# Kalman-Inspired Feature Propagation for Video Face Super-Resolution

Ruicheng Feng<sup>®</sup>, Chongyi Li<sup>®</sup>, and Chen Change Loy<sup>®</sup>

S-Lab, Nanyang Technological University, Singapore {ruicheng002, ccloy}@ntu.edu.sg lichongyi25@gmail.com

Abstract. Despite the promising progress of face image super-resolution, video face super-resolution remains relatively under-explored. Existing approaches either adapt general video super-resolution networks to face datasets or apply established face image super-resolution models independently on individual video frames. These paradigms encounter challenges either in reconstructing facial details or maintaining temporal consistency. To address these issues, we introduce a novel framework called Kalman-inspired Feature Propagation (KEEP), designed to maintain a stable face prior over time. The Kalman filtering principles offer our method a recurrent ability to use the information from previously restored frames to guide and regulate the restoration process of the current frame. Extensive experiments demonstrate the effectiveness of our method in capturing facial details consistently across video frames. Code and video demo are available at https://jnjaby.github.io/projects/KEEP/.

# 1 Introduction

The field of Face Super-Resolution (FSR), which focuses on reconstructing highresolution (HR) face images from highly degraded versions, has witnessed remarkable progress. In particular, numerous studies have successfully leveraged various types of prior information, such as geometric facial priors [7,8,53], reference priors [26–28], generative priors [2,4,45,52], and codebook priors [13,46,56]. These approaches have significantly advanced the realism and quality of generated face images. However, the majority of these studies are confined to still images, with the extension to Video Face Super-Resolution (VFSR) remaining relatively under-explored. Despite the substantial potential benefits of video face restoration in various practical domains, such as the restoration of old films, VFSR has yet to receive the same level of attention and development as its image-based counterpart.

Two main strategies emerge for implementing VFSR. The first approach involves adapting general Video Super-Resolution (VSR) networks, such as EDVR [44], BasicVSR [3], BasicVSR++ [5], and RVRT [30], to large-scale face video datasets [34, 48]. These methods exploit temporal information and propagate features across video frames. However, they are not specifically tailored for



Fig. 1: Comparing main VFSR strategies. We show seven frames with an interval of 6. Generic VSR model BasicVSR [3] fails to reconstruct facial components faithfully. Single-image FSR model CodeFormer [56] hallucinates unnatural and inconsistent face details. Our method, in contrast, enables consistent restoration of low-quality face video while preserving temporal coherence across frames.

face restoration and often fall short in reconstructing detailed facial features, particularly in severely degraded scenarios, as depicted in Fig. 1(second row). The second method involves applying existing face image SR models to process each video frame independently. This frame-by-frame approach, while straightforward, introduces temporal inconsistencies in the video, as demonstrated in Fig. 1(third row). This problem arises because FSR is inherently ill-posed and the existing priors may not suffice to maintain appearance consistency throughout the video sequence. Specifically, a single degraded image could correspond to multiple high-resolution interpretations, leading to discrepancies and inconsistent structures across independently processed video frames.

In this paper, we wish to devise an effective framework for maintaining a stable face prior over time for VFSR. We use CodeFormer [56], a representative model that exploits codebook priors for FSR, to demonstrate how face priors can be consistently preserved across different time frames. A pivotal aspect of our approach is the idea that frames previously restored can act as references, guiding and regulating the restoration process of the current frame. This strategy helps minimize the divergence between consecutive frames. Moreover, this reliance on previously restored frames naturally suggests a recurrent framework, enabling the effective use of information from past restorations. This intuition aligns closely with Kalman filtering principles, or linear quadratic estimation, which involves using a sequence of time-based measurements, despite their statistical noise and inaccuracies, to produce more accurate estimates of unknown

variables than would be possible with a single measurement. Similarly, in VFSR, faces observed over time are often noisy and inaccurate, making them suitable for refinement using Kalman filtering techniques.

Driven by these insights, we formulate a novel method, Kalman-inspired fEaturE Propagation (KEEP), which recurrently updates the current latent state in CodeFormer by incorporating information from preceding frames. This method of temporal propagation within the latent space ensures the stability of the face prior over time, thereby capturing facial details that consistently match in appearance. The effectiveness of KEEP is shown in Fig. 1, where it is evident that our method delivers high-quality restoration with superior consistency, compared to both generic video restoration methods fine-tuned for faces and approaches that restore frames independently. Please refer to the supplementary video to appreciate the superiority of our approach in terms of temporal consistency. A key advantage of KEEP is its robustness in handling severe video-based degradation, outperforming single-image models. In addition, our model demonstrates enhanced performance on non-frontal faces by providing more stable estimations of face priors.

In summary, the main contribution of this work is a novel framework for maintaining a stable and meaningful face prior for VFSR. While we demonstrate its application using the CodeFormer method as a case study, the underlying principles of our framework, inspired by the Kalman filtering approach, are applicable to other approaches. Extensive experimental results on the VFHQ dataset [48] and real-world data demonstrate the effectiveness of our approach in improving both the fidelity and coherence of VFSR outputs. Compared to other state-ofthe-art methods, our model achieves superior performance with a large margin of 0.8 dB in PSNR, while also significantly maintaining temporal coherence.

### 2 Related Work

**Blind Face Restoration.** Blind face restoration aims at recovering severely degraded face images in the wild. Unlike natural images, faces are highly structured. This property allows researchers to incorporate prior information into the restoration models, which have demonstrated remarkable progress in capabilities to restore high-quality faces. Most existing FSR methods can be categorized into four classes: geometric priors, reference priors, generative priors, and codebook priors. *Geometric priors* usually include facial elements, such as face landmarks [8], parsing maps [7], and facial component heatmaps [53]. Another major line is reference-based methods that require high-quality exemplar images. GFRNet [28] and ASFFNet [27] leverage a warped high-quality image to extract rich details to improve facial detail restoration. DFDNet [26] constructs deep dictionaries with facial components from large-scale images to recover fine details. Generative priors from pre-trained GAN, e.g., StyleGAN [20] and StyleGAN2 [21], are employed through iterative latent optimization of GAN inversion [12, 33, 37]. Still, they produce face images with low fidelity and are computationally expensive. To address this, GLEAN [2], GPEN [52], and GFP- GAN [45] integrate generative priors into encoder-decoder architectures, which estimate latent priors in one-forward pass. These methods achieve great tradeoff between quality and fidelity but usually fail when the corruption is severe. *Codebook priors* [13,46,56] can be regarded as a special case of generative priors. In contrast to continuous generative priors, they squeeze the latent space into a small finite codebook space and improve the robustness to severe degradation. However, most existing FSR methods are image-based and thus they cannot guarantee temporally consistent details for VFSR.

Video Super-Resolution. Most existing video restoration techniques can be categorized into two paradigms based on their parallelizability: parallel and recurrent methods. Parallel models estimate all frames simultaneously, and the restoration of each frame does not rely on the update of other frames. These methods typically involve feature extraction, feature alignment, feature fusion, and reconstruction. Representative works, FSTRN [25] and VESCPN [1], introduce fast spatio-temporal networks using 3D convolutions to enhance alignment, combining motion compensation and super-resolution. EDVR [44] and TDAN [42] leverage deformable convolutions for aligning adjacent frames. Besides, Transformer-based methods [29,32] are proposed to reconstruct all frames simultaneously by jointly extracting, aligning, and fusing features. In addition, RVRT [30] and TTVSR [31] integrate optical flow into Transformer and enable long-range models in videos. Empowered by the great expressive capability of Transformer, this line of work exhibits remarkable performance improvements over previous methods. However, they suffer from large model sizes and high memory consumption. Recurrent methods [3, 16, 18] do not aggregate information solely from adjacent frames. Instead, they maintain hidden states to convey relevant information from previous frames and propagate latent features sequentially, accumulating information for later restoration. For example, RBPN [14] treated each frame as a separate source, combined iteratively in a refinement framework. RSDN [17] divided the input into structure and detail components, proposing a two-stream structure-detail block to learn textures. BasicVSR [3] and BasicVSR++ [5] fused bidirectional hidden states from both past and future frames, bring significant improvements. They aim to fully utilize information from the entire sequence, synchronously updating the hidden state through the weights of the reconstruction network. Due to the recurrent nature of feature propagation, recurrent methods experience information loss.

Video Consistency. In addition to VSR, previous works also attempt to inflate image models into video models [43, 47, 50, 51] or improve temporal consistency with the implicit representation of the given videos [23, 24, 35]. For example, Stitch-it-in-Time [43] discovers locally consistent pivots in the latent space to provide spatially consistent transitions. Tune-A-Video [47] adopts cross-frame attention and fine-tune it on a single video. All-in-One deflicker [23] learns a neural atlas for each video to solve long-term inconsistency. These works either require post-processing or optimization on a video basis. In contrast, we aim at repurposing image FSR models and enforcing temporal consistency in the restored face videos.

### 3 Methodology

In a corrupted face video, local textures and facial details are irrevocably lost. Therefore, the latent codes, usually estimated by an encoder, are inaccurate to match the real underlying priors of ground-truth video. Different from image-based autoencoder, models in the video domain enable us to exploit evidence accumulated from preceeding frames for better restoration. In this paper, we repurpose the image-based CodeFormer for video FSR and propose KEEP to estimate stable face priors in latent space over time given the noisy and inaccurate estimations, which consequently enable temporally coherent restoration. An appealing idea to realize KEEP is through reformulating the method in a Kalman Filter framework.

#### 3.1 Formulation

State Space Model. We consider observations of low-quality (LQ) video sequence  $X = \{\boldsymbol{x}_t\}_{t=1}^T$  with length of T, where  $\boldsymbol{x}_t \in \mathbb{R}^{H \times W \times 3}$ , and underlying high-quality (HQ) sequences  $Y = \{\boldsymbol{y}_t\}_{t=1}^T$ . Kalman filter [19] assumes linear dynamic systems that are characterized by a state space model driven by Gaussian noise

$$\boldsymbol{y}_t = \boldsymbol{F}_t \boldsymbol{y}_{t-1} + \boldsymbol{q}_t, \tag{1}$$

where  $F_t$  is the transition matrix and  $q_t$  denotes process noise drawn from Gaussian noise. The observation  $x_t$  is measured by

$$\boldsymbol{x}_t = \boldsymbol{H}\boldsymbol{y}_t + \boldsymbol{r}_t, \tag{2}$$

where H is the measurement matrix and  $r_t$  represents measurement noise. However, the linear assumption does not hold in some complex real-world scenarios. Hence, the non-linear Kalman filter can be reformulated as

$$\boldsymbol{y}_t = d(\boldsymbol{y}_{t-1}, \boldsymbol{q}_t) \tag{3}$$

$$\boldsymbol{x}_t = h(\boldsymbol{y}_t) + \boldsymbol{r}_t, \tag{4}$$

where  $d(\cdot)$  and  $h(\cdot)$  are non-linear transition and measurement models. Specifically,  $d(\cdot)$  can be represented by any explicit motion estimation (*e.g.*, optical flow), which defines how the current frame transits to the next one.  $h(\cdot)$  commonly models the degradation in video restoration problems. As opposed to the classical assumptions in Kalman filter, the measurement function  $h(\cdot)$  is unknown in a blind setting. This is regarded as *partially* known dynamic models [39].

Inspired by VQGAN [10] and Stable Diffusion [40], we estimate underlying latent representations  $Z = \{\boldsymbol{z}_t\}_{t=1}^T$  such that  $\boldsymbol{z}_t$  can correspond to  $\boldsymbol{y}_t$  by a generative model  $g_{\theta}$ , given by

$$\boldsymbol{y}_t = g_{\boldsymbol{\theta}}(\boldsymbol{z}_t). \tag{5}$$

Instead of directly estimating individual pixels, modeling the low-dimensional latent code is computationally more efficient and focuses on more perceptually significant variations of the data. The graphical model is depicted in Fig. 2 (a).



Fig. 2: (a) Graphical model of state space. It defines the underlying dynamic system model, where f describes how the latent states  $z_t$  transit over time, g is a generative model, and h models the degradation from clean frame  $y_t$  to degraded frame  $x_t$ . (b) Block diagram of Kalman filter model. In each time step, a predictive state from previous frame  $\hat{z}_{t-1}^+$  (Blue dash box) and new observed state of current frame  $x_t$  (Red dash box) are fused by Kalman gain  $\mathcal{K}_t$  from Kalman Gain Network (KGN) to produce more accurate estimates. The combined state  $\hat{z}_t^+$  is then used to generate the estimated clean frame  $\hat{y}_t$  by  $g_{\theta}$ . Note that  $\tilde{z}_1$  goes along with  $\tilde{z}_{t-1}$  as an anchor and it is omitted in the diagram for simplicity.

**Kalman Filter Model.** The principles of Kalman filter can be formulated by a two-step procedure, *i.e.*, *state prediction* and *state update*. The overview diagram is illustrated in Fig. 2 (b). In this problem, the observation is a face image  $\boldsymbol{x}_t$ , and the state is  $\boldsymbol{z}_t$ .

In the state prediction step, the model predicts the prior estimation  $\hat{z}_t^-$  of the current state  $z_t$  based on the posterior estimation  $\hat{z}_{t-1}^+$  of the previous state and the dynamic model. Specifically, the prior estimation of latent state  $z_t$  and estimation of observation  $x_t$  are computed as

$$\hat{z}_t^- = f(\hat{z}_{t-1}^+),$$
 (6)

$$\hat{\boldsymbol{x}}_t^- = h(g_\theta(\hat{\boldsymbol{z}}_t^-)). \tag{7}$$

The system dynamics f define how the latent state Z evolves over time, and it incorporates any control inputs that might affect the current state.

In the state update step, a posterior state estimation  $\hat{z}_t^+$  is computed based on the prior estimation and new observations  $x_t$  as

$$\hat{\boldsymbol{z}}_t^+ = \hat{\boldsymbol{z}}_t^- + \mathcal{K}_t \Delta \boldsymbol{z}_t, \tag{8}$$

where  $\mathcal{K}_t$  is conceptually referred to as the Kalman gain and  $\Delta z_t$  as the innovation, *i.e.*, the residual between the prior estimation  $\hat{z}_t^-$  and approximation of current state from  $x_t$ , given by

$$\Delta \boldsymbol{z}_t = \hat{\boldsymbol{z}}_t^- - \tilde{\boldsymbol{z}}_t, \tag{9}$$

where  $\tilde{\boldsymbol{z}}_t = e(\boldsymbol{x}_t)$ . Note that this formula differs from the original Kalman filter which minimizes the innovation of  $\Delta \boldsymbol{x}_t$  (*i.e.*, residual between  $\boldsymbol{x}_t$  and  $\hat{\boldsymbol{x}}_t^-$ ), since



Fig. 3: Overview of the proposed KEEP. It consists of four modules: encoder  $\mathcal{E}_L$ , decoder  $\mathcal{D}_Q$ , Kalman filter network, and CFA. We illustrate the information flow in one timestep.

the assumption of available measurement function  $h(\cdot)$  is no longer valid in our setting. An original measurement system models how observation  $\boldsymbol{x}_t$  is derived from the latent state  $\boldsymbol{z}_t$ , formally  $\boldsymbol{x}_t = h(\boldsymbol{z}_t)$ . Inspired by KFNet [55], we directly estimate the state  $\boldsymbol{z}_t$  given new observation  $\boldsymbol{x}_t$ , *i.e.*, mapping  $\boldsymbol{x}_t$  to  $\tilde{\boldsymbol{z}}_t$  with an estimator  $\boldsymbol{e}$ .

The remaining problem is to compute the Kalman gain  $\mathcal{K}_t$ . As discussed in KalmanNet [39], covariance estimation is intractable when dealing with highdimensional signals. Additionally, the second-order statistical moments are only used for calculating Kalman gain. Inspired by this, we directly learn the gains from the data distribution and do not explicitly maintain an estimation of covariances. Additionally, we follow [51] to include  $\tilde{z}_1$  of the first frame as anchor into KGN for Kalman gain estimation. Then, the final predicted  $\hat{y}_t$  can be derived by

$$\hat{\boldsymbol{y}}_t = g_\theta(\hat{\boldsymbol{z}}_t^+). \tag{10}$$

#### 3.2 Parameterized Models

As shown in Fig. 2 (b), we will parameterize or define the system dynamics f, observation estimator e, Kalman Gain Nets (KGN), and generative model  $g_{\theta}$  in this section. The overall framework is shown in Fig. 3.

**Generative Model.** The generative model  $g_{\theta}$  is generally parameterized by a backbone of CodeFormer [56], which consists of a LQ encoder  $\mathcal{E}_L$ , a HQ encoder  $\mathcal{E}_H$ , a codebook lookup Transformer and quantization layer  $T_Q$ , and a decoder  $\mathcal{D}$ . For simplicity,  $\mathcal{D}_Q$  denotes the decoder with the codebook lookup Transformer and quantization layer absorbed. Basically, the predicted  $\hat{y}_t$  is given by  $\mathcal{D}_Q(\hat{z}_t^+)$ , and the observed state  $\tilde{z}_t$  is approximated by

$$\tilde{\boldsymbol{z}}_t = e(\boldsymbol{x}_t) = \mathcal{E}_L(\boldsymbol{x}_t). \tag{11}$$

8 R. Feng et al.

**State Dynamic System.** The system dynamics define how the system evolves over time, and it incorporates any control inputs that might affect the current state. The prediction for the state  $z_t$  at the current timestep is achieved via state extrapolation. In particular, given the posterior estimation of previous state  $\hat{z}_{t-1}^+$ , we define the dynamic model by

$$\hat{z}_{t}^{-} = f(\hat{z}_{t-1}^{+}) = \mathcal{E}_{H}(\omega(\mathcal{D}_{Q}(\hat{z}_{t-1}^{+}), \Phi_{t-1 \to t})).$$
(12)

where  $\Phi_{t-1 \to t}$  denotes the flow estimated from LQ frames  $\boldsymbol{x}_{t-1}$  to  $\boldsymbol{x}_t$  and  $\omega$  the spatial warping modules. Specifically, we first decode the estimated code  $\hat{\boldsymbol{z}}_{t-1}^+$  of the previous frame to get the estimation of  $\hat{\boldsymbol{y}}_{t-1} = \mathcal{D}_Q(\hat{\boldsymbol{z}}_{t-1}^+)$  and warp it to the current frame. Then, it was encoded back in the latent space to obtain the prediction of the current state  $\hat{\boldsymbol{z}}_t^-$ .

**Kalman Filter System.** Given the approximated observed state  $\tilde{z}_t$  and prior estimation  $\hat{z}_t^-$  from system dynamics, the filter system aims to promote temporal information propagation and maintain stable latent code priors. In particular, the filter recursively fuses both the estimation to form a more accurate posterior estimate of the current state  $\hat{z}_t^+$ , which is also known as *state update*.

According to Eqn. 8, a more intuitive way to express the updated state is a linear interpolation  $\hat{z}_t^+ = \mathcal{K}_t \hat{z}_t^- + (1 - \mathcal{K}_t) \tilde{z}_t$  for  $\mathcal{K}_t$  normalized in the range of [0, 1]. The Kalman Gain  $\mathcal{K}_t$  measures the estimated accuracy of the predicted states compared to the approximated observed states, to update the state and reduce the uncertainty. As illustrated in Fig. 3, the Kalman gain is approximated via Kalman Gain Network (KGN), consisting of two distinct parameterized modules, *i.e.*, uncertainty network and gain network. The uncertainty associated with the current prediction is implicitly estimated by an uncertainty network constructed by spatial-temporal attention [47] and temporal attention layers (or any other recurrent networks). Then, a gain network calculates Kalman gain  $\mathcal{K}_t$  for each code token. Please refer to the supplementary files for the detailed architectures.

#### 3.3 Local Temporal Consistency

Inspired by [51], we adopt cross-frame attention (CFA) layers in the decoder to further promote local temporal consistency to regularize the information propagation. Specifically, given the latent features from the previous frame  $v_{t-1}$  and current frame  $v_t$ . They are projected onto the embedding space and output the features  $v'_i$  by  $v'_i = \operatorname{Attn}(Q, K, V) = \operatorname{softmax}(\frac{QK^T}{\sqrt{d}}) \cdot V$ , where

$$Q = W_Q \cdot v_t, K = W_K \cdot v_{t-1}, V = W_V \cdot v_{t-1}.$$
(13)

Intuitively, cross-frame attention modules can be regarded as searching and matching similar patches from the previous frame and fusing them correspondingly. This module facilitates temporal information propagation in the decoder. We adopt CFA modules on features of small scale 16 and 32 to avoid introducing blur to the decoded results.

## 4 Experiments

**Dataset.** VFHQ [48] contains over 15,000 high-quality video clips of diverse interviews and talk shows, where 15,381 clips are used for training and 50 clips are reserved for testing. Each sequence consists of 100 to 900 frames of resolution  $512 \times 512$ . Following common practice [6,36,45,48], we adopt blind settings in all experiments. Specifically, we apply random blur, resize, and noise as image-based degradations. Moreover, video compression is adopted to control the video quality by changing streaming bitrate. For a comprehensive evaluation, we synthesize three splits of the VFHQ-Test dataset containing different levels of degradation, denoted as VFHQ-mild, VFHQ-medium, and VFHQ-heavy. They follow the same degradation model but differ in the degree of noise, blur, and compression. In addition to synthetic data, we also collect real corrupted video for testing. Please refer to the supplementary files for more details.

Alignment. Pre-trained image models for face restoration are trained on the FFHQ dataset [20], where each image is automatically cropped and aligned. Hence, employing pre-trained models requires a similar alignment phase on VFHQ dataset [48]. However, the discrete step (*i.e.*, cropping) is sensitive to the detected locations of landmarks, which can consequently result in unintentional temporal inconsistencies. Inspired by the work of Fox *et al.* [11] and Tzaban *et al.* [43], we employ a Gaussian lowpass filter over the landmarks. We find that this smoothing can significantly attenuate the inconsistencies induced by the alignment step. See supplementary files for more details.

**Implementations.** For all stages of training, we initialize all networks with Kaiming Normal [15] and train them using Adam optimizer [22], and a batch size of 4 for all the experiments. The learning rate is set to  $2 \times 10^{-4}$  for stages I and II, and  $1 \times 10^{-4}$  for stage III. The models are trained with 800k, 400k, and 50k iterations for three stages, respectively. We implement our models with PyTorch [38] and train them using NVIDIA Tesla V100 GPUs. Hyper-parameters  $\lambda_1$ ,  $\lambda_{VGG}$ , and  $\lambda_{GAN}$  are set to  $10^{-2}$ , 1, and  $1 \times 10^{-2}$ . We use GMFlow [49] for optical flow estimation.

Metrics. For quantitative evaluation, we evaluate the fidelity of restoration using PSNR, SSIM, and LPIPS [54]. In addition, we evaluate the identity preservation scores, termed as IDS, by cosine similarity of the off-the-shelf identity detection network ArcFace [9]. Besides, we follow Tzaban *et al.* [43] to measure the pose consistency using Average Keypoint Distance (AKD), which is quantified by the average distance of detected landmarks between the generated and ground-truth video frames. In addition to the above single-frame quality evaluation, we also measure the fluctuation of identity/landmarks across frames. Hence,  $\sigma_{IDS}$  measures the standard deviation of identity similarity over the entire video. We expect a considerable identity drift in the generated videos without local identity jitter, where  $\sigma_{IDS}$  is supposed to be low. Similarly, we use  $\sigma_{AKD}$ to measure the standard deviation of keypoint distances over the video, which quantifies the temporal consistency of the pose. 10 R. Feng et al.

**Table 1: Quantitative comparison on the VFHQ-mild.** Red and Blue indicate the best and the second best results. Full results on other test partitions (medium and heavy) are presented in the supplementary material.

Method	$\ PSNR\uparrow$	$\mathrm{SSIM}\uparrow$	LPIPS↓	$ $ IDS $\uparrow$	$\mathrm{AKD}{\downarrow}$	$\sigma_{IDS}(\times 10^{-2})$ .	$\sigma_{AKD}\downarrow$
GPEN [52]	25.5193	0.7517	0.2988	0.7142	11.4691	4.7416	3.5109
GFPGAN [45]	26.2933	0.7795	0.2482	0.7437	10.5467	4.5700	3.6482
RestoreFormer [46]	25.5720	0.7344	0.3195	0.7530	10.5354	4.7159	3.4122
CodeFormer [56]	24.6597	0.7454	0.2742	0.6272	11.4983	6.3726	3.6927
EDVR [44]	26.6051	0.7858	0.2484	0.7195	11.6220	4.8048	3.5829
BasicVSR [3]	$\ 26.0458$	0.7765	0.2496	0.6973	11.3679	5.0343	3.6054
BasicVSR++[5]	27.1996	0.8057	0.1958	0.7641	11.3136	5.2543	4.6425
KEEP (Ours)	27.9994	0.8267	0.1619	0.7960	8.8182	3.6866	3.2538

#### 4.1 Comparison with State-of-the-Art Methods

**Baselines.** We compare our method to two categories of approaches. i) Imagebased Face SR models (CodeFormer [56], GPEN [52], GFPGAN [45], Restore-Former [46]) are used to generate face videos frame-by-frame. ii) We retrain the general VSR models (EDVR [44], BasicVSR [3], BasicVSR++ [5]) on VFHQ dataset [48]. The degradation settings remain unchanged as our experiments while other training settings follow their original papers.

Quantitative Evaluation. The quantitative results are listed in Table 1. We observe that our method achieves better results than existing methods across all the metrics. The results indicate that KEEP can faithfully recover facial details while preserving the identity. Our method also maintains temporal coherence across frames, as quantified by  $\sigma_{IDS}$  and  $\sigma_{AKD}$ , which represent the fluctuation of restored face identities and facial shapes. Though exhibiting structural distortions and artifacts (See Fig. 4), general VSR models (EDVR, BasicVSR, BasicVSR++) typically achieve higher performance on fidelity metrics (PNSR, SSIM, LPIPS) than single-image FSR models. This suggests that image-based models produce high-quality but relatively low-fidelity results. Inconsistency could be introduced when latent code estimation is noisy and inaccurate.

Qualitative Evaluation. In Fig. 4, we can observe that the compared methods fail to reconstruct consistent appearances with perceptually pleasant details. For example, GFPGAN tends to hallucinate facial details (*e.g.*, ears in the second frame and incomplete glass in the last frame of the left example). CodeFormer produces unnatural facial shapes (*e.g.*, eyes), and BasicVSR leaves severe artifacts on the face images. In the last frame of the right example, both GFPGAN and CodeFormer generate unpleasant eyes (see yellow arrows). In contrast, our KEEP exploits temporal information and restores finer and coherent facial details. We refer readers to supplementary files for more video results.



**Fig. 4: Qualitative comparison on the VFHQ.** Our KEEP produces high-fidelity face videos with faithful and consistent details. See arrows for details.

 
 Table 2: Ablation study of variant networks.

Table 3:	Ablation	study	$\operatorname{on}$	optical	flow
estimator.					

Models	LPIPS↓	$\mathrm{IDS} \uparrow$	$\mathrm{AKD}\!\!\downarrow$
w/o CFA	0.1621	0.7970	8.9029
w/o KFN	0.1721	0.7773	9.1952
Full model	0.1619	0.7960	8.8182

Estimator	LPIPS↓	$\mathrm{IDS} \uparrow$	AKD↓
PWC-Net [41] GMFlow [49] (Ours)	$0.1623 \\ 0.1619$	$0.7957 \\ 0.7960$	8.7839 8.8182

#### 4.2 More Analysis

**Effectiveness of KFN.** We first investigate the effectiveness of the Kalman Filter Network (KFN). As shown in Table 2, removing KFN results in worse performance on LPIPS, IDS, and AKD. The results suggest the design of KFN is the key to promoting temporal consistency and identity preservation.

Effectiveness of CFA. Table 2 also shows that Cross-Frame Attention (CFA) can further improve the performance. Though not clearly reflected in the number, we further qualitatively show the effectiveness of KFN and CFA in the supplementary video. From the demo video, we can observe that 1) adopting KFN can ensure consistency in global style and maintain the global appearance of the recovered faces. 2) adding CFA can further coherently render local texture details (*e.g.*, hair), and suppress flicker.



Fig. 5: Comparison of temporal flicker. We select each frame's column (red lines) and show the changes across time. Image-based models (GFPGAN, CodeFormer, and RestoreFormer) have obvious discontinuity around the eyes and wrinkles, and general VSR methods leave artifacts behind. In contrast, by maintaining stable facial priors and aggregating temporal information, our method remarkably suppresses temporal jitters and promotes coherent local details.



**Fig. 6:** Illustration of predicted state and observed state through Decoder  $\mathcal{D}_Q$ . The teeth in red boxes illustrate a slight difference between  $\tilde{z}_t$  and  $\hat{z}_t^-$  when decoded in pixel space, suggesting the potential to supplement and fuse information to obtain a stable latent code.

Effectiveness of Various Flow Estimator. We compare models with different flow estimators in Table 3. The results suggest that the accuracy of estimated flows does not significantly affect the final performance. We conjecture this can be attributed to two factors: 1) Minor misalignment in pixel space can be reasonably diminished as the latent code is highly downsampled by a factor of  $32\times$ , at which level the latent representations are less sensitive to small spatial discrepancies present in the pixel space. 2) Other modules can compensate for the inaccuracy caused by flow estimators in a joint training fashion.

Analysis of Flickering. We extract a short vertical segment of pixels from each frame and stack them horizontally to visualize the jittering issues within the video. In particular, existing methods demonstrate clear jitter and flicker across time, while our method shows better temporal consistency. Fig. 5 demonstrates that other image-based models bring obvious jitters around eyes, and general VSR methods leave behind artifacts, while our method could remarkably suppress temporal jitters and promote coherent local details.



**Fig. 7: Identity similarity across frames.** Our method preserves the identity of input images and exhibits less fluctuation over time. CodeFormer results of two frames (green and red dashed line) exemplify an abrupt change of identity in the right figure, while our method maintains a stable identity both quantitatively and qualitatively.



Fig. 8: Qualitative comparison on different levels of degradation. Our KEEP maintains high-fidelity in various degradations.

Analysis of Latent Space. Since the true state  $z_t$  is unavailable, we indirectly analyze the predicted state  $\tilde{z}_t$  and observed state  $\hat{z}_t^-$  by decoding them to the pixel space through Decoder  $\mathcal{D}_Q$ . As shown in Fig. 6, the areas around teeth exhibit a slight difference, while most remaining parts in error maps show similar decoded results. This indicates that the predicted and observed states could supplement each other to obtain a more accurate estimation of  $z_t$ , which is where the power of Kalman filter lies.

**Identity Preservation.** We show the identity similarity across frames of one representative video clip in Fig. 7 As illustrated in the left figure, our method achieves better identity preservation and less identity jitter within the video, compared to the single-image methods. We also exemplify that the identity of CodeFormer results can change abruptly within several frames. The identity score of CodeFormer increases from 0.4552 to 0.6392, and turns down in later frames (See red curve), demonstrating great fluctuation across time. In contrast, our method maintains a stable identity both quantitatively and qualitatively.

Various Degradations. Fig. 9a shows that KEEP is consistently better than the compared methods across different difficulty levels. Fig. 8 demonstrates re-



(a) PNSR and IDS on different partitions. (b) Comparisons on non-frontal faces.

**Fig. 9:** (a) Our methods achieve consistently better performance on various levels of degradation. (b) While CodeFormer fails to restore eyes to these challenging cases, our method can still produce plausible facial elements.

sults on different levels of degradation. We draw the following observations. 1) Even GFPGAN and CodeFormer can restore plausible results on frames with mild degradation, the performance significantly deteriorates (e.g., eyes) upon heavier degradation. 2) Our method achieves better results by considering complementary information between adjacent frames and maintaining stable face priors. In addition, our method is appealing in handling heavy degradation.

**Non-Frontal-View Faces.** In Fig. 9b, our model shows enhanced performance on non-frontal faces by providing more stable face priors estimations. While the single-image model CodeFormer cannot recover the eyes, our KEEP is still able to show robustness to these challenging cases.

# 5 Conclusion

We present a novel framework, KEEP, aiming at resolving the challenges associated with facial detail and temporal consistency in video face restoration. The proposed method demonstrates a unique capability to maintain a stable face prior over time, which is achieved by Kalman filtering principles, where our approach recurrently incorporates information from previously restored frames to guide and regulate the restoration process of the current frame. Extensive experiments demonstrate the efficacy of KEEP in consistently capturing facial details across video frames and keeping the temporal stability of face videos.

## Acknowledgment

This study is supported under the RIE2020 Industry Alignment Fund Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

#### References

- Caballero, J., Ledig, C., Aitken, A., Acosta, A., Totz, J., Wang, Z., Shi, W.: Realtime video super-resolution with spatio-temporal networks and motion compensation. In: CVPR (2017)
- Chan, K.C., Wang, X., Xu, X., Gu, J., Loy, C.C.: GLEAN: Generative latent bank for large-factor image super-resolution. In: CVPR (2021)
- 3. Chan, K.C., Wang, X., Yu, K., Dong, C., Loy, C.C.: BasicVSR: The search for essential components in video super-resolution and beyond. In: CVPR (2021)
- 4. Chan, K.C., Xu, X., Wang, X., Gu, J., Loy, C.C.: GLEAN: Generative latent bank for large-factor image super-resolution and beyond. TPAMI (2022)
- 5. Chan, K.C., Zhou, S., Xu, X., Loy, C.C.: Basicvsr++: Improving video superresolution with enhanced propagation and alignment. In: CVPR (2022)
- Chan, K.C., Zhou, S., Xu, X., Loy, C.C.: Investigating tradeoffs in real-world video super-resolution. In: CVPR (2022)
- Chen, C., Li, X., Yang, L., Lin, X., Zhang, L., Wong, K.Y.K.: Progressive semanticaware style transformation for blind face restoration. In: CVPR (2021)
- Chen, Y., Tai, Y., Liu, X., Shen, C., Yang, J.: FSRNet: End-to-end learning face super-resolution with facial priors. In: CVPR (2018)
- Deng, J., Guo, J., Xue, N., Zafeiriou, S.: ArcFace: Additive angular margin loss for deep face recognition. In: CVPR (2019)
- Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: CVPR (2021)
- 11. Fox, G., Tewari, A., Elgharib, M., Theobalt, C.: StyleVideoGAN: A temporal generative model using a pretrained stylegan. In: BMVC (2021)
- Gu, J., Shen, Y., Zhou, B.: Image processing using multi-code gan prior. In: CVPR (2020)
- Gu, Y., Wang, X., Xie, L., Dong, C., Li, G., Shan, Y., Cheng, M.M.: VQFR: Blind face restoration with vector-quantized dictionary and parallel decoder. In: ECCV (2022)
- 14. Haris, M., Shakhnarovich, G., Ukita, N.: Recurrent back-projection network for video super-resolution. In: CVPR (2019)
- He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing humanlevel performance on imagenet classification. In: ICCV (2015)
- 16. Huang, Y., Wang, W., Wang, L.: Bidirectional recurrent convolutional networks for multi-frame super-resolution. In: NeurIPS (2015)
- 17. Isobe, T., Jia, X., Gu, S., Li, S., Wang, S., Tian, Q.: Video super-resolution with recurrent structure-detail network. In: ECCV (2020)
- Isobe, T., Zhu, F., Wang, S.: Revisiting temporal modeling for video superresolution. In: BMVC (2020)
- 19. Kalman, R.E.: A new approach to linear filtering and prediction problems (1960)
- Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: CVPR (2019)

- 16 R. Feng et al.
- 21. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: CVPR (2020)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
- 23. Lei, C., Ren, X., Zhang, Z., Chen, Q.: Blind video deflickering by neural filtering with a flawed atlas. In: CVPR (2023)
- Lei, C., Xing, Y., Chen, Q.: Blind video temporal consistency via deep video prior. In: NeurIPS (2020)
- 25. Li, S., He, F., Du, B., Zhang, L., Xu, Y., Tao, D.: Fast spatio-temporal residual network for video super-resolution. In: CVPR (2019)
- Li, X., Chen, C., Zhou, S., Lin, X., Zuo, W., Zhang, L.: Blind face restoration via deep multi-scale component dictionaries. In: ECCV (2020)
- Li, X., Li, W., Ren, D., Zhang, H., Wang, M., Zuo, W.: Enhanced blind face restoration with multi-exemplar images and adaptive spatial feature fusion. In: CVPR (2020)
- Li, X., Liu, M., Ye, Y., Zuo, W., Lin, L., Yang, R.: Learning warped guidance for blind face restoration. In: ECCV (2018)
- Liang, J., Cao, J., Fan, Y., Zhang, K., Ranjan, R., Li, Y., Timofte, R., Van Gool, L.: Vrt: A video restoration transformer. IEEE Transactions on Image Processing (2024)
- Liang, J., Fan, Y., Xiang, X., Ranjan, R., Ilg, E., Green, S., Cao, J., Zhang, K., Timofte, R., Gool, L.V.: Recurrent video restoration transformer with guided deformable attention. NeurIPS (2022)
- Liu, C., Yang, H., Fu, J., Qian, X.: Learning trajectory-aware transformer for video super-resolution. In: CVPR (2022)
- 32. Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H.: Video swin transformer. In: CVPR (2022)
- 33. Menon, S., Damian, A., Hu, S., Ravi, N., Rudin, C.: Pulse: Self-supervised photo upsampling via latent space exploration of generative models. In: CVPR (2020)
- Nagrani, A., Chung, J.S., Zisserman, A.: Voxceleb: A large-scale speaker identification dataset. In: Interspeech (2017)
- Ouyang, H., Wang, Q., Xiao, Y., Bai, Q., Zhang, J., Zheng, K., Zhou, X., Chen, Q., Shen, Y.: Codef: Content deformation fields for temporally consistent video processing. In: CVPR (2024)
- Pan, J., Bai, H., Dong, J., Zhang, J., Tang, J.: Deep blind video super-resolution. In: ICCV (2021)
- Pan, X., Zhan, X., Dai, B., Lin, D., Loy, C.C., Luo, P.: Exploiting deep generative prior for versatile image restoration and manipulation. TPAMI 44(11), 7474–7489 (2021)
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in pytorch (2017)
- Revach, G., Shlezinger, N., Ni, X., Escoriza, A.L., Van Sloun, R.J., Eldar, Y.C.: Kalmannet: Neural network aided kalman filtering for partially known dynamics. IEEE Transactions on Signal Processing (2022)
- 40. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022)
- 41. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: CVPR (2018)
- 42. Tian, Y., Zhang, Y., Fu, Y., Xu, C.: TDAN: Temporally-deformable alignment network for video super-resolution. In: CVPR (2020)

- 43. Tzaban, R., Mokady, R., Gal, R., Bermano, A., Cohen-Or, D.: Stitch it in time: Gan-based facial editing of real videos. In: SIGGRAPH Asia (2022)
- 44. Wang, X., Chan, K.C., Yu, K., Dong, C., Change Loy, C.: EDVR: Video restoration with enhanced deformable convolutional networks. In: CVPRW (2019)
- 45. Wang, X., Li, Y., Zhang, H., Shan, Y.: Towards real-world blind face restoration with generative facial prior. In: CVPR (2021)
- 46. Wang, Z., Zhang, J., Chen, R., Wang, W., Luo, P.: Restoreformer: High-quality blind face restoration from undegraded key-value pairs. In: CVPR (2022)
- 47. Wu, J.Z., Ge, Y., Wang, X., Lei, S.W., Gu, Y., Shi, Y., Hsu, W., Shan, Y., Qie, X., Shou, M.Z.: Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In: ICCV (2023)
- 48. Xie, L., Wang, X., Zhang, H., Dong, C., Shan, Y.: VFHQ: A high-quality dataset and benchmark for video face super-resolution. In: CVPR (2022)
- Xu, H., Zhang, J., Cai, J., Rezatofighi, H., Tao, D.: Gmflow: Learning optical flow via global matching. In: CVPR (2022)
- Xu, Y., AlBahar, B., Huang, J.B.: Temporally consistent semantic video editing. In: ECCV (2022)
- 51. Yang, S., Zhou, Y., Liu, Z., , Loy, C.C.: Rerender a video: Zero-shot text-guided video-to-video translation. In: SIGGRAPH Asia (2023)
- 52. Yang, T., Ren, P., Xie, X., Zhang, L.: Gan prior embedded network for blind face restoration in the wild. In: CVPR (2021)
- 53. Yu, X., Fernando, B., Ghanem, B., Porikli, F., Hartley, R.: Face super-resolution guided by facial component heatmaps. In: ECCV (2018)
- 54. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018)
- Zhou, L., Luo, Z., Shen, T., Zhang, J., Zhen, M., Yao, Y., Fang, T., Quan, L.: KFNet: Learning temporal camera relocalization using Kalman filtering. In: CVPR (2020)
- 56. Zhou, S., Chan, K.C., Li, C., Loy, C.C.: Towards robust blind face restoration with codebook lookup transformer. In: NeurIPS (2022)