

Beyond Wide-Angle Images: Structure-to-Detail Video Portrait Correction via Unsupervised Spatiotemporal Adaptation

Wenbo Nie^{1,2}, Lang Nie^{3†}, Chunyu Lin^{1,2*}, Jingwen Chen^{1,2}, Ke Xing^{1,2}, Jiyuan Wang^{1,2}, Kang Liao⁴

¹Institute of Information Science, Beijing Jiaotong University, Beijing, China

²Visual Intelligence + X International Joint Laboratory of the Ministry of Education, Beijing, China

³Chongqing University of Posts and Telecommunications, Chongqing, China

⁴Nanyang Technological University, Singapore

Abstract

Wide-angle cameras, despite their popularity for content creation, suffer from distortion-induced facial stretching—especially at the edge of the lens—which degrades visual appeal. To address this issue, we propose a structure-to-detail portrait correction model named **ImagePC**. It integrates the long-range awareness of the transformer and multi-step denoising of diffusion models into a unified framework, achieving global structural robustness and local detail refinement. Besides, considering the high cost of obtaining video labels, we then repurpose ImagePC for unlabeled wide-angle videos (termed **VideoPC**), by spatiotemporal diffusion adaptation with spatial consistency and temporal smoothness constraints. For the former, we encourage the denoised image to approximate pseudo labels following the wide-angle distortion distribution pattern, while for the latter, we derive rectification trajectories with backward optical flows and smooth them. Compared with ImagePC, VideoPC maintains high-quality facial corrections in space and mitigates the potential temporal shakes sequentially in *blind scenarios*. Finally, to establish an evaluation benchmark and train the framework, we establish a video portrait dataset with a large diversity in the number of people, lighting conditions, and background. Experiments demonstrate that the proposed methods outperform existing solutions quantitatively and qualitatively, contributing to high-fidelity wide-angle videos with stable and natural portraits.

Introduction

With the development of the self-media and videography industries, wide-angle lenses have become increasingly favored for the capability of capturing expansive scenes. However, such lenses inevitably introduce geometric distortions, particularly pronounced at the lens boundaries, resulting in background straight line curvatures and facial feature deformations in still images and video recordings.

Traditional wide-angle portrait correction methods typically require precise camera parameters (*e.g.*, focal length (Shih, Lai, and Liang 2019)) as a prerequisite, followed by stereographic projection, face and line detection, and energy optimization to achieve geometric correction (Shih,

Lai, and Liang 2019). Such a pipeline is complex in process and inapplicable when the relevant camera parameters are unknown. The same goes for video portrait correction work (Lai et al. 2022). In contrast, learning-based solutions eliminate this issue (Tan et al. 2021) by directly learning the spatial mapping from wide-angle images to rectification flows in a supervised manner. However, these methods often exhibit noticeable temporal shakes when applied to videos. Moreover, there is currently a lack of datasets for video portrait correction due to the significantly higher cost of labeling video data. To this end, we take a pioneering step in achieving video portrait correction without video annotations. First of all, we present a structure-to-detail portrait correction model, named **ImagePC**, to generate high-quality rectification flows, which integrates the advantages of both Transformer and diffusion models. The Transformer model establishes long-range geometric dependencies by capturing spatial global features across entire images, providing structural guidance for the subsequent diffusion process. The diffusion model refines high-fidelity flow patterns through iterative denoising guided by Transformer-derived structural information. This hybrid architecture enables structure-to-detail pixel-wise correction, where structural priors from the Transformer and detail-generation capacity from the diffusion model are synergistically combined.

More importantly, we propose an unsupervised spatiotemporal adaption approach to enable ImagePC the capability to rectify video sequences stably. The repurposed model, which we term **VideoPC**, addresses the issue of temporal shakes without requiring video correction labels. Concretely, we first use the pre-trained ImagePC to generate rectification flows for each video frame as pseudo-labels. When ImagePC is adapted to unseen wide-angle videos, these pseudo-labels can guide the denoising process, ensuring spatial consistency in facial corrections with the wide-angle distortion distribution prior. To enforce temporal smoothness, we derive rectification trajectories across sequential frames using backward optical flows and design a temporal smoothness constraint to mitigate sequential jitters. With these spatiotemporal constraints, VideoPC achieves stable wide-angle video correction in blind scenarios that not only preserves high-fidelity facial details but also frees from labor-intensive video annotations.

To establish an evaluation benchmark and train the pro-

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

*Corresponding author;

† project lead

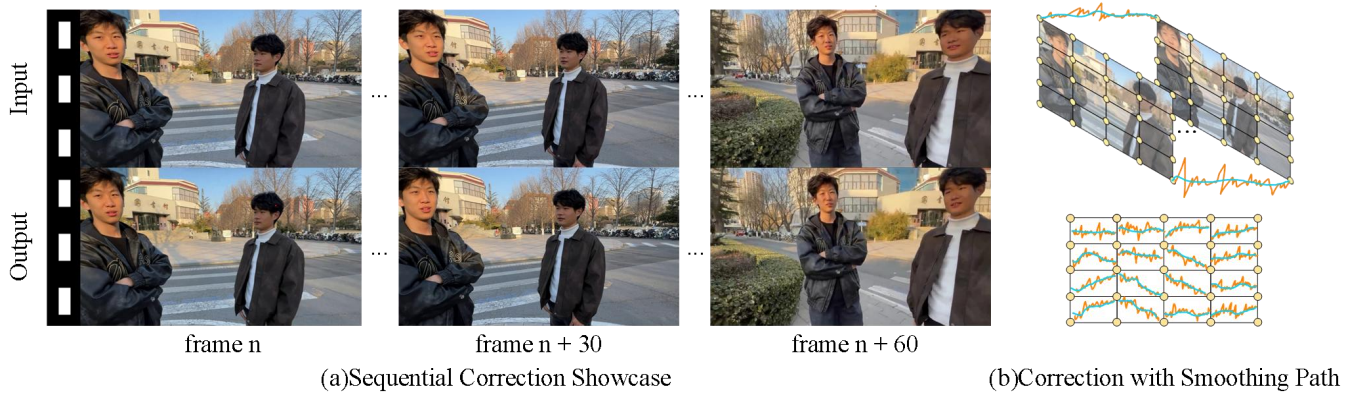


Figure 1: Our wide-angle video portrait correction results. (a) The first row shows a wide-angle phone video with edge distortion, corrected in the second row using our method while preserving spatiotemporal consistency. (b) Video correction introduces temporal shake (orange line in the figure), while our method smooths the trajectory (blue line).

posed video portrait correction model, we construct a wide-angle video dataset with a wide diversity in scene, camera, and number of people. Finally, we conduct extensive experiments about portrait correction in both image and video, demonstrating our superiority over other solutions. The contributions center around:

- We propose a structure-to-detail portrait correction model, named **ImagePC**, which integrates long-range dependences of transformer and iterative refinements of diffusion models for global-to-local rectification.
- We design an unsupervised spatiotemporal adaptation framework to transfer **ImagePC** to **VideoPC** without labeled videos. It combines spatio-temporal optical flows to track correction trajectories, establishing spatial consistency and temporal smoothness constraints.
- We present a video portrait dataset with wide-angle cameras and conduct extensive experiments to validate our superior performance over existing solutions.

Related Work

Portrait Correction

Portrait correction in wide-angle images is a significant research area. Traditional methods, such as Shih et al.’s (Shih, Lai, and Liang 2019) content-aware warping, adjust distortions by analyzing image content and camera parameters. However, these methods often rely on camera parameters or complex inputs and are difficult to extend to video processing. Stereographic projection methods (Shih, Lai, and Liang 2019) can preserve local conformality but struggle with temporal consistency in videos.

Deep learning methods (Tan et al. 2021), Zhu et al. (Zhu et al. 2022) developed a semi-supervised Transformer to reduce annotation costs. Nie et al.’s method CoupledTPS (Nie et al. 2024a), iteratively couples limited TPS models with a warping flow to precisely correct facial features.

Existing deep learning methods, designed for static images, fail to address the temporal dynamics of videos due to expensive frame-by-frame annotation and limited datasets.

Lai et al. (Lai et al. 2022) proposed a temporal consistency optimization method using spherical projection and energy minimization, but it struggles to integrate global video information with facial correction details and lacks adaptability in complex environments. To address this issue, we propose learning video correction principles from image networks.

Video stabilization

Traditional video stabilization techniques primarily stabilize the video by smoothing the feature trajectories between multiple frames or adjacent frames (Wang et al. 2013; Lee et al. 2009; Rublee et al. 2011). Subsequently, grid-based methods (Wang et al. 2013; Liu et al. 2013, 2016; Yan et al. 2024) and optical flow (Liu et al. 2014); (James, Jain, and Rajwade 2022) have also been applied to represent motion, becoming key tools for stabilizing videos. With the rapid development of deep learning, deep learning methods have been applied to video stabilization technology (Zhao and Ling 2020; Yu and Ramamoorthi 2019; Zhao et al. 2023; Peng et al. 2024; Yang et al. 2006), where the original video frames are input, and stabilized video frames are directly output. Additional methods, including unsupervised online video stitching (Nie et al. 2024b), and meta-learning-based approaches to improve full-frame stabilization (Ali et al. 2024), have been explored to enhance stability. These approaches, ranging from optical flow to deep learning models, provide valuable solutions for addressing video temporal shake in video portrait correction.

Diffusion Models

Diffusion models (Yang et al. 2024) excel in generative tasks, extending from image generation to video-related applications. In video generation, early advancements include Video Diffusion Models (VDM) (Ho et al. 2022b), which adapted the 2D U-Net architecture from image DMs to a 3D U-Net. To enhance temporal resolution, Ho et al. (Ho et al. 2022a) introduced a cascade of video super-resolution models.

More recent work focuses on video editing and enhancing temporal consistency. For example, Tune-A-Video (Wu et al. 2023) fine-tunes models on videos for coherent outputs, while FateZero (Qi et al. 2023) ensures temporal consistency using attention-based blending techniques. Similarly, Rerender-A-Video (Yang et al. 2023) and VideoControlNet (Hu and Xu 2023) introduce pixel-aware latent fusion and motion-aware adjustments to maintain visual consistency during editing.

Methodology

Overview

The overview of our framework is shown in Fig. 2, consisting of an image correction model and a video correction model. In the image model, ImagePC tackles the distorted image I_s , producing precise optical flows using a structure-to-detail rectification architecture. The predicted flow is then applied to generate a geometrically corrected image \hat{I}_s .

In the video model, directly adopting the image model to correct the video frames can easily yield temporal shakes due to temporally unsmooth warps. But we could take the correction flows from the image model as the reference and transfer the image model to the video domain through spatiotemporal adaptation.

ImagePC

As illustrated in Fig. 2(a,b), our ImagePC is a framework for high-quality portrait correction by predicting precise optical flow fields to rectify facial distortions.

Structure-to-Detail Model In the first stage of ImagePC, we focus on capturing the global geometric structure of the distorted input image I_s . The Transformer network (Zhu et al. 2022), in Fig. 2(a), with its advantage in capturing long-range information, extracts a set of intermediate layer features h that effectively reflect the global geometric structure of the distorted input image. These features h serve as a structural prior, guiding the diffusion model in the next stage to accurately align the facial geometry while preserving spatial coherence.

Building on the global structure from the first stage, the next stage, as shown in Fig. 2(b), refines local details using a diffusion-based approach. To ensure that the correction process primarily focuses on the area of the face, we use InsightFace (Ren et al. 2023; Guo et al. 2021; Guo, Zhao, and Wang 2024) to process the original facial image I_s , allowing accurate face detection and keypoint localization. Based on these detection results, a binary mask M is generated, which clearly delineates the facial region to enable better portrait perception. This mask provides spatial guidance for subsequent processing and improves the quality of the correction. Leveraging the detail-handling capability of diffusion models, we form the conditioning input for ImagePC by concatenating M , I_s , and h :

$$x_{in} = \text{cat}[M, I_s, h]. \quad (1)$$

Our approach leverages the iterative refinement capabilities of Denoising Diffusion Implicit Models (DDIM) (Song,

Meng, and Ermon 2022), enabling iterative data generation through noise-adding and denoising processes.

Adopting this concept, ImagePC redefines the diffusion process to operate directly on optical flow fields.

Unlike traditional image-to-image methods that directly generate a corrected facial image, we modify the diffusion network in this stage to predict an optical flow field F^{corr} . The repurposed image-to-flow pipeline not only ensures that the corrected content totally loyal to the original content, but also enables the derivation of video correction trajectories.

With high-quality correction flows, the wide-angle rectification can be formulated as:

$$\hat{I}_s = \mathcal{W}(F^{corr}, I_s). \quad (2)$$

The warping operation \mathcal{W} realigns the pixels in I_s according to the predicted flow, resulting in a natural facial appearance.

Objective Function The model is optimized using a comprehensive loss function:

$$\mathcal{L}_{image} = \lambda_1 \mathcal{L}_{mask} + \lambda_2 \mathcal{L}_{photo} + \lambda_3 \mathcal{L}_{flow}, \quad (3)$$

where \mathcal{L}_{flow} ensures optical flow accuracy, \mathcal{L}_{photo} enforces image quality, and \mathcal{L}_{mask} enhances facial edge details, with λ_1 , λ_2 , and λ_3 as weighting coefficients.

Optical Flow Loss \mathcal{L}_{flow} : It directly measures the motion-level difference between the predicted optical flow and the ground truth optical flow as:

$$\mathcal{L}_{flow} = \frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W W (F^{Corr}(h, w) - F_{gt}(h, w))^2. \quad (4)$$

Image Loss \mathcal{L}_{photo} : It evaluates the image-level difference between the corrected image and the ground truth image as:

$$\mathcal{L}_{photo} = \frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W W (\hat{I}_s(h, w) - I_{gt}(h, w))^2. \quad (5)$$

Mask-based Sobel Loss \mathcal{L}_{mask} : We use the Sobel kernels (G_x and G_y) to extract edge information and calculate the edge loss within the facial region mask as:

$$\mathcal{L}_{mask} = [|G_x(F^{Corr}) - G_x(F_{gt})| + |G_y(F^{Corr}) - G_y(F_{gt})|] \cdot M. \quad (6)$$

VideoPC

Videos are dynamic sequences composed of continuous multi-frame images $\mathcal{I} = \{\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_n\}$. To achieve unsupervised wide-angle video correction, we use pre-trained ImagePC to generate corrected optical flow for each frame as pseudo-labels. However, directly applying these pseudo-labels to the video sequences will result in severe temporal shakes, significantly decreasing the video’s quality and visual experience. To address this issue, we further propose to transfer ImagePC to VideoPC with the adaptability in wide-angle video scenarios.

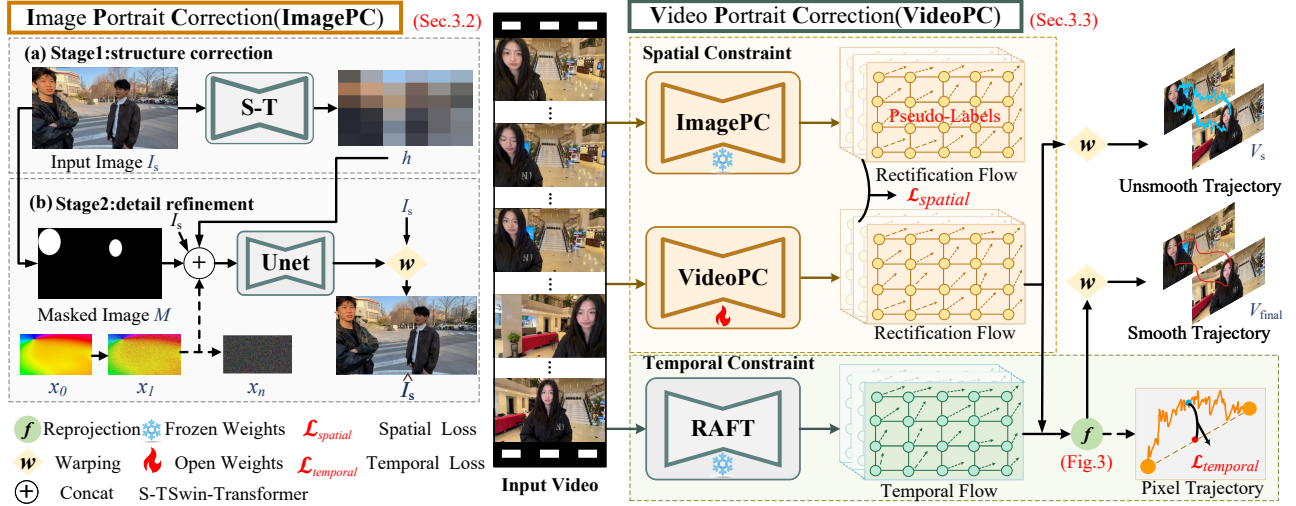


Figure 2: The proposed framework consists of two models. (1) ImagePC: We design a structure-to-detail architecture to generate high-quality rectification flows for each individual frame. (2) VideoPC: We adapt ImagePC from the image to video domain with spatial consistency and temporal smoothness constraints.

Correction Trajectory Our method aims to control the comprehensive motion between consecutive frames. To represent such motions, we derive correction trajectories that are built on inter-frame and intra-frame optical flows. For inter-frame motions, we use the RAFT (Zhang et al. 2024) to capture the temporal flows ($f_{t \rightarrow t+1}$ and $f_{t+1 \rightarrow t}$) from wide-angle video frames (I_t and I_{t+1}). For intra-frame motions, we leverage the VideoPC model (the same architecture as ImagePC) to predict the rectification flows (F_t^{corr}). The remaining key question is how to combine these flows into our desired correction trajectories.

Based on Forward Optical Flow For forward optical flow (i.e., F_t^{corr} denotes the correspondences from wide-angle images to rectified images), we can identify the differences of corresponding points within the consecutive rectified results as the following formula:

$$r(t+1) = f_{t+1 \rightarrow t} + F_t^{Corr} (G + f_{t+1 \rightarrow t}) - F_{t+1}^{Corr}, \quad (7)$$

where $G \in \mathbb{R}^{2 \times H \times W}$ denotes pixel-wise grid coordinates. The latter part is the position in the (t+1)-th frame (i.e., $F_{t+1}^{Corr} + G$), while the upper is the corresponding position in the t -th frame (i.e., $f_{t+1 \rightarrow t} + F_t^{Corr} (G + f_{t+1 \rightarrow t}) + G$).

Based on Backward Optical Flow However, in the actual portrait distortion correction process, employing forward optical flow to warp is prone to generating invalid spatial holes. Therefore, all current optical flow-based correction methods employ backward optical flows, that is, F_t^{corr} denotes the mapping flow from rectified results to original wide-angle images. The difference in optical flow direction makes the relative motion described in Eq. 7 unsuitable for practical backward optical flow transformations. To compare the rectified results of the t -th and $(t+1)$ -th frames, we need to identify the differences in corresponding positions in the

input image for the same grid location. However, directly comparing the two optical flows F_t^{corr} and F_{t+1}^{corr} is not feasible, as the two backward optical flows do not correspond to the same points. Therefore, we need to leverage $f_{t \rightarrow t+1}$ to correct the temporal misalignment.

The position difference with backward flows is derived as follows:

$$r(t+1) = f_{t \rightarrow t+1} (G + F_t^{Corr}) + F_t^{Corr} - F_{t+1}^{Corr}, \quad (8)$$

where Fig. 3 provide a visual explanation about Eq. 8.

It can be further concatenated in temporal order from the initial moment to obtain the position at each moment as follows:

$$R(t) = r(1) + r(2) + \dots + r(t), \quad (9)$$

where $r(1)$ is defined as a zero matrix. The final correction trajectory can be obtained by sequentially chaining $R(1), R(2), \dots, R(t), \dots, R(N)$ over time.

The produced trajectory contains comprehensive motions, including the shakes caused by unsmooth rectification flows and the inherent jitters in original wide-angle videos.

Spatiotemporal Adaptation To achieve unsupervised video correction, we design our objective function \mathcal{L}_{video} concerning two aspects: spatial consistency and temporal smoothness.

$$\mathcal{L}_{video} = \mathcal{L}_{spatial} + \lambda \mathcal{L}_{temporal}. \quad (10)$$

For the temporal smoothness, we encourage the correction trajectory from VideoPC to be as smooth as possible. In particular, for three consecutive frames, we require the trajectory position of the middle frame to lie between the trajectory positions of the preceding and following frames:

$$\mathcal{L}_{temporal} = \|R(t+1) + R(t-1) - 2R(t)\|. \quad (11)$$

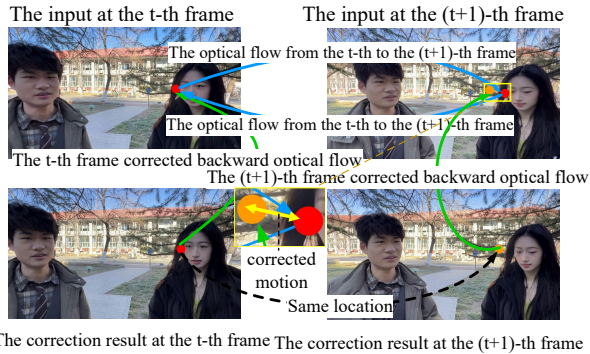


Figure 3: Illustration of the computation of $r(t+1)$ in Eq. 8, depicting the integration of RAFT optical flow $f_{t \rightarrow t+1}$, corrected mesh G , and correction flows F_t^{Corr} and F_{t+1}^{Corr} for temporal stabilization.

For spatial consistency, we leverage the pre-trained ImagePC model to ensure the quality of facial distortion correction within each frame. Concretely, we use pseudo-labels generated by ImagePC to provide spatial supervision for VideoPC, including the optical flow constraint \mathcal{L}_{flow} (Eq. 4) and mask-based edge constraint \mathcal{L}_{mask} (Eq. 6). In the transfer process of facial correction capabilities from the image to video domain, these components prevent the structural integrity and edge details of the facial regions from severe degradation.

Data preparation

Due to the lack of publicly available wide-angle video portrait datasets, we introduce a new dataset containing 136 clips, ranging from 5 seconds to 90 seconds in duration. Our dataset includes clips captured by seven smartphones covering a wide variety of scenarios and encompass diverse scenarios, including content captured by self-media creators and film clips extracted from cinematic works. Detailed information is available in the supplementary materials. To authentically replicate real-world recording conditions, we deliberately introduced artificial shaking to simulate the instability inherent in handheld video capture. Fig. 4 presents some examples of our wide-angle video portrait dataset.

Experiment

Dataset and Implementation Detail

Our methodology builds upon two core datasets: the established wide-angle image dataset (Tan et al. 2021) and our newly captured video dataset.

Image Data In our experiments, we use the dataset from Tan et al. (Tan et al. 2021), which consists of 5133 training images and 129 test images.

Video Data For both training and testing phases, we utilized the dataset described in Sec. . The training set comprised 7,134 frames that comprehensively represent diverse capture devices, environmental conditions, and scene types.



Figure 4: Our wide-angle video face dataset encompasses a broad range of scenes and showcases significant diversity in individuals. We present some examples from our dataset.

Training Process of ImagePC Our framework employs a two-stage training protocol. In the first stage, we utilize the Adam optimizer (Kingma and Ba 2017) with an initial learning rate of 1×10^{-4} , training the model for 200 epochs. The subsequent texture refinement stage leverages high-dimensional features extracted from the initial phase. This feature, combined with image I_s and mask M , acts as the conditional guidance for the diffusion model to enhance fine-grained texture details. This stage implements a learning rate of 2×10^{-4} , with optimization conducted over 500,000 iterations.

Training Process of VideoPC For VideoPC, we fine-tune the ImagePC model to maintain correction efficacy while smoothing temporal jitter. Training employs the Adam optimizer (Kingma and Ba 2017) with a learning rate of 2×10^{-4} over 300,000 iterations. The entire training process is conducted on a **single RTX 4090**.

Evaluation Metric

Following Tan et al.’s work (Tan et al. 2021), we evaluate geometric correction performance using two established metrics: **LineAcc** and **ShapeAcc**, where higher values indicate better alignment. We also employ Jin et al.’s **Stability Score** (Avg, Trans, Rot) (Choi and Kweon 2020), which uses homography and FFT, to measure VideoPC’s shake reduction, with higher scores denoting smoother videos.

Comparative Result

We evaluate our approach against existing methods across image and video modalities, providing comprehensive quantitative and qualitative results.

Method	Reference	LineAcc	ShapeAcc
Shih et al.	TOG2019	66.143	97.253
Tan et al.	CVPR2021	66.784	97.490
Zhu et al.	CVPR2022	66.825	97.491
DualPriors	ECCV2024	67.304	99.012
CoupledTPS	TPAMI2024	66.808	97.500
MOWA	TPAMI2025	-	97.475
ImagePC (Ours)	-	66.898	97.508

Table 1: Quantitative comparison of the proposed ImagePC with other portrait correction methods. Higher LineAcc and ShapeAcc indicate better straight-line and face correction performance, respectively.




Figure 5: Comparison of ImagePC with existing methods on wide-angle distorted images. ImagePC achieves superior restoration of natural facial features and reduced edge distortions compared to state-of-the-art alternatives.

Image Comparison As shown in Tab. 1, our ImagePC achieves the second-best performance across LineAcc and ShapeAcc metrics, surpassed only by DualPriors(Yao et al. 2024). This shows that our image correction model can retain the geometric integrity of straight edges and preserve facial shape well. Furthermore, in the qualitative comparison presented in Fig. 5, we can further observe that ImagePC outperforms DualPriors(Yao et al. 2024) in the aforementioned aspects. This undoubtedly demonstrates that, despite DualPriors(Yao et al. 2024) achieving state-of-the-art performance on certain metrics, their approach exhibits noticeable abrupt artifacts at stitching boundaries and inconsistent head-to-body proportions in the images. This may be attributed to their method relying on a parallel network that separately rectifies the background and portrait before stitching them into a single image, with blank areas repaired using LaMa(Suvorov et al. 2021). In contrast, Fig. 5(c) clearly illustrates that our method generates more natural and refined character appearances, highlighting the superiority of ImagePC.

Video Comparison In the video domain, our approach not only corrects individual frames but also prioritizes inter-frame continuity to mitigate severe visual jitter. As demonstrated in Table 2, VideoPC exhibits superior performance in handling rotational distortions while maintaining optimal stability. In Fig. 6 (a), the red arrows and boxes highlight the unnatural or failed corrections in previous methods when processing video sequences, and these errors typically result in video discontinuity and poor visual perception. In contrast, our framework effectively corrects distortions while preserving the spatiotemporal consistency across the entire video sequence, underscoring the superiority of VideoPC. Additional video results can be found in the supplementary materials.

Video Correction Trajectory In Fig. 6 (b), we arbitrarily select a pixel from video frames and plot its trajectory, revealing that face distortion correction introduces noticeable jitter. Our smoothing process effectively mitigates this warping shake, producing smoother and more stable trajectories

compared to CoupledTPS (Nie et al. 2024a) and DualPriors+LAMA (Yao et al. 2024; Suvorov et al. 2021), as shown in an enlarged view. This results in corrected video frames that are more faithful to the original, significantly reducing jitter and enhancing the viewing experience.

Method	Average	Translational	Rotational
CoupledTPS	0.9811	0.9785	0.9836
DualPriors	0.9477	0.9290	0.9664
VideoPC(Ours)	0.9886	0.9821	0.9951

Table 2: Quantitative stability metrics (Translation/Rotation) for video stabilization models.

Notably, the impressive performance of our proposed method **ImagePC** and **VideoPC** is primarily attributed to the two-stage structural design. This framework combines the transformer’s capacity for long-range dependency modeling with the diffusion model’s multi-step denoising process, achieving global structural consistency and local refinement. By integrating these complementary mechanisms within a unified architecture, our approach achieves remarkable generation quality while maintaining spatial-temporal consistency across various distortion scenarios.

Ablation Study

In order to demonstrate the effectiveness of our approach and verify the impact of different components, we conducted a series of ablation experiments. In particular, we focus on different diffusion conditions, collaborative architecture design, and video correction trajectories.

Structure-to-Detail Model and Conditional Guidance: Specifically, to independently analyze the impact of each component on correction quality, we designed four ablation scenarios: 1) Transformer Only: The Transformer predicts optical flow without diffusion, testing its global structure modeling. 2) Diffusion with Condition (I_s): The diffusion model exclusively uses the distorted image I_s to predict optical flow, assessing its baseline detail correction. 3) Diffu-

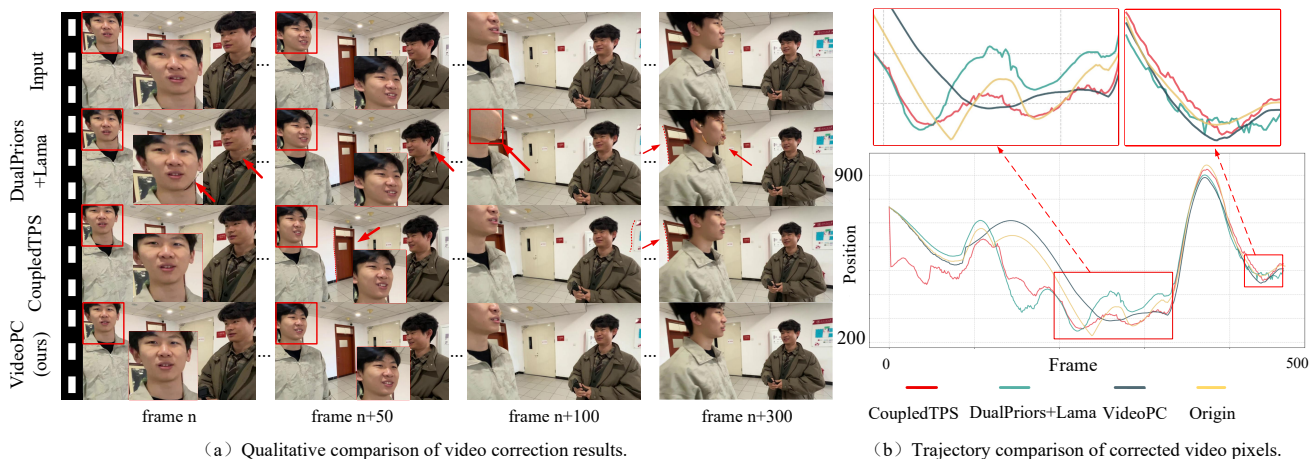


Figure 6: Comparison with existing portrait correction methods in wide-angle videos.

sion with Condition (I_s, M) : A mask is applied that marks the face region in the distorted wide-angle image, serving as the conditioning input to guide the model’s generation process. 4) Transformer + Diffusion with Condition (I_s, M) : The Transformer provides a feature vector h concatenated with I_s and M to guide the diffusion model, evaluating their combined structural and detail refinement capabilities.

Transformer	Diffusion	M	I_s	LineAcc	ShapeAcc
✓				66.617	95.533
	✓		✓	66.249	97.465
	✓	✓	✓	66.687	97.488
✓	✓	✓	✓	66.898	97.508

Table 3: Ablation experiments on the hybrid architecture and different conditions.

Tab. 3 presents the experimental results, clearly demonstrating the impact of each module. Specifically, when the Transformer model is integrated with the diffusion model, the scores of LineAcc and ShapeAcc achieve their peak values. These findings confirm our hypothesis that the joint transformer-diffusion architecture can significantly improve the performance of wide-angle lens distortion correction. Furthermore, the progressive incorporation of conditioning information highlights how each component contributes to the robustness and accuracy of portrait stability.

Method	Stability	LineAcc	ShapeAcc
ImagePC	0.9325	66.898	97.508
VideoPC	0.9886	66.692	97.480

Table 4: Ablation experiments comparing ImagePC and VideoPC on video stabilization and correction.

Spatiotemporal Smoothing on Video Correction: Our VideoPC model is a fine-tuned extension of the ImagePC model, specifically designed to address the challenges of video correction. Meanwhile, as mentioned before, the key

difference between video and photo correction is the spatiotemporal consistency. Therefore, in this section, we directly compare the video results generated by VideoPC and ImagePC to evaluate whether VideoPC alleviates the warping shake caused by the correction process.

In fact, to mitigate the video shaking, some of the original correction of ImagePC will inevitably be compromised. In our experiments, we strive to balance correction and smoothing. This balance ensures that while we reduce the warping shake, we do not sacrifice too much of the correction accuracy. Tab. 4 presents the results. ImagePC achieves excellent correction accuracy, while VideoPC, with spatiotemporal smoothing, increases Stability to 0.9886 (a 6.0% improvement), significantly reducing jitter. The minor decrease in Line/ShapeAcc is statistically insignificant ($p > 0.05$, t-test), indicating that VideoPC effectively balances correction accuracy and inter-frame smoothness.

For further video results and more detailed analysis, readers can refer to the supplementary materials.

Conclusion

In this work, we introduce a novel approach to address facial distortion correction in both images and videos, integrating of structure and detail for effective portrait distortion correction. Our method represents the first deep learning-based solution for wide-angle video portrait correction, offering a significant advancement over traditional approaches. By using image-level labels to learn the task of unlabeled wide-angle video correction, we drastically reduce the annotation burden and costs. Additionally, our proposed unsupervised framework, VideoPC, enhances spatiotemporal consistency while maintaining a smooth balance between correction and stabilization. We validate our method using a newly created wide-angle video dataset, demonstrating superior performance compared to existing approaches, showcasing its potential for wide-scale, automated, and accurate correction of wide-angle portrait distortion in still images and dynamic video sequences.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (NSFC) under Grant (62573039,U2441242).

References

- Ali, M. K.; Im, E. W.; Kim, D.; and Kim, T. H. 2024. Harnessing Meta-Learning for Improving Full-Frame Video Stabilization. *arXiv:2403.03662*.
- Choi, J.; and Kweon, I. S. 2020. Deep Iterative Frame Interpolation for Full-Frame Video Stabilization. *ACM Transactions on Graphics*, 39(1).
- Guo, J.; Deng, J.; Lattas, A.; and Zafeiriou, S. 2021. Sample and Computation Redistribution for Efficient Face Detection. *arXiv preprint arXiv:2105.04714*.
- Guo, K.; Zhao, C.; and Wang, J. 2024. A fast mask synthesis method for face recognition. *Visual Intelligence*, 2(1): 25.
- Ho, J.; Chan, W.; Saharia, C.; Whang, J.; Gao, R.; Gritsenko, A.; Kingma, D. P.; Poole, B.; Norouzi, M.; Fleet, D. J.; and Salimans, T. 2022a. Imagen Video: High Definition Video Generation with Diffusion Models. *arXiv:2210.02303*.
- Ho, J.; Salimans, T.; Gritsenko, A.; Chan, W.; Norouzi, M.; and Fleet, D. J. 2022b. Video Diffusion Models. *arXiv:2204.03458*.
- Hu, Z.; and Xu, D. 2023. VideoControlNet: A Motion-Guided Video-to-Video Translation Framework by Using Diffusion Model with ControlNet. *arXiv:2307.14073*.
- James, J. G.; Jain, D.; and Rajwade, A. 2022. GlobalFlowNet: Video Stabilization using Deep Distilled Global Motion Estimates. *arXiv:2210.13769*.
- Kingma, D. P.; and Ba, J. 2017. Adam: A Method for Stochastic Optimization. *arXiv:1412.6980*.
- Lai, W.-S.; Shih, Y.; Liang, C.-K.; and Yang, M.-H. 2022. Correcting Face Distortion in Wide-Angle Videos. *IEEE Transactions on Image Processing*, 31: 366–378.
- Lee, K.-Y.; Chuang, Y.-Y.; Chen, B.-Y.; and Ouhyoung, M. 2009. Video stabilization using robust feature trajectories. In *2009 IEEE 12th International Conference on Computer Vision*, 1397–1404.
- Liu, S.; Tan, P.; Yuan, L.; Sun, J.; and Zeng, B. 2016. MeshFlow: Minimum Latency Online Video Stabilization. In *European Conference on Computer Vision*.
- Liu, S.; Yuan, L.; Tan, P.; and Sun, J. 2013. Bundled camera paths for video stabilization. *ACM Trans. Graph.*, 32(4).
- Liu, S.; Yuan, L.; Tan, P.; and Sun, J. 2014. SteadyFlow: Spatially Smooth Optical Flow for Video Stabilization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Nie, L.; Lin, C.; Liao, K.; Liu, S.; and Zhao, Y. 2024a. Semi-Supervised Coupled Thin-Plate Spline Model for Rotation Correction and Beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12): 9192–9204.
- Nie, L.; Lin, C.; Liao, K.; Zhang, Y.; Liu, S.; Ai, R.; and Zhao, Y. 2024b. Eliminating Warping Shakes for Unsupervised Online Video Stitching. *arXiv:2403.06378*.
- Peng, Z.; Ye, X.; Zhao, W.; Liu, T.; Sun, H.; Li, B.; and Cao, Z. 2024. 3D Multi-frame Fusion for Video Stabilization. *arXiv:2404.12887*.
- Qi, C.; Cun, X.; Zhang, Y.; Lei, C.; Wang, X.; Shan, Y.; and Chen, Q. 2023. FateZero: Fusing Attentions for Zero-shot Text-based Video Editing. *arXiv:2303.09535*.
- Ren, X.; Lattas, A.; Gecer, B.; Deng, J.; Ma, C.; and Yang, X. 2023. Facial Geometric Detail Recovery via Implicit Representation. In *2023 IEEE 17th International Conference on Automatic Face and Gesture Recognition (FG)*.
- Rublee, E.; Rabaud, V.; Konolige, K.; and Bradski, G. 2011. ORB: An efficient alternative to SIFT or SURF. In *2011 International Conference on Computer Vision*, 2564–2571.
- Shih, Y.; Lai, W.-S.; and Liang, C.-K. 2019. Distortion-free wide-angle portraits on camera phones. *ACM Trans. Graph.*, 38(4).
- Song, J.; Meng, C.; and Ermon, S. 2022. Denoising Diffusion Implicit Models. *arXiv:2010.02502*.
- Suvorov, R.; Logacheva, E.; Mashikhin, A.; Remizova, A.; Ashukha, A.; Silvestrov, A.; Kong, N.; Goka, H.; Park, K.; and Lempitsky, V. 2021. Resolution-robust Large Mask Inpainting with Fourier Convolutions. *arXiv preprint arXiv:2109.07161*.
- Tan, J.; Zhao, S.; Xiong, P.; Liu, J.; Fan, H.; and Liu, S. 2021. Practical Wide-Angle Portraits Correction with Deep Structured Models. *arXiv:2104.12464*.
- Wang, Y.-S.; Liu, F.; Hsu, P.-S.; and Lee, T.-Y. 2013. Spatially and Temporally Optimized Video Stabilization. *IEEE Transactions on Visualization and Computer Graphics*, 19(8): 1354–1361.
- Wu, J. Z.; Ge, Y.; Wang, X.; Lei, W.; Gu, Y.; Shi, Y.; Hsu, W.; Shan, Y.; Qie, X.; and Shou, M. Z. 2023. Tune-A-Video: One-Shot Tuning of Image Diffusion Models for Text-to-Video Generation. *arXiv:2212.11565*.
- Yan, Y.; Zhou, Z.; Wang, Z.; Gao, J.; and Yang, X. 2024. DialogueNeRF: towards realistic avatar face-to-face conversation video generation. *Visual Intelligence*, 2(1): 24.
- Yang, J.; Schonfeld, D.; Chen, C.; and Mohamed, M. 2006. Online Video Stabilization Based on Particle Filters. In *2006 International Conference on Image Processing*, 1545–1548.
- Yang, L.; Zhang, Z.; Song, Y.; Hong, S.; Xu, R.; Zhao, Y.; Zhang, W.; Cui, B.; and Yang, M.-H. 2024. Diffusion Models: A Comprehensive Survey of Methods and Applications. *arXiv:2209.00796*.
- Yang, S.; Zhou, Y.; Liu, Z.; and Loy, C. C. 2023. Rerender A Video: Zero-Shot Text-Guided Video-to-Video Translation. *arXiv:2306.07954*.
- Yao, L.; Chen, C.; Li, X.; Yan, Z.; and Zuo, W. 2024. Combining Generative and Geometry Priors for Wide-Angle Portrait Correction. *arXiv:2410.09911*.
- Yu, J.; and Ramamoorthi, R. 2019. Robust Video Stabilization by Optimization in CNN Weight Space. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3795–3803.

Zhang, T.; Patil, S. G.; Jain, N.; Shen, S.; Zaharia, M.; Stolica, I.; and Gonzalez, J. E. 2024. RAFT: Adapting Language Model to Domain Specific RAG. arXiv:2403.10131.

Zhao, M.; and Ling, Q. 2020. PWStableNet: Learning Pixel-Wise Warping Maps for Video Stabilization. *IEEE Transactions on Image Processing*, 29: 3582–3595.

Zhao, W.; Li, X.; Peng, Z.; Luo, X.; Ye, X.; Lu, H.; and Cao, Z. 2023. Fast Full-frame Video Stabilization with Iterative Optimization. arXiv:2307.12774.

Zhu, F.; Zhao, S.; Wang, P.; Wang, H.; Yan, H.; and Liu, S. 2022. Semi-Supervised Wide-Angle Portraits Correction by Multi-Scale Transformer. arXiv:2109.08024.