setup, including a beam splitter, cameras of the same make operating at different speeds, and systems for time synchronization, geometric alignment, and color calibration [6, 10, 18]. Most cameras are not equipped to meet these stringent requirements, and for those that are, the setup is cumbersome and can be nearly impossible in specific contexts, such as for a moving dashcam or bodycam. As a result, if the images or target domain come from scenarios where collecting paired training data is impractical—which is often the case—supervised deblurring methods become unusable.

One approach that circumvents the need for paired training data in the target blur domain is the reblurring-based method [17, 26]. Instead of directly deblurring the image, it first maps the blurry input to an intermediate blur domain where a supervised deblurring model is already available, and then applies that model to recover the sharp image. This blur-to-blur-to-sharp strategy can be effective if the intermediate domain closely resembles the target blur domain, since a mapping between the two blur domains must also be learned. However, identifying such a suitable intermediate domain is often challenging, and any mismatch can negatively impact performance. Furthermore, the multi-stage pipeline introduces additional computational overhead, making it less efficient than direct deblurring approaches.

Both of the aforementioned approaches share a fundamental limitation: the domain gap. This gap arises either from

[†]Work done while at Qualcomm AI Research.

[‡]Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc.

a mismatch between the target blur domain and the domain used to train the deblurring model, or from the challenge of identifying an intermediate domain sufficiently similar for effective mapping. This raises a natural question: *Is it possible to eliminate the domain gap altogether?*

The logical place to find data without a domain gap is the target domain itself—using images drawn directly from the same distribution as the blurry images we aim to deblur. Such data is often readily available in large quantities. The main obstacle, however, is that these images typically lack paired sharp counterparts. But do we really need paired data to train a deblurring model? Our answer is no. This insight forms the core motivation of our work: we demonstrate that it is possible to train high-performing deblurring models entirely within the target domain by leveraging unpaired reference images and generating pseudo-ground truth supervision - without ever requiring paired blur/sharp datasets.

Specifically, we propose BluRef, a novel unpaired reference-guided deblurring framework that operates entirely within the target domain. BluRef learns from unpaired blurry and sharp reference images that do not need to be geometrically aligned or captured simultaneously. It leverages a Dense Matching (DM) model, pretrained on synthetically warped data, to establish correspondences between the blurred input and sharp references, generating high-quality pseudo-ground-truth images. These pseudo-sharp targets are iteratively refined and used to supervise the deblurring network, eliminating the need for paired datasets or pretrained deblurring models. This design simplifies data collection—requiring only unpaired video frames—and confines the use of synthetic data to the DM pretraining stage, ensuring strong performance on real-world blur.

A valuable byproduct of BluRef is the generation of pseudo-ground-truth pairs during its iterative refinement process. These pairs can be reused to train deblurring networks of varying capacities, including lightweight models suitable for mobile or resource-constrained devices. This extends the practical utility of BluRef beyond a single backbone, enabling flexible deployment across a range of platforms. We further analyze the pseudo-pair generation mechanism and its downstream benefits in Sec. 4.5.

The name BluRef reflects our use of reference images, but our approach should not be confused with reference-based deblurring methods [11, 34, 35], which remain within the supervised learning paradigm. These methods require paired blur-sharp data during training and incorporate sharp references to enhance texture preservation during both training and inference. For example, Lai et al. [5] employ a dual-camera setup with spatially aligned reference frames for blurry face enhancement. While such methods demonstrate the value of reference images, they do not address the core challenge of learning from unpaired, unconstrained data—central to our unsupervised BluRef framework.

In short, this paper makes three key contributions. **First**, we introduce BluRef, the first unpaired reference-guided framework for unsupervised image deblurring, eliminating the need for ground-truth supervision or paired synthetic data through iterative dense matching with unpaired references.

Second, we show that BluRef’s pseudo-ground-truth images can be reused to train models of varying capacities, including lightweight networks suitable for mobile deployment. **Third**, extensive experiments indicate that BluRef outperforms existing unsupervised methods and approaches the quality of supervised models on both synthetic and real-world datasets.

2. Related Work

Related to our image deblurring method are studies on reference-based image restoration and dense matching, which will be reviewed in this section.

Image deblurring. The proposed method stands out from other unsupervised deblurring techniques [3, 13, 29, 32] given its pioneering approach of using reference images to generate pseudo-ground truths for directly training a deblur network. A notable related approach, Pham et al. [17], converts unknown blur types into a known domain, allowing pre-trained models to handle previously unsupported blur types. However, unlike their method, which depends on a pre-trained network trained with paired data, our approach operates without this requirement. This independence is essential, as it removes the need to adapt the input blurry image to align with the capabilities of a pre-trained network, a process that can be challenging when the target deblurring domain differs significantly from those with available paired data, especially given the rigorous requirements for collecting such data. Recent supervised methods, such as ID-BLAU [26] and SAM-Deblur [7], achieve strong performance but remain constrained by the need for paired datasets. In contrast, our method attains competitive performance in a fully unsupervised setting while maintaining strong adaptability across diverse blur domains. Furthermore, it surpasses unsupervised techniques that depend on synthetic data augmentation [18, 31], whose artificial blur generation often fails to represent real-world degradation, leading to limited generalization.

Reference-based image restoration. Reference-based image restoration, particularly in super-resolution, utilizes high-resolution references to enhance restoration via feature transfer—such as Zhang et al. [33] utilizing Patch Match [1] and VGG extractors [20]—and attention mechanisms for multi-scale fusion [28]. To mitigate error-prone spatial alignment, Wang et al. [25] introduced an aligned attention method preserving high-frequency details.

In deblurring, references enhance networks using sharp video frames [27] or selection and ranking modules [8, 9]. However, these typically rely on paired datasets with short-exposure guidance Zhao et al. [34] or aligned dual-camera setups for specific domains like faces [5], whereas our approach handles complex scenes with significant motion diversity and minimal overlap. Recent efforts dividing deblurring and reference transfer [11], or focusing on similarity extraction [35], still fall behind current single-image baselines or struggle to effectively bridge blurry and sharp domains.

Most notably, prior reference-based methods primarily enhance supervised models requiring paired ground-truth data. In contrast, our work pioneers the use of reference images within a fully unsupervised framework, eliminating the

need for paired datasets. By effectively leveraging reference inputs, we introduce a novel module that outperforms existing unsupervised single-image deblurring methods, enabling high-quality restoration without relying on paired data.

Dense matching models. Dense flow regression methods, typically used for predicting correspondence maps in optical flow tasks, have recently been adapted for geometric matching to accommodate substantial variations in geometry and appearance. Notable among these are DGC-Net [14], which employs a coarse-to-fine CNN framework, and GLU-Net [23], which overcomes the resolution constraints of previous models by using global and local correlation layers, complemented by the GOCor module [22] for enhanced matching precision. Our method innovates within this landscape by being the first to apply Dense Matching for matching across domains based on the semantic similarity of objects, breaking new ground in the field of domain-specific dense correspondence.

3. BluRef – Reference-based Deblurring

We approach the problem of image deblurring through an iterative enhancement process. Our goal is to estimate a sharp image from a blurry image I_{blur} of a scene by utilizing a set of N reference sharp images $\{I_{\text{ref}}^n\}_{n=1}^N$. None of the sharp reference images needs to be the exact sharp version of the blurry input I_{blur} . While the reference images should come from scenes similar to that of the blurry image, they do not need to match precisely and can be captured from different temporal and spatial perspectives. In this section, we first describe the iterative training procedure, then provide details about the self-supervised training of a dense matching model, and finally explain the pseudo-sharp image generation method.

3.1. Iterative Target Refinement and Training

Our iterative process comprises multiple epochs, each with two main procedures: (1) **pseudo-sharp image generation**, in which we find matching details between the blurry input and the reference sharp images to create a temporary pseudo-sharp image; and (2) **deblurring network optimization**, where we use the current pseudo-sharp image as ground truth to optimize the deblurring network backbone. This is illustrated in Fig. 2 and is mathematically expressed as follows:

$$I_{\text{deblur}}^{(k)} = \mathcal{D}(I_{\text{blur}}; \Theta^{(k)}), \quad (1)$$

$$I_{\text{pseudo}}^{(k)}, M_{\text{pseudo}}^{(k)} = \mathcal{G}(I_{\text{deblur}}^{(k)}, \{I_{\text{ref}}^n\}_{n=1}^N), \quad (2)$$

$$\Theta^{(k+1)} = \underset{\Theta}{\operatorname{argmin}} \mathcal{L}_{\text{deblur}}(\mathcal{D}(I_{\text{blur}}; \Theta); I_{\text{pseudo}}^{(k)}, M_{\text{pseudo}}^{(k)}) \quad (3)$$

In the equations above, \mathcal{D} denotes the deblurring network parameterized by $\Theta^{(k)}$ at the k -th epoch, which produces a deblurred estimate $I_{\text{deblur}}^{(k)}$ from the blurry input I_{blur} . The function \mathcal{G} represents a process that takes this deblurred image and matches it with the set of unpaired sharp references $\{I_{\text{ref}}^n\}_{n=1}^N$ using a trained Dense Matching (\mathcal{DM}) model to generate a pseudo-sharp image $I_{\text{pseudo}}^{(k)}$ and a confidence

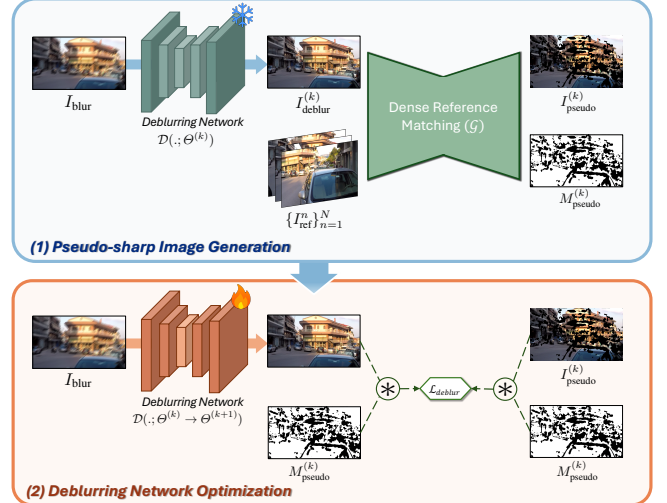


Figure 2. Iterative approach for generating a pseudo-sharp image and using it to train the deblurring model. In each epoch, the reference images $\{I_{\text{ref}}^n\}_{n=1}^N$ are matched with the deblurred result $I_{\text{deblur}}^{(k)}$ to produce the updated pseudo-sharp image and the corresponding confidence mask, which serve as the supervision targets for training the deblurring model.

mask $M_{\text{pseudo}}^{(k)}$ that suppresses unreliable correspondences (more details in Sec. 3.3). We initialize the process with $I_{\text{deblur}}^{(0)} = I_{\text{blur}}$ to provide a stable starting point for matching, before the deblurring network learns meaningful features. The pseudo-sharp image serves as the target, and the deblurring network parameters are updated by minimizing the masked reconstruction loss $\mathcal{L}_{\text{deblur}}$, where the mask $M_{\text{pseudo}}^{(k)}$ weights the valid regions during optimization. This iterative refinement leverages the progressively enhanced deblurring model to incrementally improve the pseudo-sharp target, ensuring that, after each epoch, it more closely approximates the true sharp image.

3.2. Dense Matching – Self-supervised Training

To bridge the gap between blurred and sharp images, our strategy is to extract corresponding regions across both domains. This involves training a \mathcal{DM} model to recognize patterns and similarities between blurred and sharp images, which is crucial for generating what will be referred to as a ‘pseudo-sharp’ image. In this work, the \mathcal{DM} model is defined as a function that establishes correspondences between a target image, which can be blurry or of low resolution, and a series of sharp reference images. This function is formulated as follows: $\mathcal{DM} : (I_{\text{target}}, I_{\text{ref}}) \rightarrow (I_{\text{trans}}, M_{\text{conf}})$. Here, I_{target} is the image we want to deblur, I_{ref} is a sharp reference, I_{trans} is the result of applying the inferred correspondences from I_{ref} to I_{target} , and M_{conf} is a binary confidence mask that specifies the certainty of each transformation applied.

To train the \mathcal{DM} model to effectively bridge the feature gap between blurred and sharp domains, we adopt a self-supervised scheme inspired by prior works [19, 22–24]. This approach generates training pairs, denoted as $(I_{\text{warped}}, I_{\text{gt}})$, by applying artificial warping to an initial set of sharp images from any domain. Specifically, we create each synthetic pair

by sampling a base sharp image and applying random geometric transformations, such as Homography or Thin-Plate Spline, to produce a deformed version. For producing each synthetic image pair, we first resize a given original sharp image to a larger resolution, $H' \times W'$, and apply a dense flow field of the same dimensions using random geometric transformations to create a deformed image. The warped image, I_{warped} , is then produced by centrally cropping this deformed image to the desired resolution, $H \times W$, while the ground-truth image, I_{gt} , is obtained by similarly cropping the original sharp image to the same $H \times W$ size. To enhance the model’s adaptability to match detail between sharp and blurry domains, we further augment the warped images I_{warped} by blur augmentation with random blur and noise, inspired by BSRGAN [31], ensuring the model trained on such data can effectively learn to match features across diverse visual conditions (more detail in Supplementary). In the context of our BluRef framework, I_{warped} serves as the target image I_{target} , while I_{gt} from training corresponds to the reference image I_{ref} .

3.3. Pseudo-sharp Image Generation

After acquiring a robust \mathcal{DM} model that can identify pixel correspondences between sharp and blur images, we employ it within a pseudo-sharp generation process, denoted by \mathcal{G} . This process produces a pseudo-sharp image I_{pseudo} for each blur image I_{blur} , given the set of reference sharp images $\{I_{\text{ref}}^n\}_{n=1}^N$. When the number of reference images N is greater than one, \mathcal{G} aggregates multiple reference matches to form a single pseudo-sharp image. We consider three aggregation strategies: *Weighted Average*, *Sequential Accumulation*, and *Progressive Reference Averaging*.

Weighted Average. This strategy applies the \mathcal{DM} model independently to each pair $\{I_{\text{blur}}, I_{\text{ref}}^n\}$, generating a series of pseudo-sharp images I_{pseudo}^n and corresponding confidence masks M_{pseudo}^n . The final I_{pseudo} and M_{pseudo} are obtained by averaging these intermediate results:

$$I_{\text{pseudo}}, M_{\text{pseudo}} = \frac{1}{N} \sum_{n=1}^N I_{\text{pseudo}}^n * M_{\text{pseudo}}^n, \frac{1}{N} \sum_{n=1}^N M_{\text{pseudo}}^n.$$

Sequential Accumulation. This strategy iteratively refines I_{pseudo} by using the output from the previous iteration as input for the next, ensuring the continuity of sharp details:

$$\begin{aligned} I_{\text{pseudo}}^n, M_{\text{pseudo}}^n &= \\ \mathcal{DM}(I_{\text{pseudo}}^{n-1} * M_{\text{pseudo}}^{n-1} + I_{\text{blur}} * (1 - M_{\text{pseudo}}^{n-1}), I_{\text{ref}}^n), \\ I_{\text{pseudo}}, M_{\text{pseudo}} &= I_{\text{pseudo}}^N, M_{\text{pseudo}}^N. \end{aligned} \quad (4)$$

By employing this method, we maintain consistency in pixel quality between iterations as I_{pseudo}^n builds upon the sharpness from I_{pseudo}^{n-1} , creating a continuity of detail.

Progressive Reference Averaging. This strategy combines the advantages of previous strategies by retaining sharp details from prior iterations and selectively enhancing un-

matched areas in I_{pseudo}^n :

$$\begin{aligned} I_{\text{pseudo}}^n, M_{\text{pseudo}}^n &= \mathcal{DM}(I_{\text{blur}} * (1 - M_{\text{pseudo}}^{n-1}), I_{\text{ref}}^n), \\ I_{\text{pseudo}}, M_{\text{pseudo}} &= \frac{1}{N} \sum_{n=1}^N M_{\text{pseudo}}^n * I_{\text{pseudo}}^n, \frac{1}{N} \sum_{n=1}^N M_{\text{pseudo}}^n. \end{aligned} \quad (5)$$

This strategy ensures that with each iteration, newly matched details are integrated without discarding the sharpness achieved from previous matches, culminating in a high-quality pseudo-sharp image.

3.4. Loss Function

At each epoch, our deblurring network is trained using the common image deblurring training objectives, with the main component being a reconstruction loss. However, since not all regions in the pseudo-sharp image are predicted with high certainty, we use the confidence map M_{pseudo} to weigh the contribution of each image region in the loss computation:

$$\Theta^{(k+1)} := \underset{\Theta}{\operatorname{argmin}} \mathcal{L} \left(\mathcal{D}(I_{\text{blur}}; \Theta) * \bar{M}_{\text{pseudo}}^{(k)}, I_{\text{pseudo}}^{(k)} * \bar{M}_{\text{pseudo}}^{(k)} \right).$$

Here, \mathcal{L} is any suitable loss metric, including L_1 , L_2 , or PSNR. It can be chosen based on the specific requirements for the deblurring model. $\bar{M}_{\text{pseudo}}^{(k)}$ is the binary confidence mask, obtained by binarizing the confidence map $M_{\text{pseudo}}^{(k)}$ (using a threshold of 0.7 in our experiments).

3.5. Inference

At test time, BluRef operates as a direct deblurring model. Both the dense matching module and the pseudo-sharp generation process are discarded after training. Given a trained deblurring network with parameters Θ^* , the final sharp image is produced by a single forward pass:

$$\hat{I}_{\text{sharp}} = \mathcal{D}(I_{\text{blur}}; \Theta^*).$$

This means BluRef requires only a single blurry input at inference and runs with the same computational cost as a standard deblurring backbone. Despite its powerful training framework, the inference pipeline remains simple and efficient—no extra components, no multi-stage processing, just fast and direct deblurring.

4. Experiments

4.1. Experimental Setups

Datasets. We use three datasets in our experiments, namely GoPro [15], RB2V [16], and PhoneCraft [17]. GoPro is a common image deblurring dataset that consists of paired sharp and synthetic blur images. It contains 2103 training pairs across 22 scenes and 1111 test pairs from 11 scenes. RB2V, in contrast, is a collection of real blur images from varied scenes. We specifically use the RB2V_Street subset to capture a wide range of scene types, which covers 11k training pairs and 2053 test pairs. PhoneCraft is an unpaired real-world dataset recorded with a smartphone, comprising

12 blurry and 11 sharp video clips. Each clip ranges from 30 to 40 seconds and is captured at a frame rate of 60 fps.

To adapt GoPro and RB2V datasets to our problem, we divide each dataset into separate sets for blurry and sharp reference images, ensuring distinct indices to maintain an unsupervised learning environment. For each blurry image, we pick N sharp reference images, which consist of two sets of consecutive frames shifted Δ frames from the blurry ones on both sides, as illustrated in Fig. 3. Δ is not just a measure of temporal distance but also reflects the real-world occurrence of capturing videos with both sharp and blurry frames or, in the context of dataset preparation for deblurring, capturing a scene with cameras operating at different speeds. Then, for evaluating the deblurring performance under various temporal separations, Δ is set to values of 1, 10, and 20. Next, the number of sharp reference frames N is another important hyperparameter that impacts the quality of the deblurring. We will discuss its effect in Sec. 4.3. By default, we use $N=6$ for all experiments in the other sections. Thus, in our experiment, the training dataset for GoPro includes 1050 blurry images alongside 1053 sharp images, whereas, for RB2V, it consists of 5400 blurry and 5600 sharp images.




Figure 3. Protocol to collect sharp reference images in our experiments. For each blurry image, we collect N sharp reference images as the two groups of consecutive frames from both sides, each displaced by Δ frames from the blurry image.

Unlike GoPro and RB2V, PhoneCraft consists of separate blurry and sharp video clips recorded under different motion states. Blurry frames arise from rapid camera or object motion, whereas sharp clips are captured when the camera remains static or moves slowly in a static scene. Since no temporal alignment exists between the two, we randomly select sharp references for each blurry input to eliminate temporal correlation. The PhoneCraft dataset serves as a challenging real-world benchmark. All videos are captured by a handheld smartphone in unconstrained environments, resulting in complex blur patterns and significant temporal and geometric misalignment, with no paired sharp ground truth available. This setting closely reflects practical deployment conditions. We further analyze the performance of BluRef and existing baselines on this dataset in Sec. 4.4.

Implementation Details. All model components are trained using Adam optimizer [4], with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. We employ an initial learning rate of 2×10^{-4} and implement a cosine annealing [12] schedule to gradually reduce the learning rate down to 10^{-6} with a total of 400K iterations. During training, we randomly crop these images to obtain a square shape of 256×256 pixels, and we use an identical cropping box for the blurry and its reference images. For data

augmentation, we apply random flips along both the horizontal and vertical axes, with a probability of 0.5. In addition, the \mathcal{DM} model is built upon the PDC-Net+[24] framework, with the integration of GLU-Net-GOCor [22, 23] as its base architecture. For the GOCor modules [22], the training encompasses three iterations of both local and global optimizations. For training the \mathcal{DM} model, we use only the sharp image set from the GoPro dataset, applying random degradations inspired by BSRGAN [31] and adapted for blind image deblurring with motion blur augmentation to create I_{warped} . This prevents blur patterns from the original dataset from leaking into training, ensuring the integrity of our unsupervised framework. We use this pre-trained \mathcal{DM} model for all BluRef experiments in this paper. Lastly, we use the network architectures of NAFNet [2] and Restormer [30] as the backbone of our deblurring model \mathcal{D} , which is trained from scratch using BluRef.

Baselines. We evaluate BluRef’s effectiveness with a comprehensive list of baseline methods, encompassing both unsupervised (DualGAN[29], UID-GAN [13], UAUD [21]) and supervised approaches (NAFNet [2], Restormer [30]) on RB2V and GoPro dataset. Specifically, for the BluRef framework, we assessed the performance across its three distinct strategies: Weighted Average (**Avg.**), Sequential Accumulation (**Seq.**), and Progressive Reference Averaging (**Prog.**), each integrated with the aforementioned supervised techniques. To assess performance in fully unpaired real-world scenarios, we evaluate reblurring-based adaptation methods (Blur2Blur [17], BSRGAN [31]) on the PhoneCraft dataset.

To ensure fair comparisons, all models were trained under unpaired data configurations segmented according to the Δ value. Notably, the results from the supervised models serve as upper bound benchmark, as they are trained on paired blur-sharp image pairs from each subset.

4.2. Experimental Results

Quantitative Comparison. As shown in Tab. 1, existing unsupervised methods perform significantly worse than BluRef on both synthetic and real-world datasets. In contrast, BluRef—when paired with strong backbones such as NAFNet or Restormer—achieves performance close to, and in some settings even beyond, supervised models. Notably, our proposed method with the NAFNet backbone and the Progressive Averaging strategy achieves a PSNR of 31.94 dB on GoPro for $\Delta = 1$ frame, which is close to the supervised NAFNet’s score of 33.32 dB.

On the RB2V dataset, our method with Restormer and the Progressive Averaging strategy outperforms the supervised Restormer’s upper bound of 27.43 dB with a PSNR of 27.87 dB and 27.72 dB for $\Delta = 1$ frame and 10 frames, respectively, showcasing the potential of our approach to exceed supervised methods in conditions where we can leverage information from surrounding frames. Results for the Weighted Average and Sequential Accumulation strategies follow similar trends with NAFNet and are included in the supplementary material.

When increasing the temporal distance Δ from 1 to 20

frames, BlurRef only shows a minor decrease in PSNR and SSIM scores on both datasets and with both deblurring backbones. This confirms BluRef’s robustness when reference frames differ significantly from the blurry input due to motion or content changes. The ability to remain stable under large Δ values highlights BluRef’s practicality for constructing pseudo pairs from real-world, unaligned video data.

Qualitative results. We provide a qualitative comparison between our method and baseline methods in Fig. 4. As can be seen, the unsupervised approaches, including DualGAN, UID-GAN, and UAUD, produce poor results on both datasets. Particularly, they fail to recover the lady’s body in the first example and retain strong motion blur artifacts in the second one. In contrast, BluRef, when combined with either NAFNet or Restormer backbone, can produce deblurred images with much clearer details, close to upperbound result from NAFNet and approximating the sharp ground truth. It recovers correctly the body shape, face, and hand position of the lady in the first example and the tree branches in the second example. This result is satisfactory, given the challenging inputs with extreme motion blurs.

4.3. Ablation Studies

Number of referenced images. We investigate how the number of reference images influences deblurring performance by evaluating BluRef with 4, 6, 8, and 10 reference frames. For this study, we use a NAFNet backbone, $\Delta = 1$ and the Progressive Averaging strategy based on the superior results presented in Tab. 1. As shown in Tab. 2, BluRef performs robustly across a wide range of reference counts, with 6 and 8 frames yielding similarly strong results - 6 achieving slightly higher PSNR, and 8 providing marginally better SSIM. Using only 4 references provides limited scene coverage and results in a mild drop, whereas using 10 references may introduce redundant or strongly misaligned content that degrades the aggregated pseudo-ground truth. Importantly, BluRef is not sensitive to the exact choice of N : all settings from 4 to 10 frames achieve above 31 dB PSNR, indicating that a moderate number of references is sufficient and that the method remains stable even when the reference set is not perfectly aligned. In practice, videos naturally provide many usable frames, and BluRef does not require users to manually select “good” references—its confidence masks automatically suppress unreliable matches.

Matching content between blur and sharp. The datasets used in our experiments feature continuous camera and object movements; hence, frames with a small index discrepancy Δ can exhibit significant content differences. As shown in Fig. 5, we present examples of the substantial content gap between the training blurry image and reference images when $\Delta = 10$ and $\Delta = 20$. To demonstrate the challenge of our setting, we conducted experiments using PDC-Net+ to extract correlations with a threshold of 0.7. We calculated the percentage of matched regions between the ground truth and the reference images separated by Δ frames. As reported in Tab. 3, all correlation percentages are below 40%, indicating significant content variation. Despite this, our method

performs effectively, showing only a slight difference in deblurring performance compared to when $\Delta = 1$.

4.4. Evaluating the Real-World Scenarios

Unpaired real-blur dataset – PhoneCraft. We further examine BluRef’s practical impact using the PhoneCraft dataset, which comprises videos recorded by a single moving camera across everyday scenarios. This setup lacks paired ground truth, posing a challenge for supervised deblurring while providing a rigorous test environment for unsupervised approaches. For BluRef training, we utilize blurred and sharp videos from the PhoneCraft dataset that capture the same scene but are temporally and geometrically misaligned. For baselines, we select BSRGAN as generalized image deblurring model and the state-of-the-art unsupervised deblurring method Blur2Blur [17], both combined with a pretrained NAFNet backbone. Blur2Blur is an unsupervised method that transforms a blurry image into an intermediate blur domain where a pretrained deblurring network, trained on paired data in that domain, can be applied. In this context, when deblurring an image from PhoneCraft using Blur2Blur, we can choose either RSBlur [18] or GoPro as the intermediate blur domain.

Tab. 4 reports NIQE and FID scores. Although Blur2Blur achieves strong results, it requires closely matched source and intermediate blur kernel distributions. This dependence becomes evident when using GoPro as the intermediate blur: its synthetic blur characteristics fail to match the real-world blur in PhoneCraft, causing a performance drop compared to the better-aligned RSBlur. In contrast, BluRef circumvents precise domain alignment by utilizing easily obtainable reference frames, robust to significant temporal and geometrical misalignments. This allows BluRef to maintain consistently high-quality results in real-world conditions, without relying on specific blur characteristics or domain-specific adaptations.

Challenging the supervised paradigm. Although Blur2Blur faces limitations when the blur kernels of the source and target datasets differ significantly, it provides a key advantage in transferring blur kernel characteristics between domains. We can leverage this advantage by combining BluRef with Blur2Blur to enhance the performance of both. Specifically, Blur2Blur first adapts the intermediate blur dataset, RSBlur, aligning its blur characteristics with those of PhoneCraft. Instead of training BluRef’s deblurring model from scratch, we initialize it with a pretrained NAFNet on RSBlur.

This approach allows Blur2Blur to reduce the domain gap initially, while BluRef’s iterative refinement further enhances the model’s robustness, enabling it to handle a broader range of blurs effectively. As can be seen in Tab. 4, this combined approach significantly improves performance when deblurring previously unseen blurs, such as those in PhoneCraft. Fig. 6 presents the qualitative results of the combined approach, along with the individual contributions of each model.

On the GoPro and RB2V datasets, where ground truth references are available, we use PSNR and SSIM metrics to

Delta (Δ)	GoPro			RB2V		
	1 frame	10 frames	20 frames	1 frame	10 frames	20 frames
Unsupervised Deblurring						
DualGAN	22.23/0.721	22.10/0.719	21.24/0.702	21.01/0.512	20.87/0.500	20.92/0.505
UID-GAN	23.42/0.732	23.18/0.724	22.38/0.724	22.22/0.578	22.01/0.551	22.13/0.569
UAUD	24.25/0.792	24.02/0.750	23.77/0.745	22.87/0.590	22.29/0.581	22.28/0.581
BluRef (Ours)						
NAFNet - BluRef (Avg.)	29.32/0.933	29.21/0.915	29.15/0.911	25.97/0.783	25.96/0.783	25.65/0.775
NAFNet - BluRef (Seq.)	29.82/0.947	29.68/0.940	29.60/0.940	26.14/0.790	26.02/0.787	25.93/0.780
NAFNet - BluRef (Prog.)	31.94/0.960	31.87/0.955	31.52/0.947	27.87/0.821	27.72/0.820	27.24/0.812
Restormer - BluRef (Prog.)	<u>31.02/0.950</u>	<u>30.97/0.949</u>	<u>30.95/0.938</u>	<u>26.82/0.839</u>	<u>26.76/0.832</u>	<u>26.13/0.829</u>
Supervised - Upperbound						
NAFNet		33.32/0.962			28.54/0.824	
Restormer		32.92/0.961			27.43/0.849	

Table 1. Comparison of different deblurring methods on GoPro and RB2V datasets. For each test, we report PSNR \uparrow /SSIM \uparrow scores as evaluation metrics. The best scores are in **bold** and the second best score are in underline. For a supervised method, NAFNet or Restormer, we assess its upperbound of deblurring performance by training it on the paired dataset.




Figure 4. Qualitative results of several methods on GoPro and RB2V data. This figure presents four examples, two from each dataset. For each example, a zoomed-in portion of the input image is displayed, accompanied by its ground truth counterpart and the results from BluRef and three other baselines. Additional results are included in the supplementary material.

Number of Refs.	4 frames	6 frames	8 frames	10 frames
PSNR \uparrow /SSIM \uparrow	31.42/0.942	31.94/0.960	31.93/ 0.961	31.05/0.924

Table 2. Effect of the number of reference images on BlurRef performance on the GoPro dataset.




Figure 5. Visualization of blur and sharp images along with corresponding reference frames used for training BluRef.

Dataset	GoPro	RB2V
$\Delta = 10$	36.1%	33.7%
$\Delta = 20$	28.4%	25.2%

Table 3. Percentage of matching content between blur and sharp images across the dataset.

compare the combined BluRef+Blur2Blur method against the supervisedly trained NAFNet model. In this experiment, for Blur2Blur, we use GoPro and RB2V as the source do-

Models	NIQE \downarrow /FID \downarrow
BSRGAN	13.34/10.25
Blur2Blur (GoPro as intermediate blur)	12.01/8.93
Blur2Blur (RSBlur as intermediate blur)	<u>10.07/6.28</u>
BluRef	10.43/6.45
BluRef + Blur2Blur (RSBlur as intermediate blur)	8.47/5.62

Table 4. Performance comparison of unsupervised models on the PhoneCraft dataset. The best scores are in **bold** and the second best scores are underline.

mains and RSBlur as the intermediate blur domain. For BluRef training, we set the temporal offset to $\Delta = 10$ to capture challenging reference frames that reflect realistic conditions. The results in Tab. 5 show that Blur2Blur-BluRef improves robustness and outperforms the supervised method on RB2V, a real-world blur dataset. This demonstrates that unsupervised models can perform effectively in real-world conditions, even surpassing a supervised method that requires difficult-to-collect paired training data.

4.5. Paired Dataset Building – Potential Benefits

A useful byproduct of BluRef is the set of pseudo-ground-truth images generated during iterative refinement. Paired




Figure 6. Qualitative comparison of BluRef, Blur2Blur, and BSRGAN methods for real-world scenario with PhoneCraft dataset.




Figure 7. PSNR scores (by training iteration) of pseudo-sharp images generated by BluRef when comparing to the ground-truths.

	GoPro	RB2V
BluRef	31.87/0.955	27.72/0.808
Blur2Blur (RSBlur)	31.20/0.932	27.97/0.812
BluRef + Blur2Blur (RSBlur)	33.30 / 0.963	29.62 / 0.872
Supervised - upperbound	33.32 / 0.962	28.54 / 0.824

Table 5. Comparison performance between combining BluRef and Blur2Blur with supervised deblurring network. For each test, we report $\text{PSNR}\uparrow/\text{SSIM}\uparrow$ scores as evaluation metrics.

Annotation for training	GoPro	RB2V
Real ground truth	33.32/0.962	28.54/0.824
BluRef pseudo-ground truth	31.97/0.949	27.73/0.817

Table 6. Comparison between the models trained on original ground-truth data (upperbound) and the models trained on paired datasets generated by BluRef. For each test, we report $\text{PSNR}\uparrow/\text{SSIM}\uparrow$ scores as evaluation metrics.

with their corresponding blurry inputs, these pseudo pairs can serve as training data for deblurring networks of various capacities, including lightweight models. Below, we evaluate the quality and utility of these pseudo pairs.

PSNR of pseudo-sharp images. We compute the PSNR between pseudo-sharp images and the ground truth over masked high-confidence regions on RB2V and GoPro (Fig. 7). Initially, the PSNR is low as the deblurring model is still learning the blur kernel, causing dense matching to struggle. However, after 100K iterations, the pseudo-sharp images clearly converge towards the ground truth, demonstrating progressive refinement. Visualizations are in the supplementary material

Training with generated paired data. To assess the efficacy of our pseudo-sharp images in constructing a paired dataset for image deblurring training, we trained a NAFNet [2] model on this new dataset and compared it with the model trained on the original ground truth. We set $\Delta = 10$ and

NAFNet (light)	BSRGAN + NAFNet	RSBlur + NAFNet
10.25/6.85	13.34/10.25	11.4/7.87

Table 7. Comparison between the model trained on the paired dataset generated by BluRef and two generalized image deblurring baselines on the PhoneCraft dataset, using $\text{NIQE}\downarrow/\text{FID}\downarrow$ metrics.

use 6 reference images. The results, as shown in Tab. 6, indicate that the deblurring effectiveness achieved using our pseudo-sharp images closely matches that of the original dataset, with a PSNR discrepancy < 1 dB in RB2V.

Evaluating on ‘BluRef dataset’. We further investigate the impact of the ‘BluRef dataset’ in realistic applications using the unpaired real-world dataset PhoneCraft [17]. Since the paired ground truth is unknown, we pick two generalized image deblurring models as the baselines for comparison, including BSRGAN [31] and RSBlur [18]. These models are based on the NAFNet64 backbone, the default configuration for the released NAFNet. For BluRef, we opt for a more compact NAFNet version with a width of 32 and reduced to 18 blocks. The model outputs are evaluated using no-reference quality metrics, NIQE and FID, which measure the perceptual quality and fidelity of deblurred outputs without requiring paired ground truth. As shown in Tab. 7, the results reveal that in unsupervised settings with real-world blurry datasets like PhoneCraft, the BluRef pipeline exhibits significant performance improvements over other data synthesis approaches. This underscores the efficacy of BluRef in generating data that can help enhance the quality of deblurred image.

5. Conclusions

We have introduced **BluRef**, a novel unsupervised image deblurring approach that eliminates the need for paired training data by generating high-quality pseudo-ground truths through dense matching with reference images. BluRef achieves state-of-the-art performance on the GoPro and RB2V datasets, matching the quality of supervised models, and demonstrates adaptability in real-world scenarios such as the unpaired PhoneCraft dataset. Furthermore, combining BluRef with Blur2Blur enhances robustness and bridges domain gaps, achieving superior results on diverse real-world blur conditions. These findings establish BluRef as a powerful and practical solution for unsupervised image deblurring in both controlled and real-world settings.

References

- [1] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3):24, 2009. 2
- [2] Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *ECCV*, pages 17–33. Springer, 2022. 5, 8
- [3] Lufei Chen, Xiangpeng Tian, Shuhua Xiong, Yinjie Lei, and Chao Ren. Unsupervised blind image deblurring based on self-enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25691–25700, 2024. 2
- [4] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [5] Wei-Sheng Lai, Yichang Shih, Lun-Cheng Chu, Xiaotong Wu, Sung-Fang Tsai, Michael Krainin, Deqing Sun, and Chia-Kai Liang. Face deblurring using dual camera fusion on mobile phones. *ACM Transactions on Graphics (TOG)*, 41(4):1–16, 2022. 2
- [6] Anat Levin, Yair Weiss, Fredo Durand, and William T Freeman. Understanding and evaluating blind deconvolution algorithms. In *2009 IEEE conference on computer vision and pattern recognition*, pages 1964–1971. IEEE, 2009. 1
- [7] Siwei Li, Mingxuan Liu, Yating Zhang, Shu Chen, Haoxiang Li, Zifei Dou, and Hong Chen. Sam-deblur: Let segment anything boost image deblurring. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2445–2449. IEEE, 2024. 2
- [8] Yaowei Li, Ye Luo, and Jianwei Lu. Reference-guided deep deblurring via a selective attention network. *Applied Intelligence*, pages 1–13, 2022. 2
- [9] Yaowei Li, Jinshan Pan, Ye Luo, and Jianwei Lu. Deep ranking exemplar-based dynamic scene deblurring. *IEEE Transactions on Image Processing*, 31:2245–2256, 2022. 2
- [10] Chih-Hung Liang, Yu-An Chen, Yueh-Cheng Liu, and Winston H Hsu. Raw image deblurring. *IEEE Transactions on Multimedia*, 24:61–72, 2020. 1
- [11] Cunzhe Liu, Zhen Hua, and Jinjiang Li. Reference-based dual-task framework for motion deblurring. *The Visual Computer*, pages 1–15, 2023. 2
- [12] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 5
- [13] Boyu Lu, Jun-Cheng Chen, and Rama Chellappa. Uid-gan: Unsupervised image deblurring via disentangled representations. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 2(1):26–39, 2019. 2, 5
- [14] Iaroslav Melekhov, Aleksei Tiulpin, Torsten Sattler, Marc Pollefeys, Esa Rahtu, and Juho Kannala. Dgc-net: Dense geometric correspondence network. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1034–1042. IEEE, 2019. 3
- [15] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *CVPR*, 2017. 4
- [16] Bang-Dang Pham, Phong Tran, Anh Tran, Cuong Pham, Rang Nguyen, and Minh Hoai. Hypercut: Video sequence from a single blurry image using unsupervised ordering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9843–9852, 2023. 4
- [17] Bang-Dang Pham, Phong Tran, Anh Tran, Cuong Pham, Rang Nguyen, and Minh Hoai. Blur2blur: Blur conversion for unsupervised image deblurring on unknown domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 1, 2, 4, 5, 6, 8
- [18] Jaesung Rim, Geonung Kim, Jungeon Kim, Junyong Lee, Seungyong Lee, and Sunghyun Cho. Realistic blur synthesis for learning image deblurring. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 1, 2, 6, 8
- [19] Ignacio Rocco, Relja Arandjelovic, and Josef Sivic. Convolutional neural network architecture for geometric matching. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6148–6157, 2017. 3
- [20] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2
- [21] Xiaole Tang, Xile Zhao, Jun Liu, Jianli Wang, Yuchun Miao, and Tiejong Zeng. Uncertainty-aware unsupervised image deblurring with deep residual prior. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9883–9892, 2023. 5
- [22] Prune Truong, Martin Danelljan, Luc V Gool, and Radu Timofte. Gocor: Bringing globally optimized correspondence volumes into your neural network. *Advances in Neural Information Processing Systems*, 33:14278–14290, 2020. 3, 5
- [23] Prune Truong, Martin Danelljan, and Radu Timofte. Glu-net: Global-local universal network for dense flow and correspondences. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6258–6268, 2020. 3, 5
- [24] Prune Truong, Martin Danelljan, Radu Timofte, and Luc Van Gool. Pdc-net+: Enhanced probabilistic dense correspondence network. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 3, 5
- [25] Tengfei Wang, Jiaxin Xie, Wenxiu Sun, Qiong Yan, and Qifeng Chen. Dual-camera super-resolution with aligned attention modules. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2001–2010, 2021. 2
- [26] Jia-Hao Wu, Fu-Jen Tsai, Yan-Tsung Peng, Chung-Chi Tsai, Chia-Wen Lin, and Yen-Yu Lin. Id-blau: Image deblurring by implicit diffusion-based reblurring augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25847–25856, 2024. 1, 2
- [27] Xinguang Xiang, Hao Wei, and Jinshan Pan. Deep video deblurring using sharpness features from exemplars. *IEEE Transactions on Image Processing*, 29:8976–8987, 2020. 2
- [28] Fuzhi Yang, Huan Yang, Jianlong Fu, Hongtao Lu, and Bainig Guo. Learning texture transformer network for image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5791–5800, 2020. 2
- [29] Zili Yi, Hao Zhang, Ping Tan, and Minglun Gong. Dualgan: Unsupervised dual learning for image-to-image translation. In *Proceedings of the IEEE international conference on computer vision*, pages 2849–2857, 2017. 2, 5
- [30] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *CVPR*, pages 5728–5739, 2022. 5

- [31] Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4791–4800, 2021. [2](#), [4](#), [5](#), [8](#)
- [32] Youjian Zhang, Chaoyue Wang, and Dacheng Tao. Neural maximum a posteriori estimation on unpaired data for motion deblurring. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. [2](#)
- [33] Zhifei Zhang, Zhaowen Wang, Zhe Lin, and Hairong Qi. Image super-resolution by neural texture transfer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7982–7991, 2019. [2](#)
- [34] Yuzhi Zhao, Yongzhe Xu, Qiong Yan, Dingdong Yang, Xuehui Wang, and Lai-Man Po. D2hnet: Joint denoising and deblurring with hierarchical network for robust night image restoration. In *European Conference on Computer Vision*, pages 91–110. Springer, 2022. [2](#)
- [35] Han Zou, Masanori Suganuma, and Takayuki Okatani. Reference-based motion blur removal: Learning to utilize sharpness in the reference image. *arXiv preprint arXiv:2307.02875*, 2023. [2](#)