

# Editing Physiological Signals in Videos Using Latent Representations

Tianwen Zhou  
University College London  
tianwenzhou0521@gmail.com

Akshay Paruchuri  
UNC Chapel Hill  
akshay@cs.unc.edu

Josef Spjut  
NVIDIA  
josef.spjut@gmail.com

Kaan Akşit  
University College London  
kaanaksit@kaanaksit.com

## Abstract

Camera-based physiological signal estimation provides a non-contact and convenient means to monitor Heart Rate (HR). However, the presence of physiological signals in facial videos raises significant privacy concerns, as they can reveal sensitive personal information related to the health and emotional states of an individual. To address this, we propose a learned framework that edits physiological signals in videos while preserving visual fidelity. First, we encode an input video into a latent space via a pretrained 3D Variational Autoencoder (3D VAE), while a target HR prompt is embedded through a frozen text encoder. We fuse them using a set of trainable spatio-temporal layers with Adaptive Layer Normalizations (AdaLN) to capture the strong temporal coherence of remote Photoplethysmography (rPPG) signals. We apply Feature-wise Linear Modulation (FiLM) in the decoder with a fine-tuned output layer to avoid the degradation of physiological signals during reconstruction, enabling accurate modulation in the reconstructed video. Empirical results show that our method preserves visual quality with an average PSNR of 38.96 dB and SSIM of 0.98 on selected datasets, while achieving an average HR modulation error of 10.00 bpm MAE and 10.09% MAPE using a state-of-the-art rPPG estimator. Our design’s controllable HR editing is useful for applications such as anonymizing biometric signals in real videos or synthesizing realistic videos with desired vital signs.

## 1. Introduction

Physiological signals such as HR serve as valuable indicators of a person’s physical and emotional states [1]. With the rise of camera-based physiological sensing technologies [20], these signals can now be extracted from facial videos in a non-contact [5] and scalable manner [39], enabling applications in remote health monitoring and human-

computer interaction [22]. However, this growing capability also introduces a significant and under-addressed privacy threat. Unlike visible identity cues, physiological signals are often leaked involuntarily while being invisible in human-centric videos [15]. When extracted without consent, they can be misused to infer sensitive health conditions, stress levels, or emotional states, raising concerns over covert surveillance, profiling, and medical discrimination. Most existing visual privacy methods conceal visible identity cues (e.g., face blurring or synthetic replacement) [14, 30]. While such approaches may unintentionally distort or even suppress the original PPG signal embedded in subtle skin color changes, they often do so at the cost of introducing or amplifying artificial identity-related visual cues, which may be undesirable in privacy-sensitive scenarios.

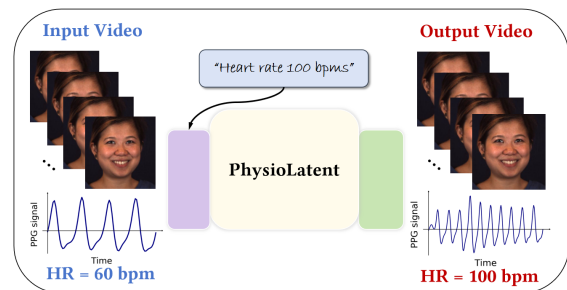


Figure 1. Our learned framework modifies physiological signals in videos by editing them using fused latent representations, maintaining good visual fidelity in the output videos. The input video and HR text prompt are encoded and fused to produce the latent representation, which is compatible with foundational generative processes. Source images are from MMSE-HR [52].

We propose a framework for editing physiological signals in videos while preserving visual fidelity. The input video is first encoded into a latent space using a pretrained 3D Variational Autoencoder (3D VAE), and a tar-

get HR prompt is embedded via a frozen text encoder. We deliberately choose a 3D VAE latent representation to establish a common latent space that is naturally compatible with foundational generative processes [19], more specifically, Latent Diffusion Models [31] and Video Diffusion Models [41, 48]. This compatibility ensures that our approach can flexibly integrate with, or serve as a post-processing stage for, emerging generative video frameworks. To explicitly exploit the strong temporal coherence of rPPG signals, we utilize trainable spatio-temporal fusion layers with temporal self-attention module augmented by AdaLN [25], allowing precise, long-range temporal conditioning. Moreover, given the subtle amplitude of rPPG variations, we introduce FiLM [26] in the decoder and fine-tune its output layer, enabling the reconstruction of videos that accurately match the desired physiological modulation. This design ensures both controllable and temporally coherent editing of HR signals without sacrificing perceptual quality. Empirical results show that our method preserves visual quality with an average PSNR of 38.96 dB and SSIM of 0.98 on selected datasets, while achieving an average HR modulation error of 10.00 bpm MAE and 10.09% MAPE on selected rPPG estimator. Our contributions are as follows:

- **Temporal-Aware Latent Fusion.** We design a spatio-temporal latent fusion block with temporal self-attention enhanced by AdaLN. This block helps maintain temporal coherence with the conditioned rPPG signals.
- **Subtle-Signal-Aware Decoder Conditioning.** We address the requirement of preserving small amplitude variations in rPPG signals by introducing FiLM into the 3D VAE decoder and fine-tuning its output layer, enabling accurate physiological modification of target waveforms while maintaining high visual quality.

Our framework further advances privacy-preserving video sharing, synthetic biometric data generation, and editing of physiological signals in human avatars [6]. Moreover, it can serve as a post-processing stage for generative video models or privacy tools, enabling the synthesis of realistic facial videos with controllable vital signs without degrading fidelity. Our extensive evaluations validate the robustness of the approach, and we will release the codebase publicly upon acceptance.

## 2. Related Works

We review three relevant categories: (1) Camera-based Physiological Measurement; (2) Manipulating Vital Signs in Videos; and (3) Biometric Privacy Protection. Table. 1 compares our work with the relevant state-of-the-art works in the literature.

### 2.1. Camera-based Physiological Measurement

Monitoring physiological signals, such as HR and respiration, is crucial for assessing human emotions and health

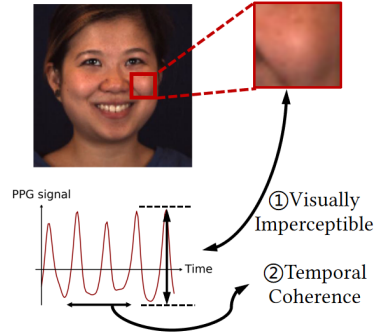


Figure 2. The two key characteristics of physiological signals in video: temporal coherence and visually imperceptible change. Source images are from MMSE-HR [52].

[39]. rPPG enables non-contact acquisition of these signals from facial videos. The fundamental mechanism of rPPG relies on the interaction of light with the multi-layered structure of the skin. Hemoglobin in the dermis partially absorbs light, and a camera reflects and captures the remainder [5]. As blood volume fluctuates with each heartbeat, these changes cause periodic variations in reflected light, which can be captured as subtle color changes in facial regions [12]. Importantly, such physiological signals possess two characteristics: they exhibit **strong temporal coherence**, as their variations evolve gradually and rhythmically over time, and they are typically **visually imperceptible to the human eye**, making them difficult to detect or manipulate directly in pixel space, as visually described in Fig. 2.

Table 1. Comparison of video-based vital sign manipulation methods. Columns indicate whether a method supports HR editing, removing physiological signals, preserving temporal coherence in edited physiological signals, protecting biometric privacy, focusing on motion robustness.

	HR Edit	Signal Removal	Temporal Coherency	Privacy	Motion Robustness
<b>This work</b>	Yes	Yes	High	Yes	No
Chen et al. [6]	Yes	Yes	Low	Yes	No
Wang et al. [45]	Yes	No	High	No	No
Paruchuri et al. [23]	No	No	High	No	Yes
Sun et al. [37]	Yes	Yes	Low	Yes	No

Poh et al. [27] were among the first to demonstrate an accurate and low-cost rPPG method using a webcam, which combined automatic face tracking with blind source separation to extract HR. Since then, various signal extraction pipelines and benchmark tools, such as the rPPG toolbox [18], have been developed to evaluate algorithms under standardized settings. Recent surveys by McDuff [20] and Xiao et al. [46] have reviewed progress in camera-based physiological sensing, covering signal origin, extraction pipelines, and application domains. For complete-

ness, we also briefly review estimation techniques, as they are used for evaluation in our experiments. Conventional rPPG pipelines typically begin with face detection, followed by Region of Interest (ROI) selection (e.g., forehead or cheeks), temporal averaging, and blind source separation such as PCA [13], POS [44], and CHROM [11]. More recently, deep learning models including TSCAN [17], DeepPhys [7], PhysNet [49], and PhysFormer++ [50] have been proposed to directly regress HR from video using CNNs or transformers. *In our work, we adopt well-established methods [7, 11, 13, 35, 44, 49, 50] to validate the accuracy of physiological signal editing in videos.*

## 2.2. Manipulating Vital Signs in Videos

Recent works have explored ways to manipulate or synthesize such signals, especially for privacy, controllability, and data augmentation. PulseEdit [6] introduced a method that perturbs the underlying rPPG signal in videos—either suppressing it or modifying it toward a target waveform—while preserving the visual appearance. Similarly, Wang et al. [45] proposed a generative adversarial framework that synthesizes realistic face videos conditioned on rPPG signals, enabling control over the generated physiological patterns. Most recently, Paruchuri et al. [23] proposed a motion augmentation framework that reasonably preserves HR while using natural videos (e.g., interviews on YouTube) to inject or modify rigid and non-rigid facial motions, enabling more robust physiological sensing under realistic motion conditions. The physiological signals’ successful manipulation requires models to preserve continuity and consistency across time, while ensuring that modifications remain visually unobtrusive. *Unlike prior methods replacing or removing HR on the image without any intermediate representations, our work edits HR directly in latent space, enabling temporally coherent and visually faithful physiological signal manipulation. Usage of 3D VAE latent space reflects a growing interest towards physiologically-conditioned generative modeling.*

## 2.3. Biometric Privacy Protection

Traditional biometric privacy protection algorithms [2] focus on de-identifying facial or fingerprint features to prevent identity theft. Recent work [3] has shown that physiological signals such as HR and stress level can also be inferred from facial videos without consent, raising concerns about profiling and emotional inference. PulseEdit [6] introduced a perturbation-based way that removes or alters physiological signals in facial videos while preserving visual appearance, thereby reducing privacy leakage. De-id [33] proposed a learning-based facial video de-identification method that eliminates facial identity from videos while retaining the embedded [37] signals and visual quality. Privacy-Phys [37] presented a 3D CNN-based framework for modifying rPPG

signals in facial videos to prevent physiological information leakage. *Our work shares the same privacy-preserving motivation of these studies.*

## 3. Editing Physiological Signs Using Latent Representations

We begin by defining our research problem, followed by an overview of our proposed framework for editing and anonymizing rPPG signals in video clips while preserving visual fidelity.

### 3.1. Problem Definition

Let  $v_t \in \mathbb{R}^{C \times H \times W}$  denote the  $t$ -th frame of a video, where  $C$ ,  $H$ , and  $W$  are the number of color channels, height, and width. Each frame  $v_t$  encodes both visual appearance  $v_t^{\text{vis}}$  and biological signal components  $v_t^{\text{bio}}$ , such that:

$$v_t = v_t^{\text{vis}} + v_t^{\text{bio}}. \quad (1)$$

where  $v_t^{\text{vis}}$  denotes the dominant visual appearance component of the frame, including facial identity, texture, and illumination, while  $v_t^{\text{bio}}$  represents the latent physiological signal component manifested as subtle, temporally varying color fluctuations correlated with biological rhythms (e.g., HR). Given a video  $v$  and a condition vector  $c \in \mathbb{R}^m$  specifying the target HR, our goal is to modify the biological component  $v_t^{\text{bio}}$  while preserving the visual component  $v_t^{\text{vis}}$ . Formally, we aim to learn a transformation function  $F(\cdot)$  that produces an anonymized video  $\hat{v} = \{\hat{v}_1, \hat{v}_2, \dots, \hat{v}_T\}$  such that

$$\hat{v}_t = F(v_t, c) = v_t^{\text{vis}} + \hat{v}_t^{\text{bio}}. \quad (2)$$

The anonymized biological signal  $\hat{v}_t^{\text{bio}}$  should satisfy these criteria:

- **Privacy Preservation:** The modified signal should reflect the desired HR  $HR_d$ , i.e.,  $HR(\hat{v}_t^{\text{bio}}) \approx HR_d$ .
- **Visual Fidelity:** The perceptual quality of the visual content must remain close to the original, i.e.,  $\mathcal{F}(v_t^{\text{vis}}, \hat{v}_t^{\text{vis}}) \leq k$ , where  $\mathcal{F}$  is a visual similarity measure (e.g., Mean Squared Error (MSE), LPIPS [51]).
- **Computational Efficiency:** The transformation function  $F(\cdot)$  should be efficient enough to support real-time applications.

To this end, we formulate the following optimization problem:

$$\hat{v}_t \leftarrow \arg \min_{v_t} \mathcal{L}(F(v_t, c), \hat{v}_t), \quad (3)$$

where the loss function  $\mathcal{L}$  will be detailed in Sec. 3.2.4. In addition, it is also important to note that rPPG signals exhibit two key characteristics visually described as in Fig 2: (1) **strong temporal dependency**, as the periodic rhythm of heartbeat manifests across multiple frames; and (2) **subtle pixel-level variations**, typically of less than 1% change in

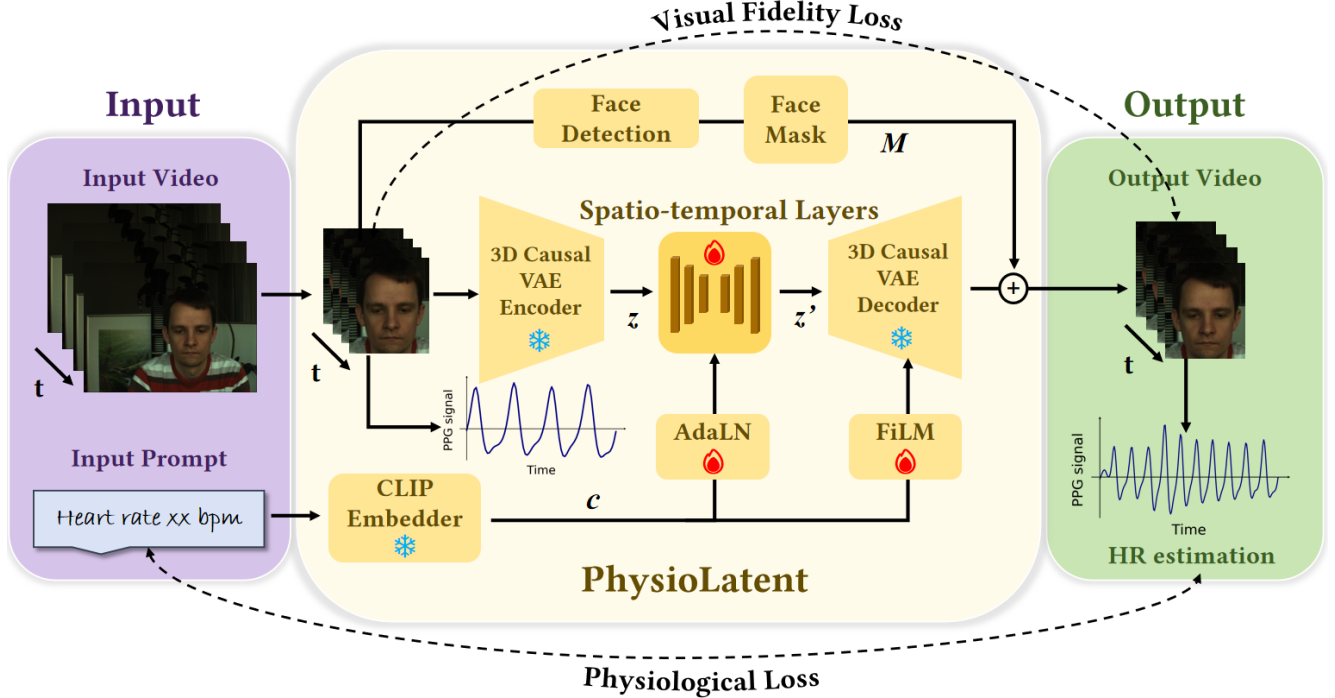


Figure 3. Overview of the proposed framework. An input facial video and a HR prompt are encoded into latent representations  $z$  and  $c$  using a 3D Causal VAE [48] encoder and a CLIP text encoder, respectively. The fused latent features are processed by spatio-temporal layers with AdaLN to inject temporal coherence and subtle rPPG variations. The 3D Causal VAE decoder reconstructs the output video with FiLM conditioning, which is supervised by visual fidelity and physiological losses to ensure perceptual quality and accurate HR modulation. To enhance visual fidelity, we incorporate a face detection module that generates a face mask  $M$ , replacing only the facial region of the input video with the decoder output. Source images are from PURE [36].

intensity. Our study reveals that maintaining required qualities in the final videos dictates a dedicated spatio-temporal modeling, see Sec. 3.2.2 and Sec. 3.2.3 for details.

### 3.2. PhysioLatent Model Overview

We provide an illustration of our overall framework *PhysioLatent* in Figure 3. An input video clip is first encoded into a latent representation by a frozen 3D Causal VAE [48] encoder, while a target HR prompt is projected via a CLIP text encoder. These two embeddings are fused and processed by trainable spatio-temporal layers designed to capture the temporal coherence of rPPG and preserve spatial fidelity. The modified latent code is decoded by a pretrained 3D Causal VAE decoder, where FiLM [26] conditioning injects the desired physiological variation and the output layer is fine-tuned to adapt to subtle rPPG edits. To further preserve visual fidelity, we use an off-the-shelf Haar Cascade [40] face detector to localize faces and replace the facial region of the input with the decoded output, keeping the rest of the frame unchanged. To guide the transformation, the output is supervised by a visual fidelity loss and a physiological loss that enforces alignment with the desired HR. This design provides a clear separation between identity-preserving vi-

sual content and condition-controlled physiological signals, which will be detailed in the following subsections.

#### 3.2.1. Latent Space Conditioning

We firstly construct a HR prompt in natural language with a fixed prompt template—e.g., “Heart rate  $xx$  bpm”, where  $xx \in [60, 120]$ —and encode it using a frozen CLIP [29] text encoder. We choose CLIP embeddings over a one-hot encoding scheme not only because CLIP provides stronger representational capacity, but also because it is more naturally compatible with future generative editing pipelines, as illustrated in Sec. 1.1. The resulting 512-dimensional embedding is projected via a learnable linear layer to match the spatial dimensions of the video latent features. The projected prompt token is concatenated channel-wise with the video latent tensor and passed to the spatio-temporal layers for conditional transformation.

#### 3.2.2. Spatio-temporal Layers

Biologically plausible modification in latent space demands modelling both spatial structure and temporal dynamics (see Fig. 2). Spatial features encode facial geometry and texture, while temporal cues, such as color fluctuations and micro-motions, play a crucial role in HR modulation.

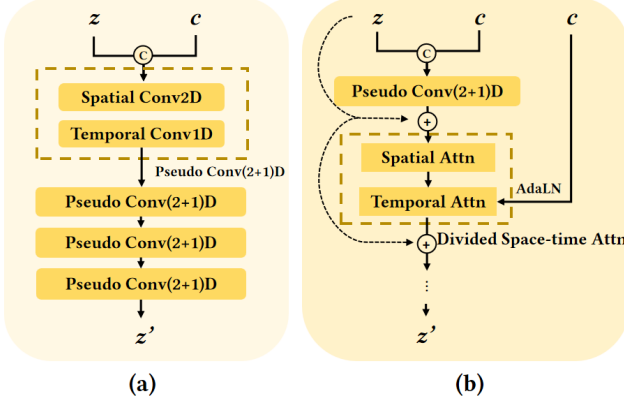


Figure 4. Comparison of proposed spatio-temporal layer settings. (a) A naïve baseline that fuses features using only stacked pseudo (2+1)D convolutions, which captures local patterns but fails to model long-range dependencies. (b) Our improved design that addresses the temporal correlation of rPPG signals by introducing decomposable space-time self-attention, with AdaLN injecting the desired HR signal into the temporal stream for fine-grained modulation.  $z$  and  $z'$  denote the input and output latent vectors, and  $c$  corresponds to the embedded text prompt.

The model sketched in Fig. 4(a) represents a naïve yet commonly used architecture for spatio-temporal fusion, relying solely on stacked pseudo (2+1)D convolutions [28]. While it captures local patterns, it lacks the ability to explicitly model long-range dependencies critical to rPPG signal propagation. In contrast, setting shown in Fig. 4(b) addresses the **temporal correlation** requirement of rPPG signals. It incorporates a decomposable space-time self-attention module, allowing the model to learn long-range dependencies effectively. Moreover, to **inject the desired HR signal into the temporal stream**, we employ AdaLN within the time-attention block, enabling fine-grained modulation of latent features conditioned on the HR prompt.

### 3.2.3. Visually Imperceptible Learned Modifications

Given the visually imperceptible subtle nature of rPPG signals (see Fig. 2), we further enhance the decoder’s sensitivity to the desired physiological signal by introducing explicit conditioning. We employ a FiLM [26] layer within the 3D Causal VAE decoder. The CLIP-projected HR embedding is used to generate scaling and shifting parameters, which modulate the intermediate activations of the decoder. This conditioning mechanism allows the framework to better reconstruct videos that conform to the specified HR. Furthermore, we unfreeze the output layer of the decoder and fine-tune it during training. This choice stems from our observation that the pre-trained decoder, while powerful for generic video reconstruction, is not optimized for capturing the subtle rPPG dynamics. Partial fine-tuning improves adaptability without losing the pretrained knowledge.

### 3.2.4. Loss Function

The training objective combines both visual and physiological constraints to ensure that the edited video maintains perceptual realism while aligning with the desired HR.

**Visual Fidelity Loss ( $\mathcal{L}_F$ ):** We adopt a hybrid loss combining pixel-wise MSE and LPIPS to preserve visual quality:

$$\mathcal{L}_F = \text{MSE}(v, \hat{v}) + \text{LPIPS}(v, \hat{v}). \quad (4)$$

**Wave Loss ( $\mathcal{L}_{\text{wave}}$ ):** To enforce temporal coherence and guide the model toward the desired waveform, we compute the negative Pearson correlation between the anonymized rPPG signal and a target sinusoidal reference with frequency  $f = HR_d/60$ :

$$\mathcal{L}_{\text{wave}} = 1 - \text{Pearson}(rPPG(\hat{v}), \sin(2\pi ft)). \quad (5)$$

From a morphological perspective, the Pearson correlation emphasizes the similarity of the waveform shape, thereby guiding the anonymized rPPG toward a smooth, periodic morphology while suppressing subject-specific details.

**Frequency Loss ( $\mathcal{L}_{\text{freq}}$ ):** To align the peak frequency  $f_{\text{pred}}$  of the estimated signal with the target frequency, we define:

$$\mathcal{L}_{\text{freq}} = |f - f_{\text{pred}}|, \quad (6)$$

where  $f_{\text{pred}}$  is obtained via Fast Fourier Transform (FFT) [10]. In a curriculum learning manner,  $\mathcal{L}_{\text{freq}}$  is excluded during the first ten epochs, allowing the model to focus on reconstructing coarse visual structures. From epoch ten onward, its contribution is gradually increased via a weight function  $\beta(t)$ , as shown in Eq. 8.

**Total Loss:**

$$\mathcal{L}_{\text{total}} = \underbrace{\alpha\mathcal{L}_{\text{wave}} + \beta(t)\mathcal{L}_{\text{freq}}}_{\text{Physiological Loss}} + \lambda\mathcal{L}_F, \quad (7)$$

where  $\alpha$  controls the wave loss,  $\beta(t)$  controls the frequency loss for epoch  $t$ , and  $\lambda$  balances the fidelity term against the physiological terms. In our experiment, we set the weights as  $\alpha = 0.2$ ,  $\lambda = 1.0$ , and use a warm-up schedule for the frequency loss:

$$\beta(t) = \begin{cases} 0, & t \leq 10, \\ 0.005(t - 10), & t > 10, \end{cases} \quad (8)$$

so that the frequency loss is disabled in the first ten epochs, then linearly ramped from 0 over the subsequent epochs.

## 4. Evaluation

We detail our experimental settings used for all evaluations. Then, we provide a detailed analysis in qualitative and quantitative terms. Lastly, we provide benchmark comparison and an ablation study to justify our design choices.

## 4.1. Implementation

We adopt the publicly available 3D Causal VAE [48] model as our video compressor. This model provides a four times downsampling in spatial resolution and an eight times downsampling in the temporal dimension. For text embedding, we utilize the pretrained CLIP [29] model ViT-B/32 as the prompt encoder. Our framework is trained using four NVIDIA RTX 4090 GPUs, each with 24 GB of memory. We use a batch size of 4 and conduct the training for 30 epochs. The training is optimized using the AdamW [24] optimizer, and a OneCycle learning rate scheduler [34] is employed, with base learning rate set to  $5 \times 10^{-4}$ . Prior to being fed into the 3D Causal VAE encoder, each input video is temporally chunked into clips of 128 frames and spatially cropped to retain only the facial region. Following the resolution used in the state of the art literature [18, 49], we resize the resolution of facial region to 72 px by 72 px. We adopt the well-established POS estimator [44] throughout all stages of our framework, following its widespread use in recent rPPG works [18]. POS provides stable and consistent pulse estimations compared with alternative estimators [7, 50] as evidenced evaluations in our Table 4.

**Dataset.** We adopt multiple publicly available rPPG datasets to enable both within-dataset and cross-dataset evaluation. The PURE dataset [36] consists of 59 facial video recordings captured at a resolution of  $640 \times 480$  pixels and a frame rate of 30 frames per second (fps). All recordings were collected in well-lit indoor environments from a total of 10 subjects, each recorded under six different conditions: steady, talking, slow translation, fast translation, small rotation, and medium rotation. We use an 80/20 split of the dataset for training and testing, respectively. In addition, we include the UBFC-rPPG dataset [4], which contains 42 videos from 42 subjects recorded under relatively constrained conditions. To enhance diversity in subject appearance, we also evaluate on the MMPD dataset [38], which includes subjects across different skin tone types and a variety of illumination and motion conditions. This combination of datasets allows us to assess the robustness of our method under varying recording and camera conditions, subject diversity, and domain shift. Our face detection, cropping, and mask generation follow the standardized implementation in the rPPG toolbox [18], which specifies detector configuration, thresholds, and failure-handling strategies; therefore, we directly adopt this pipeline without further modification.

**Evaluation Metrics.** Our assessments measure how well the modified videos align with the desired objectives. For visual fidelity, we employ PSNR and SSIM. To evaluate the effectiveness of HR modulation, we employ several established rPPG estimation algorithms. Specifically, we utilize CHROM [11], PCA [13], and POS [44] as unsupervised methods, and adopt neural network-based su-

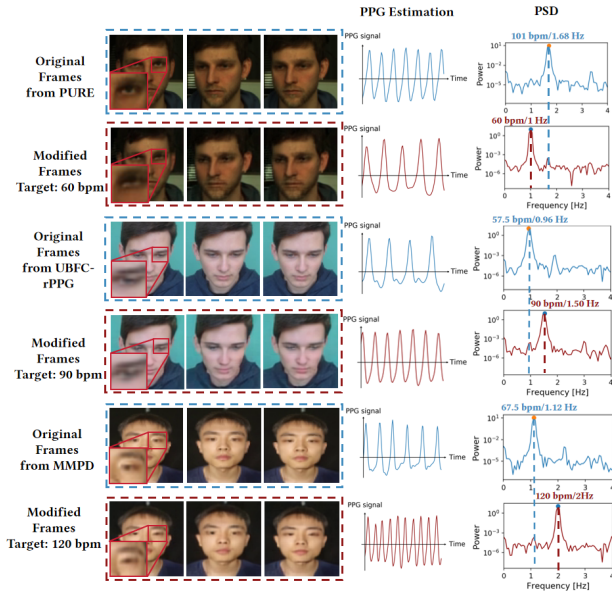


Figure 5. Qualitative comparison results of our proposed framework. We visualize representative novel frames from multiple datasets before and after modification, together with the estimated rPPG signals under the POS estimator and the corresponding Power Spectrum Density (PSD) curves. We target HR values commonly encountered in real-life scenarios. In addition, we perform zooming on selected regions of the frames to better reveal fine-grained details. From zoomed-in results, it can be observed that due to the encoding–decoding pipeline of the 3D VAE backbone, local distortions appear in high-frequency areas such as edges and textures. Source images are from MMPD [38].

pervised approaches including TSCAN [17], PhysNet [49], PhysFormer++ [50], and DeepPhys [7]. The HR estimated from each modified video is then compared with the corresponding desired HR  $HR_d$  using standard error metrics, including Mean Absolute Error (MAE) and the Mean Absolute Percentage Error (MAPE).

## 4.2. Qualitative Results

Figure 5 shows qualitative comparisons across multiple datasets. As can be observed from the original frames and the modified frames across multiple datasets, our method introduces minimal visual distortion, thereby preserving high visual fidelity in the reconstructed videos. The facial appearance and background remain consistent with the input, demonstrating that the proposed framework effectively preserves visual fidelity and edits HR. For these experiments, we target three representative HR values commonly encountered in practice: 60 bpm, 90 bpm, and 120 bpm. Moreover, from the zoomed-in regions shown in the figure, it can be observed that, due to the encoding–decoding pipeline of the 3D VAE backbone, local distortions may occasionally appear in high-frequency areas such as edges and

Table 2. Comparison of video quality and rPPG estimation accuracy between input videos and outputs produced by our framework across three benchmark datasets. PSNR and SSIM demonstrate that the visual fidelity of videos is well preserved. For physiological signal analysis, the HRs are estimated under the POS estimator. We report MAE and MAPE as mean  $\pm$  standard deviation with respect to the selected target HRs. Results highlight that the input signals exhibit large errors from the target HR, while our outputs consistently achieve substantially lower MAE/MAPE, indicating the success of the HR modification towards the target.

Dataset	Target HR(bpm)	PSNR(dB) $\uparrow$	SSIM $\uparrow$	Input rPPG Accuracy		Output rPPG Accuracy	
				MAE(bpm) $\downarrow$	MAPE(%) $\downarrow$	MAE(bpm) $\downarrow$	MAPE(%) $\downarrow$
PURE [36]	60	39.04	0.98	38.95 ( $\pm$ 13.8)	31.10 ( $\pm$ 12.9)	9.22 ( $\pm$ 4.4)	8.20 ( $\pm$ 4.1)
	80	39.02	0.98	41.80 ( $\pm$ 13.5)	35.88 ( $\pm$ 12.7)	9.98 ( $\pm$ 4.0)	10.55 ( $\pm$ 3.6)
	100	38.85	0.98	44.67 ( $\pm$ 12.9)	38.92 ( $\pm$ 12.1)	10.41 ( $\pm$ 4.1)	10.29 ( $\pm$ 3.9)
	120	38.94	0.98	50.63 ( $\pm$ 14.2)	42.20 ( $\pm$ 13.7)	10.36 ( $\pm$ 4.5)	11.34 ( $\pm$ 4.0)
	<b>Avg</b>	<b>38.96</b>	<b>0.98</b>	<b>44.01</b>	<b>37.02</b>	<b>10.00</b>	<b>10.09</b>
UBFC-rPPG [4]	60	39.86	0.98	31.91 ( $\pm$ 7.6)	33.82 ( $\pm$ 7.3)	11.05 ( $\pm$ 4.0)	10.08 ( $\pm$ 3.6)
	80	39.98	0.98	27.17 ( $\pm$ 9.0)	24.01 ( $\pm$ 8.7)	12.28 ( $\pm$ 4.5)	11.05 ( $\pm$ 3.9)
	100	40.11	0.98	23.65 ( $\pm$ 8.9)	19.02 ( $\pm$ 8.3)	9.82 ( $\pm$ 3.0)	10.98 ( $\pm$ 2.8)
	120	40.04	0.98	24.40 ( $\pm$ 7.9)	20.34 ( $\pm$ 7.5)	11.18 ( $\pm$ 4.1)	10.15 ( $\pm$ 3.7)
	<b>Avg</b>	<b>40.09</b>	<b>0.98</b>	<b>26.77</b>	<b>24.30</b>	<b>11.08</b>	<b>10.57</b>
MMPD [38]	60	36.95	0.94	47.57 ( $\pm$ 16.2)	33.65 ( $\pm$ 15.5)	10.62 ( $\pm$ 4.3)	7.88 ( $\pm$ 3.5)
	80	37.56	0.95	42.85 ( $\pm$ 15.8)	35.45 ( $\pm$ 15.0)	10.38 ( $\pm$ 4.1)	9.82 ( $\pm$ 3.6)
	100	37.60	0.94	46.85 ( $\pm$ 16.9)	38.82 ( $\pm$ 16.0)	7.62 ( $\pm$ 3.4)	6.70 ( $\pm$ 2.9)
	120	37.87	0.95	41.03 ( $\pm$ 16.5)	34.19 ( $\pm$ 15.9)	10.75 ( $\pm$ 4.4)	7.96 ( $\pm$ 3.6)
	<b>Avg</b>	<b>37.50</b>	<b>0.95</b>	<b>44.58</b>	<b>35.52</b>	<b>9.84</b>	<b>8.09</b>

textures. These qualitative observations highlight the effectiveness of our approach in editing physiological signals without noticeably degrading perceptual quality.

### 4.3. Quantitative Results

We consider a range of representative HRs (60, 80, 100, and 120 bpm) as target values, and prompt our model to adjust the input video accordingly, modifying its HR toward the specified target in the output. Table 2 presents a comparison of video quality and rPPG estimation accuracy between the input videos and the outputs produced by our framework across three benchmark datasets. For video quality, PSNR and SSIM are reported on the outputs relative to the inputs. Achieving around 40 dB PSNR and above 0.95 SSIM demonstrates that our method preserves high visual fidelity. For physiological signal analysis, we report MAE, and MAPE with respect to the target HR. The input signals exhibit large deviations from the target. In contrast, the output signals consistently achieve markedly reduced MAE and MAPE, confirming the success of our HR modification in aligning estimations with the desired target. We further observe that the outputs also reduce variability across subjects and trials, as indicated by smaller standard deviations, suggesting that our method not only corrects the mean error but also stabilizes estimation outcomes. In addition, we report in the supplementary material section 2 the best achievable visual quality with a plain 3D Causal VAE baseline without any modification layers, which serves as an upper bound for reconstruction fidelity under such a framework. Furthermore, Table 4 compares the results across different

rPPG estimators that we trail on, where we use 120 bpm as the target HR on PURE dataset. The results show that all unsupervised and supervised estimators are successfully misled by our framework into perceiving the manipulated signals as being close to 120 bpm. This highlights the robustness of our approach against diverse estimators.

### 4.4. Benchmark Comparison

We compare our framework against PulseEdit [6] across three benchmark datasets in terms of both video quality and rPPG estimation accuracy, as shown in Table 3 with a fixed target HR of 120 bpm. Since the authors of PulseEdit did not release official code, we implemented their method following the description in their paper to the best of our ability. Overall, PulseEdit achieves slightly higher PSNR values, which can be attributed to its purely signal processing-based design. In contrast, our approach relies on a lossy encoding-decoding pipeline, which naturally leads to slightly lower PSNR values compared to purely signal processing-based methods. Due to the lossy compression inherent in the 3D VAE encoding-decoding process, our method achieves slightly lower PSNR values compared to direct signal processing approaches (see Appendix 2). Both methods achieve comparable SSIM values, suggesting that our framework is able to preserve structural consistency of the video content despite the lossy compression. More importantly, our method consistently delivers superior performance on HR modification. Specifically, the output signals produced by our framework yield lower MAE and MAPE with respect to the reference HR, indicating that rPPG esti-

Table 3. Benchmark comparison between our method and PulseEdit [6] on video quality and rPPG estimation accuracy across multiple datasets. PSNR and SSIM are reported for the output videos, while MAE and MAPE are computed for both input and output signals **relative to a reference HR of 120 bpm** under the POS estimator. Since the authors did not release their code base, did not release official code, we implement PulseEdit based on the descriptions provided in their work, adopting similar face detector, resizing, and temporal chunking procedures as in our pipeline for fairness. Our codebase will include implementation of PulseEdit. **Bold** denotes the best result.

Dataset	Method	PSNR(dB) ↑	SSIM ↑	Input rPPG Accuracy		Output rPPG Accuracy	
				MAE(bpm) ↓	MAPE(%) ↓	MAE(bpm) ↓	MAPE(%) ↓
PURE [36]	Our Work	38.94	<b>0.9761</b>	50.63	42.20	<b>10.36</b>	<b>11.34</b>
	PulseEdit [6]	<b>42.68</b>	0.9720			16.71	12.26
UBFC-rPPG [4]	Our Work	40.04	0.9803	24.40	20.34	<b>11.18</b>	<b>10.15</b>
	PulseEdit [6]	<b>43.08</b>	<b>0.9867</b>			15.07	15.56
MMPD [38]	Our Work	37.87	0.9542	41.03	34.19	<b>10.75</b>	<b>7.96</b>
	PulseEdit [6]	<b>41.72</b>	<b>0.9664</b>			20.36	18.30

Table 4. MAE and MAPE result for input/output videos w.r.t. a fixed reference HR of 120 bpm, showing our method successfully fools a range of unsupervised and supervised rPPG estimators.

Dataset	rPPG Estimator	Input rPPG Accuracy		Output rPPG Accuracy	
		MAE(bpm) ↓	MAPE(%) ↓	MAE(bpm) ↓	MAPE(%) ↓
PURE [36]	POS	50.63	42.20	10.36	11.34
	PCA	57.08	47.57	14.54	16.62
	CHROM	49.61	41.34	14.14	11.45
	TSCAN	42.23	35.19	13.68	12.03
	DeepPhys	49.37	41.14	11.75	11.84
	PhysNet	37.40	41.17	10.07	11.58
	PhysFormer++	38.95	32.46	10.29	10.53

mators are more effectively guided toward the desired target. In addition, compared to PulseEdit, our framework demonstrates stronger robustness across datasets with diverse characteristics: even for MMPD [38] and UBFC-rPPG [4], which involve significant head movements or varying skin tones, our method maintains stable performance. This highlights the trade-off between strict pixel-level fidelity and successful physiological signal editing.

#### 4.5. Ablation Study

Table 5 reports the ablation study evaluating the contribution of each component in our framework. We begin with a baseline configuration (Setting a), which employs a spatio-temporal layer setting as in Fig. 4(a). This setting achieves limited performance since it does not account for the temporal coherence of the rPPG signal. Incorporating temporal attention with AdaLN improves both video quality and physiological consistency, as temporal modeling better preserves the temporal coherence of rPPG signals. Adding FiLM-based conditioning in the decoder further enhances performance by injecting the target condition into the generative process, resulting in improved fidelity. Finally, introducing curriculum-based frequency loss yields the best overall results, boosting PSNR/SSIM while reducing MAE/MAPE, demonstrating that progressive frequency loss effectively

guides the network to first reconstruct visual content and then gradually align with the desired HR pattern.

Table 5. Ablation study of key components in our framework. MAE/MAPE for HR estimation is conducted under POS estimator in PURE dataset [36]. Each row shows the effect of adding a component to baseline (Setting a). **Bold** denotes the best result.

Configuration	PSNR ↑	SSIM ↑	MAE ↓	MAPE ↓
Baseline Setting (a)	31.84	0.8807	27.63	25.63
+ Temporal Attn. + AdaLN	34.82	0.9242	18.28	19.84
+ Decoder FiLM conditioning	35.45	0.9153	14.92	13.91
+ Curriculum-based freq. loss	<b>38.94</b>	<b>0.9761</b>	<b>10.36</b>	<b>11.20</b>

## 5. Conclusion

We presented *PhysioLatent*, a framework for editing camera-based physiological signals in facial videos while preserving visual fidelity. By conditioning a pretrained 3D VAE latent space with text-embedded HR prompts and combining spatio-temporal attention with AdaLN and FiLM-based conditioning, our method achieves controllable rPPG manipulation. On multiple datasets, *PhysioLatent* attains high perceptual quality and low modulation error, and also supports an HR removal mode that suppresses periodic components without harming appearance. Finally, its text-driven design integrates naturally with VLM- and diffusion-based pipelines, enabling both privacy-preserving editing and the creation of physiologically consistent synthetic datasets.

## References

- [1] B. M. Appelhans and L. J. Luecken. Heart rate variability as an index of regulated emotional responding. *Review of General Psychology*, 10:229–240, 2006. 1
- [2] M. Barni, G. Droandi, and R. Lazzaretti. Privacy protection in biometric-based recognition systems: A marriage between cryptography and signal processing. *IEEE Signal Processing Magazine*, 32:66–76, 2015. 3

- [3] S. Bhutani, M. Elgendi, and C. Menon. Preserving privacy and video quality through remote physiological signal removal. *Communications Engineering*, 4:66, 2025. 3
- [4] S. Bobbia, R. Macwan, Y. Benezeth, A. Mansouri, J. Dubois, et al. Unsupervised skin tissue segmentation for remote photoplethysmography. *Pattern Recognition Letters*, 124:82–90, 2017. 6, 7, 8, 1
- [5] G. Cennini, J. Arguel, K. Akşit, A. van Leest, et al. Heart rate monitoring via remote photoplethysmography with motion artifacts reduction. *Optics Express*, 18:4867–4875, 2010. 1, 2
- [6] M. Chen, X. Liao, and M. Wu. Pulseedit: Editing physiological signals in facial videos for privacy protection. *IEEE Transactions on Information Forensics and Security*, 17:457–471, 2022. 2, 3, 7, 8
- [7] W. Chen and D. McDuff. Deepphys: Video-based physiological measurement using convolutional attention networks. In *ECCV*, pages 349–365, 2018. 3, 6
- [8] U. A. Ciftci, A. Demir, and L. Yin. How do the hearts of deep fakes beat? deep fake source detection via interpreting residuals with biological signals. In *IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–10, 2020. 1
- [9] U. A. Ciftci, I. Demir, and L. Yin. Fakecatcher: Detection of synthetic portrait videos using biological signals. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 1
- [10] J. W. Cooley and J. W. Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics of Computation*, 19:297–301, 1965. 5
- [11] G. de Haan and V. Jeanne. Robust pulse rate from chrominance-based rppg. *IEEE Transactions on Biomedical Engineering*, 60:2878–2886, 2013. 3, 6
- [12] E. d’Eon and D. Luebke. Advanced techniques for realistic real-time skin rendering, 2007. <https://developer.nvidia.com/gpugems/gpugems3/>. 2
- [13] F. L. Gewers, G. R. Ferreira, H. F. D. Arruda, F. N. Silva, C. H. Comin, D. R. Amancio, L. D. F. Costa, et al. Principal component analysis: A natural approach to data exploration. *ACM Computing Surveys*, 54, 2021. 3, 6
- [14] R. Gross, L. Sweeney, J. Cohn, F. De la Torre, S. Baker, et al. *Face De-identification*, pages 129–146. 2009. 1
- [15] D. Gupta and A. Etemad. Privacy-preserving remote heart rate estimation from facial videos. *arXiv preprint arXiv:2306.01141*, 2023. 1
- [16] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis. 3d gaussian splatting for real-time radiance field rendering. *arXiv preprint arXiv:2308.04079*, 2023. 2
- [17] X. Liu, J. Fromm, S. Patel, and D. McDuff. Multi-task temporal shift attention networks for on-device contactless vitals measurement. *arXiv preprint arXiv:2006.03790*, 2021. 3, 6
- [18] X. Liu, G. Narayanswamy, A. Paruchuri, X. Zhang, J. Tang, Y. Zhang, Y. Wang, S. Sengupta, S. Patel, D. McDuff, et al. rppg-toolbox: Deep remote ppg toolbox. *arXiv preprint arXiv:2210.00716*, 2022. 2, 6
- [19] X. Liu, T. Zhou, C. Wang, Y. Wang, Y. Wang, Q. Cao, W. Du, et al. Toward the unification of generative and discriminative visual foundation model: a survey. *The Visual Computer*, 41:3371–3412, 2025. 2
- [20] D. McDuff. Camera measurement of physiological vital signs. *ACM Computing Surveys*, 55:176:1–176:40, 2023. 1, 2
- [21] A. Nazaret, S. Tonekaboni, G. Darnell, S. Y. Ren, G. Sapiro, A. C. Miller, et al. Modeling personalized heart rate response to exercise and environmental factors with wearables data. *NPJ Digital Medicine*, 6:207, 2023. 2
- [22] J. Orlosky, M. Sra, K. Bektaş, H. Peng, J. Kim, N. Kos’myna, T. Höllerer, A. Steed, K. Kiyokawa, K. Akşit, et al. Telelife: The future of remote living. *Frontiers in virtual reality*, 2:763340, 2021. 1
- [23] A. Paruchuri, X. Liu, Y. Pan, S. Patel, D. McDuff, and S. Sengupta. Motion matters: Neural motion transfer for better camera physiological measurement. In *WACV*, pages 5933–5942, 2024. 2, 3
- [24] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, A. Lerer, et al. Automatic differentiation in pytorch, 2017. 6
- [25] W. Peebles and S. Xie. Scalable diffusion models with transformers. *arXiv preprint arXiv:2212.09748*, 2023. 2
- [26] E. Perez, F. Strub, H. de Vries, V. Dumoulin, and A. C. Courville. Film: Visual reasoning with a general conditioning layer. *CoRR abs/1709.07871*, 2017. 2, 4, 5
- [27] M.-Z. Poh, D. J. McDuff, and R. W. Picard. Non-contact, automated cardiac pulse measurements using video imaging and blind source separation. *Optics Express*, 18:10762–10774, 2010. 2
- [28] Z. Qiu, T. Yao, and T. Mei. Learning spatio-temporal representation with pseudo-3d residual networks. *arXiv preprint arXiv:1711.10305*, 2017. 5
- [29] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021. 4, 6
- [30] Z. Ren, Y. J. Lee, and M. S. Ryoo. Learning to anonymize faces for privacy preserving action detection. *CoRR abs/1803.11556*, 2018. 1
- [31] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. *arXiv preprint arXiv:2112.10752*, 2022. 2
- [32] F. E. Rosas, D. Candia-Rivera, A. I. Luppi, Y. Guo, and P. A. Mediano. Bayesian at heart: Towards autonomic outflow estimation via generative state-space modelling of heart rate dynamics. *Computers in Biology and Medicine*, 170:107857, 2024. 2
- [33] M. R. Savic and G. Zhao. De-identification of facial videos while preserving remote physiological utility. In *BMVC*, 2023. 3
- [34] L. N. Smith and N. Topin. Super-convergence: Very fast training of residual networks using large learning rates. *CoRR abs/1708.07120*, 2017. 6
- [35] R. Spetlik, J. Cech, V. Franc, and J. Matas. Visual heart rate estimation with convolutional neural network. 2018. 3
- [36] R. Stricker, S. Müller, and H.-M. Gross. Non-contact video-based pulse rate measurement on a mobile service robot. In *RO-MAN*, pages 1056–1062, 2014. 4, 6, 7, 8, 1

- [37] Z. Sun and X. Li. Privacy-phys: Facial video-based physiological modification for privacy protection. *IEEE Signal Processing Letters*, 29:1507–1511, 2022. 2, 3
- [38] J. Tang, K. Chen, Y. Wang, Y. Shi, S. Patel, D. McDuff, X. Liu, et al. Mmpd: Multi-domain mobile video physiology dataset. In *EMBC*, pages 1–5, 2023. 6, 7, 8, 1, 3
- [39] W. Verkruysse, L. O. Svaasand, and J. S. Nelson. Remote plethysmographic imaging using ambient light. *Optics Express*, 16:21434–21445, 2008. 1, 2
- [40] P. Viola and M. J. Jones. Rapid object detection using a boosted cascade of simple features. In *CVPR*, pages 511–518, 2001. 4
- [41] T. Wan, A. Wang, B. Ai, B. Wen, C. Mao, C.-W. Xie, D. Chen, F. Yu, H. Zhao, J. Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. 2
- [42] L. Wang, Y. Shi, and W. T. Ooi. Gsvc: Efficient video representation and compression through 2d gaussian splatting. *arXiv preprint arXiv:2501.12060*, 2025. 2
- [43] T.-C. Wang, A. Mallya, and M.-Y. Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *CVPR*, 2021. 3
- [44] W. Wang, A. C. den Brinker, S. Stuijk, and G. de Haan. Algorithmic principles of remote ppg. *IEEE Transactions on Biomedical Engineering*, 64:1479–1491, 2017. 3, 6
- [45] Z. Wang, Y. Ba, P. Chari, O. D. Bozkurt, G. Brown, P. Patwa, N. Vaddi, L. Jalilian, A. Kadambi, et al. Synthetic generation of face videos with plethysmograph physiology. In *CVPR*, pages 20555–20564, 2022. 2, 3
- [46] H. Xiao, T. Liu, Y. Sun, Y. Li, S. Zhao, and A. Avolio. Remote photoplethysmography for heart rate measurement: A review. *Biomedical Signal Processing and Control*, 88: 105608, 2024. 2
- [47] Yazhou Xing, Yang Fei, and Y. H. J. C. J. X. X. C. Q. C. Videovae+: Large motion video autoencoding with cross-modal video vae. In *ICCV*, 2025. 2
- [48] Z. Yang, J. Teng, W. Zheng, M. Ding, S. Huang, J. Xu, Y. Yang, W. Hong, X. Zhang, G. Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. 2, 4, 6
- [49] Z. Yu, X. Li, and G. Zhao. Remote photoplethysmograph signal measurement from facial videos using spatio-temporal networks. In *BMVC*, page 277, 2019. 3, 6
- [50] Z. Yu, Y. Shen, J. Shi, H. Zhao, Y. Cui, J. Zhang, P. Torr, G. Zhao, et al. Physformer++: Facial video-based physiological measurement with slowfast temporal difference transformer. *Int. J. Comput. Vision*, 131:1307–1330, 2023. 3, 6
- [51] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018. 3
- [52] Z. Zhang, J. M. Girard, Y. Wu, X. Zhang, P. Liu, U. Ciftci, S. Canavan, M. Reale, A. Horowitz, H. Yang, et al. Multimodal spontaneous emotion corpus for human behavior analysis. In *CVPR*, pages 3438–3446, 2016. 1, 2