

# Recovering Pulse Waves from Video Using Deep Unrolling and Deep Equilibrium Models

Shenoy, Vineet; Lohit, Suhas; Mansour, Hassan; Chellappa, Rama; Marks, Tim K.

TR2026-031 March 19, 2026

## Abstract

Camera-based contactless monitoring of vital signs, also known as imaging photoplethysmography (iPPG), has seen applications in driver-monitoring, perfusion assessment, affective computing, and more. iPPG involves sensing the underlying cardiac pulse from video of the skin and estimating vital signs such as the pulse rate or a full pulse waveform. Some previous iPPG methods impose model-based sparse priors on the pulse signals and use iterative optimization for pulse wave recovery, while others use end-to-end black-box deep learning methods. In contrast, we introduce methods that combine signal processing and deep learning methods in an inverse problem framework. Our methods estimate the underlying pulse signal, pulse rate, and pulse rate variability from facial video by learning deep- network-based denoising operators that leverage deep algorithm unfolding and deep equilibrium models. Experiments show that our methods can denoise an acquired signal from the face and infer the correct underlying pulse rate and pulse rate variability, achieving pulse rate estimation performance consistent with the state-of-the-art on well-known benchmarks, all with less than one-fifth the number of learnable parameters as the closest competing method.

*IEEE Transactions on Image Processing 2026*



# Recovering Pulse Waves from Video Using Deep Unrolling and Deep Equilibrium Models

Vineet R. Shenoy, Suhas Lohit, Hassan Mansour, Rama Chellappa, Tim K. Marks

**Abstract**—Camera-based contactless monitoring of vital signs, also known as imaging photoplethysmography (iPPG), has seen applications in driver-monitoring, perfusion assessment, affective computing, and more. iPPG involves sensing the underlying cardiac pulse from video of the skin and estimating vital signs such as the pulse rate or a full pulse waveform. Some previous iPPG methods impose model-based sparse priors on the pulse signals and use iterative optimization for pulse wave recovery, while others use end-to-end black-box deep learning methods. In contrast, we introduce methods that combine signal processing and deep learning methods in an inverse problem framework. Our methods estimate the underlying pulse signal, pulse rate, and pulse rate variability from facial video by learning deep-network-based denoising operators that leverage deep algorithm unfolding and deep equilibrium models. Experiments show that our methods can denoise an acquired signal from the face and infer the correct underlying pulse rate and pulse rate variability, achieving pulse rate estimation performance consistent with the state-of-the-art on well-known benchmarks, all with less than one-fifth the number of learnable parameters as the closest competing method.

**Index Terms**—Pulse Rate Estimation, fixed-point methods, algorithm unrolling.

## I. INTRODUCTION

Physiological measurements using cameras [1], [2] have emerged as an exciting research topic, with many recent techniques focusing on imaging photoplethysmography (iPPG) for pulse rate estimation. In this task, given a camera recording of the skin (e.g., the face), the pulse wave is reconstructed and subsequently pulse rate is determined—all without contact with the skin. Additional studies have looked at pulse rate variability (PRV), which measures the variability in the beat-to-beat interval of the pulse wave and is an indicator of neurocardiac function [3]. The blood volume pulse signal, manifested as reflections of incident light on pulsatile blood flow, is weak in the presence of noise signals, and the main challenge is to separate the signal from the motion, lighting, and quantization noise. Pixel intensity changes due to the pulsatile signal are weak compared to those due to noise sources such as motion, lighting variations, and camera sensor quantization. Developing robust solutions to iPPG waveform recovery could improve monitoring of sick patients [4], [5], monitor drivers in automotive applications [6], and improve security in face recognition [7].

Vineet R. Shenoy and Rama Chellappa are with the Johns Hopkins University, Baltimore, MD, USA (e-mail: vshenoy4@jhu.edu, rchella4@jhu.edu). VS Performed this work as an intern at MERL.

Suhas Lohit, Hassan Mansour, and Tim K. Marks are with Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA (e-mail: {slohit, mansour, tmarks}@merl.com).

Early methods to solve this problem were modular and based on signal processing techniques; these methods cropped the face and defined one or more regions of interest, followed by a signal extraction technique. These include blind source separation methods [8], [9] and later improvements used chrominance information [10] and skin reflection properties [11]. Another class of model-driven methods was based on sparse optimization [6], [12]; the authors assumed that the pulse signal was composed of a sparse set of frequency coefficients and that the noise was present in a sparse set of facial regions. The authors then solved for a pulse vector that is sparse in the frequency domain. While these hand-crafted signal-processing methods set a strong foundation for future work, they could not learn signal representations from data to improve pulse signal estimation, especially under challenging recording scenarios.

Almost all recent data-driven iPPG methods are based on deep learning. After initial work [13] modeled the nonlinear interactions between the pulse signal and noise, deep learning methods were further refined with temporal shift models [14], spatio-temporal models [15], [16], and attention [17]. Other deep neural network methods use meta-learning [18], [19], efficient computation [20], and federated learning [21]. With recent interest in learning without labels, unsupervised/self-supervised methods such as [22]–[24] have used contrastive learning and masked autoencoders to update network parameters without a supervisory signal.

Compared to earlier iPPG methods, the deep-learning-based methods have a disadvantage in that they fundamentally lack interpretability. In this paper, we show that integrating both signal processing and learning-based techniques not only achieves strong iPPG performance, but also puts deep learning techniques on sound signal-processing foundations.

To this end, our work expands on the optimization-based signal recovery algorithms while integrating lightweight deep learning components that play specialized roles for pulse signal recovery. Similar to [6], we model pulse signal recovery as an optimization problem that minimizes a data fidelity term and a prior-promoting penalty term, and solve the resulting optimization problem using ISTA-style [25] forward-backward splitting algorithms. However, using hand-crafted priors and their corresponding proximal operators do not always capture the complex and nonlinear interactions between frequencies in the pulse signal or the noise that is present in the environment. Therefore, we impose *implicit* priors on the heartbeat signal and the noise by training deep neural networks that function as denoising operators. Our approach therefore maintains the algorithmic structure of optimization-based methods while

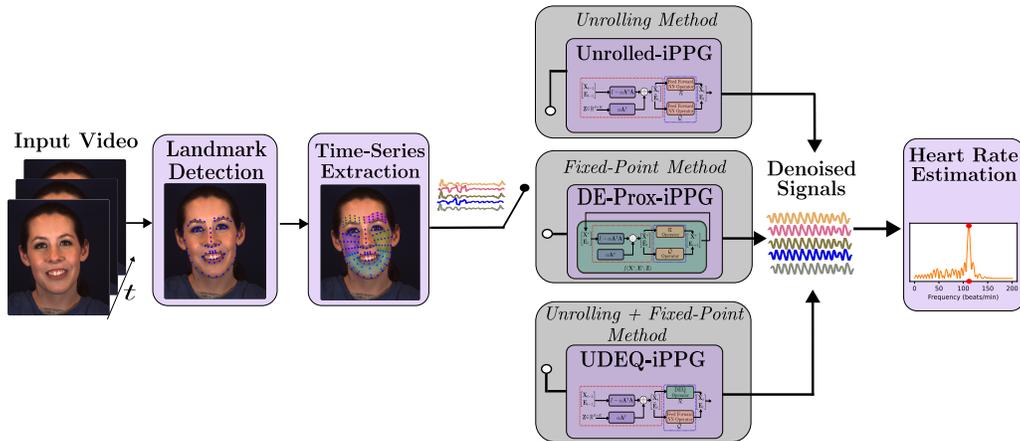


Fig. 1: Our full pipeline. We pass the input video frames through a face and landmark detector, extrapolate additional facial landmarks, and obtain a time series based on pixel intensities from five segmented face regions. We then pass our raw signals to either Unrolled iPPG, DE-Prox-iPPG, or UDEQ-iPPG which separates the signal from the noise. Finally, we sum the power spectrum coefficients across individual bins for all five signals, and select the bin with the highest power as our pulse rate estimate.

benefiting from data-driven deep learning.

We apply our insights to develop three different algorithms. First, we present **Unrolled iPPG**, which learns the denoising operators of the signal and noise priors by “unrolling” proximal gradient descent [26] for a *predetermined* number of iterations, applying a single forward pass through the structure-promoting operators of each of the signal and noise at each iteration. A potential drawback of Unrolled iPPG is that it applies only a single denoising step in each of the unrolled iterations, and the number of unrolling iterations must be decided in advance. This limitation can be overcome using fixed-point theory and deep equilibrium (DEQ) models [27], which use fixed-point solvers to effectively iterate the unrolling steps until the output converges. Incorporating these ideas, we develop **Deep Equilibrium-Proximal iPPG (DE-Prox-iPPG)**, which iterates *both* the gradient update and learned proximal/denoising steps through a solver until a fixed-point is reached [28].

Our experiments show that this method, which makes fixed-point assumptions on both the signal and noise, has its own weaknesses. Given that the pulsatile signal is a realization of the quasi-periodic nature of the beating heart [29], we assume it is sufficiently structured to be represented as a fixed point of some operator. However, the noise signal, which includes noise due to numerous sources including motion and illumination variations, may not admit a fixed-point representation. Consequently, we introduce **Unrolled-DEQ-iPPG (UDEQ-iPPG)**. In each unrolling iteration, Unrolled-DEQ-iPPG effectively applies the denoising operator on the pulse signal until a fixed-point is achieved, but applies the operator on the noise only once through a single forward pass. We show that UDEQ-iPPG’s performance exceeds that of traditional unrolling methods in cross-dataset tests.

Extensive experiments demonstrate that our models accurately estimate the underlying pulse signal from facial video, achieving results consistent with the state-of-the-art pulse

rate estimation algorithms on four publicly available datasets, and even generalizing to unseen datasets at test time. We also include a pulse rate variability study and suggest future improvements for pulse rate variability recovery and analysis.

To summarize, our main contributions are as follows:

- We model iPPG pulse signal recovery as an optimization problem that minimizes the sum of a data fidelity term and an *implicit* structure-promoting prior. We then learn the denoising neural network operator corresponding to the implicit prior.
- We present the first unrolling techniques for iPPG signal recovery. Using the ideas of proximal gradient descent, we alternate between gradient descent and neural network denoising on both the signal and noise to recover the pulsatile signal. This algorithm is called Unrolled iPPG.
- This work is the first to propose iPPG signal recovery through fixed-point methods and to incorporate deep equilibrium models into algorithms for pulse signal recovery and pulse rate estimation. We do so by encapsulating both the gradient and denoising steps in an iteration map that runs until a fixed-point is obtained for both the pulse and noise signals and call this DE-Prox-iPPG.
- We further propose UDEQ-iPPG, which integrates DE-Prox-iPPG and Unrolled iPPG. Assuming that the pulse signal can be represented as the fixed-point of a learned denoising operator, while the noise admits no such fixed-point representation, we apply the denoiser on the pulse signal until fixed-point convergence, but we apply the denoiser on the noise a single time at each iteration in an unrolling framework.
- We achieve results consistent with the state-of-the-art pulse rate estimation algorithms on four publicly available datasets, while using less than 35% of the number of parameters of competing deep neural network methods.

A preliminary version of Unrolled iPPG appeared in [30]. We expand upon this work by: 1) Incorporating fixed-point

methods for signal recovery (DE-Prox-iPPG); 2) showing that a combination of fixed-point recovery algorithms and traditional single-pass feed-forward neural networks effectively models the signal and noise (UDEQ-iPPG); 3) testing all three methods on three additional datasets; 4) performing extensive cross-dataset model evaluations; 5) evaluating on the task of pulse rate variability; and 6) providing a more thorough empirical and statistical analysis of all of our methods.

## II. RELATED WORK

### A. Imaging PPG and Physiological Analysis

Pulse rate estimation algorithms can be broadly categorized into signal-processing-based (or model-based) algorithms and deep learning algorithms. Blind source separation (BSS) methods [8], [9], [31], a subset of the signal processing methods, assume that the intensity signal extracted from video is the weighted sum of a pulse signal and other source or error signals. By assuming the source signals are statistically independent, the work of [9] uses independent component analysis (ICA) to extract the pulse signal, while [8] demixes signals along directions of maximum variance using principal components analysis (PCA). These methods were further developed to include skin reflection properties. The method of [11] does so by assuming all color variation in skin pixels is due to the pulsatile signal, and projects the extracted time series onto this color vector, the resultant signal of which is proportional to the underlying pulse. CHROM [10] first implements a skin tone correction, followed by a projection that eliminates specular reflection to extract only the pulse signal. POS [32], similar to CHROM, eliminates the specular reflection and performs skin tone correction, but defines a projection plane using a data-driven approach. Other methods have used low-rank [33] and sparsity-based methods. The method presented in [6] is most similar to our work; the authors define the extracted video signal as the sum of the pulse signal and noise, and assume the Fourier coefficients of the pulse signal and the noise signal to be sparse. They then solve a sparse optimization problem for the frequency coefficients and noise using proximal gradient descent. We propose to relax these sparsity assumptions by incorporating a learnable prior, which is realized through deep neural networks.

Deep learning methods [13]–[20], [23], [34]–[41] show significant performance improvements over model-based approaches. Many of these methods accept RGB frames directly; PhysNet [15] inputs RGB frames to spatiotemporal neural networks to reconstruct the pulse signal. DeepPhys [13] uses the temporal difference of frames to eliminate the stationary skin reflection color and reconstruct the pulse wave. Both TS-CAN [14] and InverseCAN [17] extend DeepPhys, with the former including temporal shift modules [42] and the latter estimating corruption in the background to remove noise in the foreground and improve signal-to-noise ratio (SNR). Some recent work evaluates algorithms under stricter data constraints such as the few-shot scenario [18], while EfficientPhys [20] develops pulse signal estimation algorithms assuming a resource-constrained environment. Recent work [22], [23], [36], [43] has developed neural methods with no supervisory training signal, showing good performance.

While pulse rate estimation has been the primary goal of the preceding methods, once the pulse signal is recovered a pulse rate variability analysis is also possible. Pulse Rate Variability (PRV) measures the time interval between successive beats of the heart and can be obtained by finding the peaks of a pulse waveform and the time interval between them [3], [44]. Once calculated, clinicians can use PRV metrics to assess some functions of the cardiac and nervous systems.

### B. Deep Equilibrium Models and Unrolling Algorithms

Conventional deep learning models learn by applying explicit functions (such as linear and convolutional operators, followed by nonlinearities) to inputs, building a computation graph along the way, and back-propagating through this graph to update parameters after a loss computation. Implicit models such as Neural ODEs [45] and Deep Equilibrium Models [27] do not build computation graphs; instead, they solve a nonlinear system prescribed by a criterion such as the endpoint of an ODE flow or the root of an equation, and backpropagate through this point. Thus, they decouple the choice of solver from the computation of the endpoint or root, and backpropagate through this point analytically and independently of the forward pass.

Deep Equilibrium Models (DEQs) [27] approximate “infinitely deep” models by solving for the fixed point (or equilibrium point)  $\mathbf{z}^*$  of a function:  $\mathbf{z}^* = f_\theta(\mathbf{z}^*, \mathbf{x})$ , where  $\mathbf{x}$  is the input. The choice of solver for  $\mathbf{z}^*$ —whether it be Broyden’s Method [46], Peaceman-Rachford splitting [47], Anderson Acceleration [48], or any other method—is independent of the backpropagation procedure. As a result, the Implicit Function Theorem [49] can be used to compute the gradients of the parameters. These implicit models have shown competitive performance on word-level language modeling [27] as well as ImageNet classification and Cityscapes semantic segmentation [50]. Further developments showed that these models can simultaneously perform forward inference as well as input optimization on problems such as generative modeling, adversarial training, and gradient-based meta-learning [51]. Recent work has applied DEQs to optical flow estimation [52], diffusion models [53], facial landmark estimation [54], and imaging inverse problems [28], [55]–[59]. In addition, significant research has explored the theoretical foundations of DEQs. After [27] demonstrated training instabilities in DEQs, inexact gradient methods [60], [61] and regularization techniques [62] were developed that improved fixed-point convergence, while the work of [63], [64] explored interval bound propagation methods to compute lower and upper bounds on the output of a layer.

Algorithms that truncate after a fixed number of iterations  $T$  (instead of the “infinite” iterations of a DEQ), and backpropagate through the iterations during training, are called unrolling algorithms [26]. Deep unrolling has been used in magnetic resonance imaging [65], compressive image recovery [66], and multispectral image fusion [67]. Recently, unrolling has also been applied to Interferometric Hyperspectral Image (IHI) recovery; the work of IHRUT [68] treats Interferometric Hyperspectral Image (IHI) recovery as a regularized optimization problem with priors designed to recover stripe-wise spectral

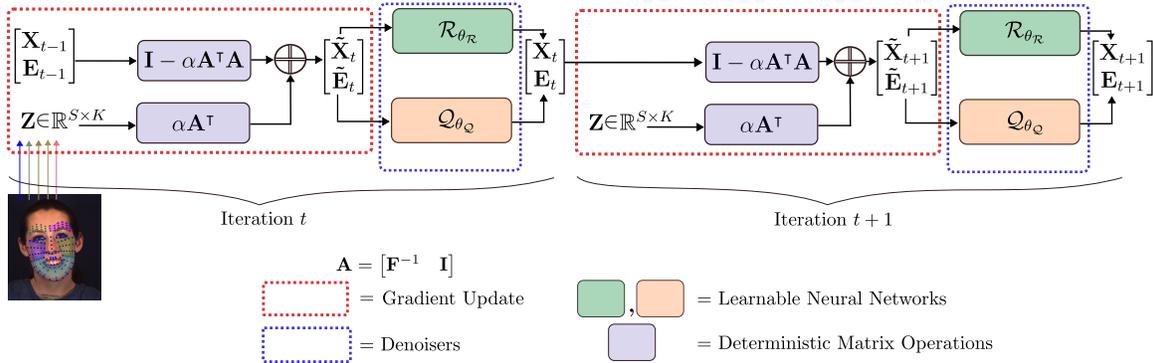


Fig. 2: Our unrolling algorithm for both Unrolled iPPG and UDEQ-iPPG. We replace the proximal operators of proximal gradient descent by deep denoisers  $\mathcal{R}_{\theta_{\mathcal{R}}}$  and  $\mathcal{Q}_{\theta_{\mathcal{Q}}}$  on the frequency coefficients  $\mathbf{X}$  and noise  $\mathbf{E}$ , respectively. Unrolled iPPG implements a single forward pass through  $\mathcal{R}_{\theta_{\mathcal{R}}}$  and  $\mathcal{Q}_{\theta_{\mathcal{Q}}}$ , while UDEQ-iPPG effectively applies  $\mathcal{R}$  to  $\tilde{\mathbf{X}}_t$  until the output converges to a fixed-point.

characteristics, and uses a custom unrolled transformer to recover the IHI signal. We propose similar Unrolling and DEQ algorithms for iPPG signal recovery below.

### III. SIGNAL MODEL

The circulation of blood in the human body can be described as a pump—the heart—that rhythmically applies a force on a branched system of elastic tubes to move blood [29]. Optical sensors such as those that perform photoplethysmography (PPG) capture blood volume changes by measuring transmission or reflection of illumination that is incident on the skin. While PPG measures this based on contact with the skin, we model the pulse signal based on non-contact imaging, which uses reflected light from the skin captured via frames of an RGB camera.

Given a video segment with  $S$  video frames,  $C$  color channels, and height and width dimensions  $H \times W$  pixels, assume that a time series  $\mathbf{Z} \in \mathbb{R}^{S \times K}$  has been extracted from  $K$  facial regions (described in Section V-C), where the set of  $K$  time series is assumed to capture a combination of the pulsatile signal and noise:

$$\mathbf{Z} = \mathbf{F}^{-1} \mathbf{X} + \mathbf{E}. \quad (1)$$

Here,  $\mathbf{Z} \in \mathbb{R}^{S \times K}$  is a time series of length  $S$  frames extracted from the image intensities of  $K$  facial regions of the video. These facial regions are defined as the “left cheek”, “right cheek”, “left forehead”, “right forehead”, and “chin”. The matrix  $\mathbf{F}^{-1} \in \mathbb{C}^{S \times N}$  is the oversampled inverse Fourier transform matrix,  $\mathbf{X} \in \mathbb{C}^{N \times K}$  represents the frequency coefficients of the underlying pulse signal in the  $K$  face regions, and  $\mathbf{E} \in \mathbb{R}^{S \times K}$  is a non-pulsatile noise matrix that captures noise due to multiple sources such as specular reflections, motion, and camera quantization error. The additive noise model is an approximate model: specular reflections, motion, and camera quantization error are often non-linear sources of noise that may not perfectly fit the additive noise model. However, previous work [6] has shown this to be a good approximation.

We aim to recover the pulse signal  $\mathbf{X}$  and extract the noise  $\mathbf{E}$  from the signal  $\mathbf{Z}$ . We model this as the minimization of the sum of a data fidelity term and regularizing penalty function:

$$\min_{\mathbf{X}, \mathbf{E}} \underbrace{\frac{1}{2} \left\| \mathbf{Z} - \mathbf{A} \begin{bmatrix} \mathbf{X} \\ \mathbf{E} \end{bmatrix} \right\|_F^2}_{\text{Data fidelity } = D} + \lambda \cdot \underbrace{\rho(\mathbf{X}, \mathbf{E})}_{\text{Regularization}}, \quad (2)$$

where  $\mathbf{A} = [\mathbf{F}^{-1} \ \mathbf{I}]$ . The data fidelity term,  $D$ , ensures consistency of the recovered signal with the measurements, while the regularization term imposes a prior on the signal and noise components. This optimization problem can be solved via proximal gradient descent, which alternates between a gradient update with respect to  $D$  and a proximal mapping operation corresponding to the regularizer  $\rho(\mathbf{X}, \mathbf{E})$ , such that

$$\begin{aligned} \begin{bmatrix} \tilde{\mathbf{X}}_{t+1} \\ \tilde{\mathbf{E}}_{t+1} \end{bmatrix} &\stackrel{(a)}{=} \begin{bmatrix} \mathbf{X}_t \\ \mathbf{E}_t \end{bmatrix} - \alpha \nabla_{\begin{bmatrix} \mathbf{X} \\ \mathbf{E} \end{bmatrix}} D \left( \begin{bmatrix} \mathbf{X}_t \\ \mathbf{E}_t \end{bmatrix} \right) \\ &\stackrel{(a)}{=} \alpha \mathbf{A}^T \mathbf{Z} + (\mathbf{I} - \alpha \mathbf{A}^T \mathbf{A}) \begin{bmatrix} \mathbf{X}_t \\ \mathbf{E}_t \end{bmatrix}, \\ \begin{bmatrix} \mathbf{X}_{t+1} \\ \mathbf{E}_{t+1} \end{bmatrix} &\stackrel{(b)}{=} \text{prox}_{\lambda \rho} \left( \begin{bmatrix} \tilde{\mathbf{X}}_{t+1} \\ \tilde{\mathbf{E}}_{t+1} \end{bmatrix} \right), \end{aligned} \quad (3)$$

where “(a)” defines the gradient step and “(b)” defines the proximal step for an iteration  $t$ . The proximal operator is defined as

$$\text{prox}_f(v) = \arg \min_x \left( f(x) + \frac{1}{2} \|x - v\|_2^2 \right). \quad (4)$$

When  $\rho(\cdot, \cdot)$  is explicitly defined, we may be able to calculate a closed form solution for the proximal operator. In this paper, we instead define the regularization term *implicitly*, and we present three paradigms that instantiate the proximal operators as denoising neural networks.

### IV. ALGORITHMS

Unlike previous work [6], which explicitly defines the regularizer  $\rho(\mathbf{X}, \mathbf{E})$  as in (2), we assume that the penalty term is a complicated function that can not be written in closed form. In the framework of proximal gradient descent,  $\rho(\mathbf{X}, \mathbf{E})$  is not minimized directly, but is minimized through its proximal operator as shown in Equation (4). In this paper, we

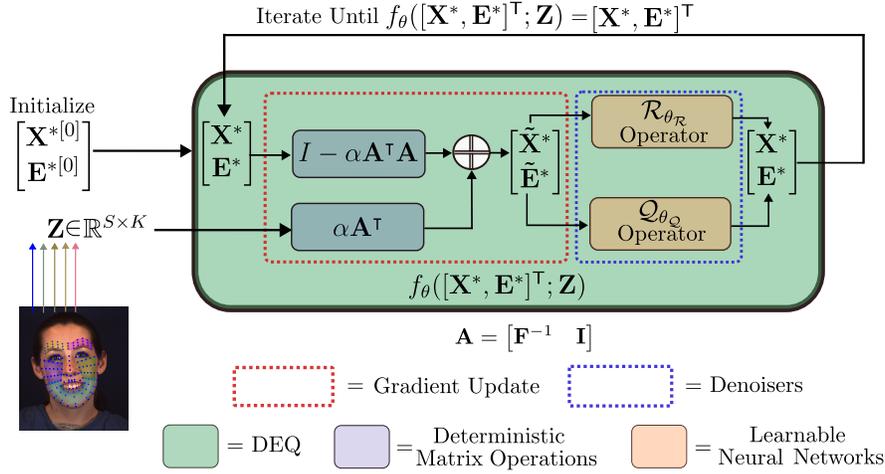


Fig. 3: Our DE-Prox-iPPG algorithmic paradigm. The video-extracted time series  $\mathbf{Z}$  and the two concatenated optimization variables  $[\mathbf{X}^*, \mathbf{E}^*]^\top$  are passed through the block outlined in red which performs gradient updates  $\frac{\partial \mathbf{D}}{\partial \mathbf{X}^*}, \frac{\partial \mathbf{D}}{\partial \mathbf{E}^*}$ . The results  $\tilde{\mathbf{X}}^*$  and  $\tilde{\mathbf{E}}^*$  are then passed through denoising operators  $\mathcal{R}$  and  $\mathcal{Q}$ . This process is repeated until fixed-point convergence.

approximate the proximal operator as a deep denoiser, proposing three methods: an unrolling algorithm (Section IV-A), a fixed-point algorithm (Section IV-B), and a combination of unrolling and fixed-point algorithms (Section IV-C).

#### A. Unrolling Algorithm: Unrolled iPPG

Previous work [6] defined explicit sparsity priors on the frequency coefficients of the pulse signal, making an assumption that the frequency coefficients are independent of each other. However, we hypothesize that sparsity (or joint sparsity) is not sufficiently expressive to characterize the inherent structure of the heartbeat signal. Rather, we explore how the complex, nonlinear combination of the frequency coefficients can be learned from data. Maintaining the algorithmic structure of proximal gradient descent, we “unroll” [26] the iteration map from Equation (2)

$$\begin{aligned} \begin{bmatrix} \tilde{\mathbf{X}}_{t+1} \\ \tilde{\mathbf{E}}_{t+1} \end{bmatrix} &\stackrel{(a)}{=} \begin{bmatrix} \mathbf{X}_t \\ \mathbf{E}_t \end{bmatrix} - \alpha \nabla_{\begin{bmatrix} \mathbf{X}_t \\ \mathbf{E}_t \end{bmatrix}} D \left( \begin{bmatrix} \mathbf{X}_t \\ \mathbf{E}_t \end{bmatrix} \right), \\ \mathbf{X}_{t+1} &\stackrel{(b)}{=} \mathcal{R}_{\theta_{\mathcal{R}}}(\tilde{\mathbf{X}}_{t+1}) \\ \mathbf{E}_{t+1} &\stackrel{(b)}{=} \mathcal{Q}_{\theta_{\mathcal{Q}}}(\tilde{\mathbf{E}}_{t+1}) \end{aligned} \quad (5)$$

where  $\mathcal{R}_{\theta_{\mathcal{R}}}(\tilde{\mathbf{X}}_{t+1})$  and  $\mathcal{Q}_{\theta_{\mathcal{Q}}}(\tilde{\mathbf{E}}_{t+1})$  are denoising functions implemented as neural networks in step (b). We graphically show the algorithm in Figure 2. Similar to proximal gradient descent, the gradient descent step ensures the recovered signal’s consistency with the measurements, while the denoising operators impose the learned signal prior. We unroll for  $T$  iterations, after which we compute a loss between the predicted signal and ground-truth (during training), or output the estimated pulse signal (during inference).

#### B. Fixed-Point Algorithm: DE-Prox-iPPG

A potential drawback of Unrolled iPPG is that the gradient and denoising steps are applied for a predetermined number of

iterations  $T$ , where  $T$  is a hyperparameter. It may be possible to apply these steps until convergence instead, obviating the need to choose  $T$ . Noting that proximal gradient descent can be viewed as a fixed-point iteration [28], [69], we propose Deep Equilibrium Proximal Gradient Descent iPPG (DE-Prox-iPPG). This method, illustrated in Figure 3, encapsulates the entire iteration map (i.e., gradient descent followed by denoising) in a DEQ and iterates these steps until a fixed point is found for both optimization variables. Compared to unrolling methods, we assume that we can find a joint fixed point for the frequency coefficients  $\mathbf{X}^*$  and noise  $\mathbf{E}^*$ . Therefore, we optimize for the fixed point of the concatenation of the two variables  $\mathbf{X}^*$  and  $\mathbf{E}^*$ , and define a DEQ  $f_{\theta}(\cdot, \cdot)$  that iterates through (a) gradient descent steps and (b) denoising operators until convergence:

$$f_{\theta} \left( \begin{bmatrix} \mathbf{X}^* \\ \mathbf{E}^* \end{bmatrix}, \mathbf{Z} \right) = \begin{cases} \begin{bmatrix} \tilde{\mathbf{X}}^* \\ \tilde{\mathbf{E}}^* \end{bmatrix} & \stackrel{(a)}{=} \begin{bmatrix} \mathbf{X}^* \\ \mathbf{E}^* \end{bmatrix} - \alpha \nabla_{\begin{bmatrix} \mathbf{X}^* \\ \mathbf{E}^* \end{bmatrix}} D \left( \begin{bmatrix} \mathbf{X}^* \\ \mathbf{E}^* \end{bmatrix} \right) \\ \mathbf{X}^* & \stackrel{(b)}{=} \mathcal{R}_{\theta_{\mathcal{R}}}(\tilde{\mathbf{X}}^*), \\ \mathbf{E}^* & \stackrel{(b)}{=} \mathcal{Q}_{\theta_{\mathcal{Q}}}(\tilde{\mathbf{E}}^*). \end{cases} \quad (6)$$

DE-Prox-iPPG computes the gradient and denoising steps until a fixed point of the *global system* is found. While Unrolled iPPG unrolls for  $T$  iterations, the number of DE-Prox-iPPG iterations varies depending on the input window, as well as on the joint fixed-point of the signal and noise variables.

#### C. Unrolling + Fixed-Point Algorithm: UDEQ-iPPG

While DE-Prox-iPPG computes the fixed-point iteration for both the pulse signal and noise, we recognize that fixed-point optimization for both signals may be suboptimal. To address this, we propose Unrolled-DEQ-iPPG (UDEQ-iPPG), which is based on the following two assumptions: 1) Given that the pulsatile signal is a realization of the quasi-periodic nature of the beating heart [29], we assume it is sufficiently structured

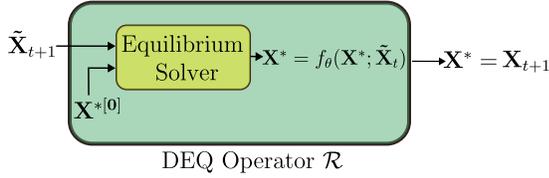


Fig. 4: The DEQ operator used in UDEQ-iPPG. At each unrolling iteration, we apply the denoising operator until a fixed-point is found.

to be represented as a fixed point of some neural network-based denoising operator; and 2) we assume that the variation in the source of the noise, such as motion or illumination, *does not* admit a structure that would be appropriately represented as the fixed point of an operator. To that end, we instantiate  $\mathcal{R}$  in Figure 2 as a DEQ but restrict  $\mathcal{Q}$  to a single forward pass of a neural network, and we embed both operators in the Unrolled iPPG structure. The corresponding iteration map is summarized as follows:

$$\begin{aligned} \begin{bmatrix} \tilde{\mathbf{X}}_{t+1} \\ \tilde{\mathbf{E}}_{t+1} \end{bmatrix} &\stackrel{(a)}{=} \begin{bmatrix} \mathbf{X}_t \\ \mathbf{E}_t \end{bmatrix} - \alpha \nabla_{\begin{bmatrix} \mathbf{X} \\ \mathbf{E} \end{bmatrix}} D \left( \begin{bmatrix} \mathbf{X}_t \\ \mathbf{E}_t \end{bmatrix} \right), \\ &\stackrel{(a)}{=} \alpha \mathbf{A}^\top \mathbf{Z} + (\mathbf{I} - \alpha \mathbf{A}^\top \mathbf{A}) \begin{bmatrix} \mathbf{X}_t \\ \mathbf{E}_t \end{bmatrix}, \\ \mathbf{X}^* &\stackrel{(b)}{=} \mathcal{R}_{\theta_{\mathcal{R}}}(\mathbf{X}^*; \tilde{\mathbf{X}}_{t+1}) \\ \mathbf{X}_{t+1} &= \mathbf{X}^*, \\ \mathbf{E}_{t+1} &\stackrel{(b)}{=} \mathcal{Q}_{\theta_{\mathcal{Q}}}(\tilde{\mathbf{E}}_{t+1}). \end{aligned} \quad (7)$$

The DEQ operator  $\mathcal{R}(\cdot, \cdot)$  for UDEQ-iPPG is illustrated in Figure 4. In each unrolling iteration, we input the gradient-updated frequency coefficients  $\tilde{\mathbf{X}}_{t+1}$  to the DEQ to produce the fixed-point variable  $\mathbf{X}^*$  using a fixed-point solver. On the other hand, only a single forward pass through  $\mathcal{Q}$  occurs for the noise variable  $\tilde{\mathbf{E}}$ . The overall procedure is unrolled over  $T$  iterations, similar to Unrolled iPPG.

#### D. Parameter Updates for Fixed-Point Networks

The fixed point in DE-Prox-iPPG can be found naively by repeated application of  $f_\theta(\cdot, \cdot)$  until  $\|f_\theta(\begin{bmatrix} \mathbf{X}^* \\ \mathbf{E}^* \end{bmatrix}; \mathbf{Z}) - \begin{bmatrix} \mathbf{X}^* \\ \mathbf{E}^* \end{bmatrix}\|$  is below some tolerance. This method, unfortunately, builds large computation graphs in memory. Equivalently, we can find the root of

$$g \left( \begin{bmatrix} \mathbf{X}^* \\ \mathbf{E}^* \end{bmatrix}; \mathbf{Z} \right) = \mathbf{X}^* - f_\theta \left( \begin{bmatrix} \mathbf{X}^* \\ \mathbf{E}^* \end{bmatrix}; \mathbf{Z} \right) = \mathbf{0}. \quad (8)$$

Similarly, for UDEQ-iPPG, we would like to find the root of  $h(\mathbf{X}^*, \tilde{\mathbf{X}}_{t+1}) = \mathbf{X}^* - f_\theta(\mathbf{X}^*, \tilde{\mathbf{X}}_{t+1}) = \mathbf{0}$ . This technique permits efficient root-solving algorithms such as Broyden’s Method [46] to solve for the fixed point.

During training, we iterate for  $T$  iterations for UDEQ-iPPG, or until fixed-point convergence for DE-Prox-iPPG, as shown in Figures 2 and 3, after which both models’ parameters must be updated after computing some loss  $\mathcal{L}$ . While  $\theta_{\mathcal{Q}}$  (denoting the parameters of  $\mathcal{Q}$ ) in UDEQ-iPPG can be updated via the standard auto-differentiation modules, updating the parameters

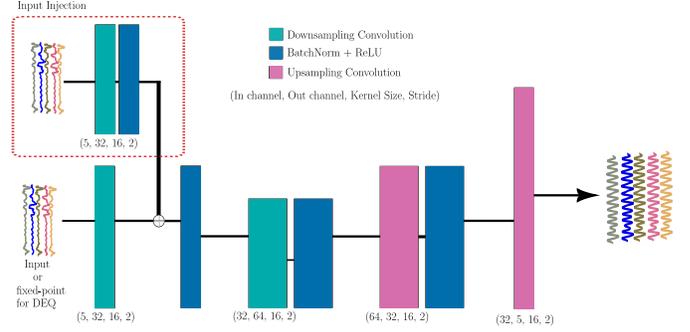


Fig. 5: The architecture diagram of both  $\mathcal{R}$  (complex-valued) and  $\mathcal{Q}$  (real-valued).

$\theta_{\mathcal{R}}$  in UDEQ-iPPG and both  $\theta_{\mathcal{R}}, \theta_{\mathcal{Q}}$  in DE-Prox-iPPG requires computing gradients with respect to the parameters. While a naïve implementation would backpropagate through iterations of the solver, a memory-efficient backpropagation scheme uses the Implicit Function Theorem (see [27] for a proof):

**Theorem 1.** Implicit Function Theorem (IFT) [27], [49], [52]. *Given the fixed-point representation  $\mathbf{p}^*$  of the optimization variables, and the corresponding loss  $\mathcal{L}(\mathbf{Z}, \mathbf{Z}_{\text{gt}})$ , where  $\mathbf{Z} = \mathbf{F}^{-1}\mathbf{X}^*$ , the gradient of the DEQ flow is given by*

$$\frac{\partial \mathcal{L}}{\partial \theta} = \frac{\partial \mathcal{L}}{\partial \mathbf{p}^*} \left( \mathbf{I} - \frac{\partial f_\theta(\mathbf{p}^*)}{\partial \mathbf{p}^*} \right)^{-1} \frac{\partial f_\theta(\mathbf{p}^*, \mathbf{Z})}{\partial \theta}. \quad (9)$$

Note that for UDEQ-iPPG,  $\mathbf{p}^* = \mathbf{X}^*$ , while for DE-Prox-iPPG it is  $\mathbf{p}^* = \begin{bmatrix} \mathbf{X}^* \\ \mathbf{E}^* \end{bmatrix}$ . We compute the mean squared error (MSE) Loss between the final output  $\mathbf{Z} = \mathbf{F}^{-1}\mathbf{X}^*$  and the ground-truth PPG signal  $\mathbf{Z}_{\text{gt}}$ . Here,  $\mathbf{X}^*$  is  $\mathbf{X}_T^*$  in UDEQ-iPPG (i.e. the output of the fixed-point operator at iteration  $T$ ) and is  $\mathbf{X}^*$  (the fixed point) in DE-Prox-iPPG. To ensure stability of the solution, we add a regularizing penalty function given by the Frobenius norm of the Jacobian with respect to the fixed-point  $\mathbf{p}^*$ , which is given by  $\mathbf{J}_{f_\theta} = \frac{\partial f_\theta(\mathbf{p}^*)}{\partial \mathbf{p}^*}$ , where  $\mathbf{J}_{f_\theta} \in \mathbb{R}^{d \times d}$  and  $d$  is the length of the vector  $\mathbf{p}^*$  [60], [62]. To compute  $\|\mathbf{J}_{f_\theta}\|_F = (\text{tr}(\mathbf{J}_{f_\theta} \mathbf{J}_{f_\theta}^\top))^{1/2}$ , we use the Hutchinson estimator [70]:

$$\text{tr}(\mathbf{J}_{f_\theta} \mathbf{J}_{f_\theta}^\top) = \mathbb{E}_{\epsilon \in \mathcal{N}(\mathbf{0}, \mathbf{I}_d)} \left[ \|\epsilon^\top \mathbf{J}_{f_\theta}\|_2^2 \right] \quad (10)$$

Our final loss is composed of the mean squared error (MSE) between the recovered waveform and the ground-truth waveform plus the Jacobian regularization penalty:

$$\mathcal{L} = \mathcal{L}_{\text{mse}}(\mathbf{Z}_{\text{final}}, \mathbf{Z}_{\text{gt}}) + \lambda \frac{\|\epsilon^\top \mathbf{J}_{f_\theta}(\mathbf{p}^*)\|_2^2}{d}, \epsilon \in \mathcal{N}(\mathbf{0}, \mathbf{I}_d), \quad (11)$$

where  $\lambda$  determines the strength of regularization.

To compute the pulse rate, we sum the power of the frequency coefficients across all regions for each bin of the signal  $\mathbf{X} \in \mathbb{C}^{N \times K}$ , where  $N$  is number of frequency bins. We select the bin with the highest power as the pulse rate.

TABLE I: Pulse rate estimation results on MMSE-HR, PURE, UBFC-rPPG, and BP4D+ datasets. Best results in each column are **bold**; second-best are underlined. Below our results are the 95% confidence intervals for the MAE and RMSE.

Type	Method	MMSE-HR			PURE			UBFC-rPPG			BP4D+			No. Params ( $\times 10^9$ ) $\downarrow$	
		MAE $\downarrow$	RMSE $\downarrow$	$\rho \uparrow$	MAE $\downarrow$	RMSE $\downarrow$	$\rho \uparrow$	MAE $\downarrow$	RMSE $\downarrow$	$\rho \uparrow$	MAE $\downarrow$	RMSE $\downarrow$	$\rho \uparrow$		
Model-based	ICA [8]	5.44	12.00	-	-	-	-	7.50	14.41	0.62	-	-	-	-	
	CHROM [10]	3.74	8.11	0.55	2.07	9.92	-	2.37	4.91	0.89	7.25	9.91	0.72	-	
	POS [32]	3.90	9.61	-	5.44	12.00	-	4.05	8.75	0.78	-	6.82	9.70	0.77	
	AutoSparsePPG('20) [6]	4.55	14.42	-	-	-	-	-	-	-	-	-	-	-	
Data-driven	HR-CNN('18) [34]	-	-	-	1.84	2.37	-	-	-	-	-	-	-	-	17.4
	SynRhythm('18) [71]	-	-	-	-	-	-	5.59	6.82	0.72	4.33	6.15	0.85	-	-
	CAN('18) [13]	4.06	9.51	-	-	-	-	-	-	-	4.08	6.11	0.86	5.3	
	CVD ('20) [72]	-	6.04	0.84	1.29	2.01	0.98	2.19	3.12	0.99	-	-	-	123.3	
	PulseGAN('21) [73]	-	-	-	-	-	-	1.19	2.10	0.98	3.84	5.37	0.89	7.6	
	InverseCAN('21) [17]	2.27	4.90	-	-	-	-	-	-	-	-	-	-	10.7	
	Gideon('21) [22]	-	-	-	2.10	2.60	0.99	3.60	4.60	0.95	4.09	5.60	0.85	8.5	
	DualGAN('22) [74]	-	-	-	0.82	1.31	0.99	<b>0.44</b>	<b>0.67</b>	<b>0.99</b>	-	-	-	76.8	
	Physformer('22) [16]	2.84	5.36	-	-	-	-	-	-	-	-	-	-	73.8	
	Federated('22) [21]	2.99	-	0.79	-	-	-	2.00	4.38	0.93	-	-	-	5.3	
	EfficientPhys-C('23) [20]	2.91	5.43	0.86	-	-	-	-	-	-	-	-	-	4.6	
	Yue('23) [24]	-	-	-	1.23	2.01	0.99	0.58	0.94	<b>0.99</b>	3.22	<b>4.47</b>	<b>0.91</b>	8.5	
	Contrastphys+('24) [23]	1.11	3.83	0.96	0.48	0.98	0.99	0.50	0.84	<b>0.99</b>	-	-	-	8.5	
Hybrid	<b>Unrolled iPPG (Ours)</b>	<u>1.11</u>	2.97	0.97	0.41	0.88	0.99	0.92	1.84	0.98	3.25	5.78	0.88	<b>1.4</b>	
	<b>DE-Prox-iPPG (Ours)</b>	[0.707, 1.60]	[1.86, 3.65]		[0.28, 0.63]	[0.68, 1.10]		[0.78, 1.32]	[1.64, 2.03]		[2.88, 4.03]	[4.91, 7.01]		<b>1.4</b>	
	<b>UDEQ-iPPG (Ours)</b>	1.22	<u>2.80</u>	<u>0.95</u>	<b>0.13</b>	<b>0.28</b>	0.99	2.86	3.84	0.97	5.01	6.95	0.77	<b>1.4</b>	
		[1.00, 1.99]	[2.03, 3.61]		[0.10, 0.40]	[0.20, 0.89]		[1.65, 4.52]	[2.68, 5.73]		[4.25, 5.33]	[5.99, 7.87]			
		<b>0.95</b>	<b>2.53</b>	<b>0.98</b>	<u>0.18</u>	<u>0.35</u>	0.99	<u>0.46</u>	<u>0.82</u>	<b>0.99</b>	<b>3.14</b>	<u>5.35</u>	0.89	<b>1.4</b>	
		[0.580, 1.38]	[1.55, 3.36]		[0.11, 0.42]	[0.21, 0.94]		[0.12, 0.93]	[0.34, 1.28]		[2.73, 3.96]	[4.71, 6.22]			

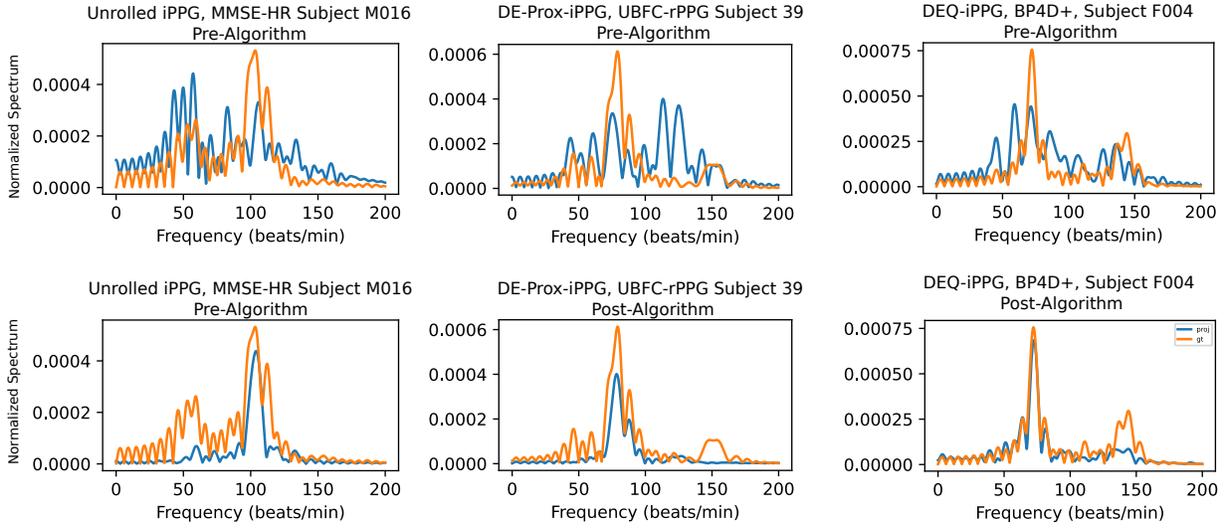


Fig. 6: Comparing the input signal spectrum to results from Unrolled iPPG, UDEQ-iPPG, and DE-Prox-iPPG on the MMSE-HR dataset for one time window. The **orange** waveform in all graphs is the ground truth pulse signal spectrum. The input signals (produced by our time-series extraction module) are shown in **blue** in the top row, while the recovered signal spectrum output by our algorithms are shown in **blue** in the bottom row. All signal recovery algorithms recover the correct peak frequency.

## V. IMPLEMENTATION AND EXPERIMENTAL RESULTS

### A. Datasets

We evaluate our algorithm on four datasets, described below.

**MMSE-HR** [75], [76]: The MMSE-HR dataset is a multi-modal dataset with synchronized facial video and physiological monitoring. Facial video was captured at a resolution of  $1040 \times 1392$  at 25 FPS (frames per second) using the Di3D dynamic imaging system, while the blood pressure waveform was captured by a finger cuff using the Biopac NIBP100D recording at 1000 Hz (downsampled to 25Hz) and calibrated using an arm cuff. A total of 102 videos from 17 male and 23 female subjects were captured as various emotions (happiness, surprise, embarrassment, fear, etc.) were elicited, resulting in substantial head motion.

We train on 10 s windows, shifting the time window by a stride of 2.4 s to obtain each new window. For evaluation on the MMSE-HR dataset, we follow the evaluation protocol of

[17], [77] concatenating three 10 s windows to obtain a 30 s window for evaluation. We perform leave-one-subject-out cross-validation and evaluate on 30 s, non-overlapping time windows.

**PURE** [78]: The PURE dataset consists of 10 subjects who perform 6 different motion tasks, from “no motion” to “fast translation” and “rotation”. The camera recorded at 30 FPS at a resolution of  $640 \times 480$ , while the corresponding pulse oximeter recorded at 60 Hz. The pulse oximeter was down-sampled to match the rate of the video. We followed the train, validation, and test splits of [34] to evaluate our algorithm, using 30 s, non-overlapping time windows as in [23] during test time. We train on 10 s windows, shifting the time window by a stride of 2.4 s to obtain each new window.

**UBFC-rPPG** [79]: Captured using a webcam recording at 30 FPS with a resolution of  $640 \times 480$ , the UBFC-rPPG dataset contains 43 videos of subjects playing a time-sensitive

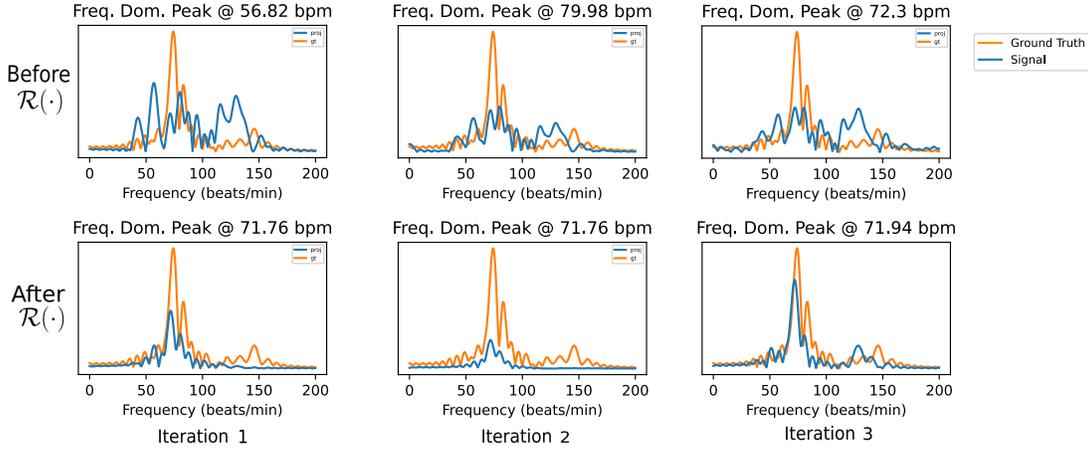


Fig. 7: We show the spectrum of our recovered signal from facial video (blue) and the ground truth signal (orange) over three unrolling iterations on the PURE dataset using UDEQ-iPPG. For each iteration, the blue signal in the top row corresponds to  $\tilde{\mathbf{X}}_t$  before  $\mathcal{R}(\cdot)$  acts upon it, while the blue signal in the bottom row corresponds to  $\mathbf{X}_t$  after  $\mathcal{R}(\cdot)$  as in Figure 2. After multiple iterations, our algorithm captures the peak frequency and a harmonic while attenuating noise.

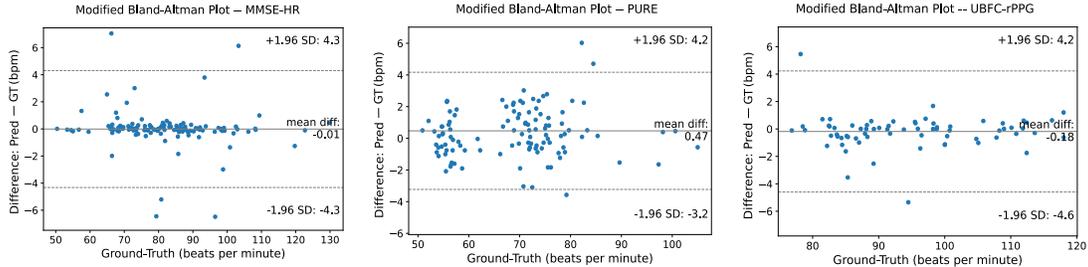


Fig. 8: Bland-Altman Analysis between predicted and ground-truth pulse rate estimates for the MMSE-HR, PURE, and UBFC-rPPG datasets using the UDEQ-iPPG algorithm. Each time window is represented by one point, which plots the difference between the predicted and ground-truth pulse rates vs. the ground-truth pulse rate. Note that the 95% limits of agreement, defined as the mean difference  $\pm 1.96 \times$  the standard deviation, are within 5 bpm of the ground-truth pulse rate on all datasets.

mathematical game to induce pulse rate changes. The pulse signal was captured using a pulse oximeter finger clip sensor recording at 30 Hz. To evaluate our algorithm on this dataset, we follow the protocols of [74], which evaluated on non-overlapping 10 s time windows per video, and averaged these windows to get a single pulse rate estimate for each video. They then computed the pulse rate metrics over the 12 pulse rate estimates (corresponding to the 12 videos in the test set). We train on 10 s windows, shifting the time window by a stride of 2.4 s to obtain each new window.

**BP4D+** [75], [76]: The BP4D+ dataset, a superset of the MMSE-HR dataset, contains 140 subjects (82 female, 58 male) who recorded a total of 548 videos while various emotions were elicited. The video was recorded at  $1092 \times 1340$ px and 25 FPS. This dataset contains significant head motion and illumination changes. We follow the evaluation procedure of [16] and evaluate on 5-fold subject-wise cross validation.

### B. Evaluation Metrics

To compute the pulse rate, we first multiply our pulse signal window by a Hanning window. Next, letting  $L = 100 \times$  signal length, we compute the  $L$ -point FFT, and square

the magnitudes of the frequency coefficients to get the power spectra. We sum the power for each frequency bin across all regions, and select the bin with the highest power as our pulse rate estimate. For pulse rate estimation, we report the mean absolute error (MAE) and root mean squared error (RMSE) between the predicted pulse rate and ground-truth pulse rate in a time window, measured in beats per minute as in [2]. We define the MAE and RMSE as:

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |R_i - \hat{R}_i|, \text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (R_i - \hat{R}_i)^2}, \quad (12)$$

where  $N$  is the number of time windows,  $R_i$  is the ground-truth pulse rate, and  $\hat{R}_i$  is the predicted pulse rate in the time window. In addition, we measure the correlation coefficient  $\rho$  between the predicted and ground-truth pulse rate estimates as in [2]. As it is also important to know what percentage of our estimates are sufficiently close to the ground-truth pulse rates, we measure the PTE6 metric [77], defined as the percentage of the time that the estimated pulse rate is within 6 bpm of the ground-truth pulse rate:

$$\text{PTE6} = \frac{100}{N} \sum_{i=1}^N P_i, \text{ where } P_i = \begin{cases} 1, & \text{if } |R_i - \hat{R}_i| < 6 \\ 0, & \text{otherwise.} \end{cases} \quad (13)$$

For a fair comparison, we use 30 s time windows for evaluation on the MMSE-HR and PURE datasets as in [23], and 10 s time windows on the UBFC-rPPG dataset as in [74]. Note that the results for [23] on the UBFC-rPPG dataset were reevaluated to match the evaluation protocol of other methods such as [74] on UBFC-rPPG.

We compute PRV metrics following the work of [23], in which we measure the low-frequency power, high-frequency power, and ratio of the low-frequency and high-frequency power of the interbeat interval signal [3]. We compute the PRV metrics on 60 sec time windows.

### C. Time Series Extraction

We generate the  $S$ -frame,  $K$ -channel time series  $\mathbf{Z} \in \mathbb{R}^{S \times K}$  from the video. To do so, we first detect the face using the FaceBoxes [80] and crop the face; next, we pass these crops to the LUVLi landmark detection [81] module, which generates 68 landmark points on the face as shown in Figure 1. These landmarks correspond to the eyebrows, nose, chin, and lips. To extract a pulsatile signal from the skin, we first interpolate new landmarks from the initial 68 landmarks to obtain an additional 77 landmarks that span the cheeks, chin, and forehead. This is shown in Figure 1; for example, using a landmark on the left cheek and corresponding lateral landmark on the nose, we define a line and interpolate 5 landmarks between them on the left cheek. We do this for landmarks across the chin, cheeks, and eyebrows. We then incorporate two design choices: 1) We segment the face into  $K = 5$  super-regions, and average each RGB color channel’s pixel intensity values in these regions as described in [77], and 2) we compute the ratio of the red color channel to green color channel as described in [10]. Both of these steps improve signal strength and reduce motion noise. Doing this for all frames, we obtain a time series  $\mathbf{Z} \in \mathbb{R}^{S \times K}$ .

After extracting the signals, each signal from the 5 spatial regions is convolved with a fifth-order Butterworth filter with cutoff frequencies of 0.7 Hz and 2.5 Hz as in [17]. Finally, the signal is AC/DC normalized by dividing by the mean signal after subtracting the mean signal:  $\hat{z}_k = \frac{z_k - \mu_k}{\mu_k}$ , where  $z_k$  is the signal in super-region  $k$ , and  $\mu_k$  is the mean of that same signal. These signals are segmented into windows of  $S$  samples and arranged as  $\mathbf{Z} \in \mathbb{R}^{S \times K}$  for pulse rate estimation.

### D. Implementation Details

The architecture diagram for  $\mathcal{R}$  and  $\mathcal{Q}$  is shown in Figure 5. Deep equilibrium models (DEQs) are instantiated as a collection of convolutions, linear operators, non-linearities, etc. organized as a single DEQ layer, and it is through this layer that the input is passed until convergence of a root-finding algorithm. To ensure a fair comparison, we adopt the same architecture throughout this paper. The DEQ models contain an extra MLP as an input injection network to ensure stability during training [60], [62]. Note that both  $\mathcal{R}(\cdot)$  and

$\mathcal{Q}(\cdot)$  have an identical architecture, with the only difference that the parameters of  $\mathcal{R}(\cdot)$  are complex (since  $\mathcal{R}(\cdot)$  operates in the frequency domain) while those of  $\mathcal{Q}(\cdot)$  are real.

All three models are trained using a batch size of 100. On the MMSE-HR and BP4D+ datasets, we train for 10 epochs using UDEQ-iPPG and 25 epochs with DE-Prox-iPPG, with an initial learning rate set to  $3 \times 10^{-4}$  and reduced by half at the 10th epoch. For the PURE dataset, we train for 10 epochs on each model. On the UBFC dataset, we found using cross-validation that training for 45 epochs provided the optimal performance. We specifically use the Broyden solver [46] to calculate the gradients instead of calculating the inverse Jacobian from Theorem 1. The backpropagation is viewed through the lens of a vector-Jacobian product: instead of explicitly constructing the Jacobian matrix, we formulate the backward pass as solving a linear system. We solve for the adjoint vector  $v$  that satisfies  $v^\top (I - J_{f_\theta}) = \frac{\partial \mathcal{L}}{\partial p^*}$ . This system can be solved efficiently and iteratively via Broyden’s method, and maintains the inverse Jacobian using a low-rank approximation. This allows us to compute gradients efficiently using vector-Jacobian products. The Jacobian regularization strength  $\lambda$  in Equation (11) is set stochastically [62] during the training iterations using a random variable  $x \sim \text{Bernoulli}(0.5)$ , where  $\lambda(x) = 5$  if  $x = 1$ , otherwise  $\lambda(x) = 0$ . All parameters including learning rates (range:  $[10^{-6}, 10^{-1}]$ ), training epochs (range:  $[0, 50]$ ), Jacobian regularization strength (range:  $\lambda \in [0, 10]$ ), frame stride between adjacent overlapping training samples (range:  $[0.1, 5]$  seconds), learning rate decay epoch (range:  $[5, 50]$  epochs), selection of solver (Anderson vs Broyden), and weight decay were determined by cross validation using the grid search procedure. On the MMSE-HR dataset we used one-vs-all cross validation in accordance with [17]; on the PURE dataset, we performed cross-validation using the validation set defined in [34]; and for both the UBFC-rPPG and BP4D+ datasets we used 5-fold cross-validation, in accordance with [23] and [24] respectively.

### E. Experimental Results

1) *Comparison to the State-of-the-Art*: The pulse rate results on the MMSE-HR, PURE, UBFC-rPPG, and BP4D+ are shown in Table I. We see strong denoising performance for all three of our models. Our models yield results consistent with the state-of-the-art on both the MMSE-HR and PURE datasets, exhibiting good performance at a lower number of trainable parameters as compared to deep learning algorithms such as [20], [23]. On the UBFC-rPPG dataset, UDEQ-iPPG nearly matches the performance of DualGAN [74], but does so with 50 times fewer trainable parameters. On the larger BP4D+ dataset, our performance is similar to previous methods, but uses a much fewer parameters. We achieve performance consistent with the state-of-the-art for the MAE between predicted and ground-truth pulse rates using UDEQ-iPPG, but occasionally our method has large error resulting in a higher RMSE. We notice larger errors when we compare UDEQ-iPPG and Unrolled iPPG against DE-Prox-iPPG on the UBFC-rPPG and BP4D+ dataset. We surmise that performance is less accurate on a select few examples due to the difficulty

TABLE II: Cross-dataset model evaluation

Train	Test	Model	MAE (bpm) ↓	RMSE (bpm) ↓	PTE6 (%) ↑
MMSE-HR	PURE	Unrolled iPPG	<b>0.25</b>	<b>0.58</b>	<b>100</b>
		DE-Prox-iPPG	1.46	2.36	96.49
		UDEQ-iPPG	0.44	1.04	100
PURE	MMSE-HR	Unrolled iPPG	3.88	11.54	86.82
		DE-Prox-iPPG	4.47	12.02	86.26
		UDEQ-iPPG	<b>2.38</b>	<b>8.18</b>	<b>89.92</b>
UBFC-rPPG	MMSE-HR	Unrolled iPPG	5.62	10.23	73.64
		DE-Prox-iPPG	5.86	14.18	82.01
		UDEQ-iPPG	<b>3.72</b>	<b>8.78</b>	<b>85.01</b>
MMSE-HR	UBFC-rPPG	Unrolled iPPG	2.32	8.46	90.69
		DE-Prox-iPPG	2.34	8.69	91.01
		UDEQ-iPPG	<b>0.42</b>	<b>1.26</b>	<b>100</b>

of learning the fixed point of the signal and noise together. We argue that learning the fixed-point of the signal is easier due to its inherent structure, while the noise is less constrained and does not admit a fixed point structure. Learning the fixed point of these variables together during training is a challenging optimization problem that results in inferior performance compared to UDEQ-iPPG or Unrolled-iPPG. Our algorithms’ strengths are shown qualitatively in Figure 6, in which we show the spectrum of the ground-truth (orange) and one facial region of our measurements  $\mathbf{Z}$  (top), and recovered signal from our algorithm (bottom). In each case, we show that Unrolled iPPG, DE-Prox-iPPG, and UDEQ-iPPG are able to recover a clean signal from a noisy signal.

During inference, our methods show rapid denoising and convergence to a fixed point, as shown in Figure 7, in which the orange signal is the frequency spectrum of the ground-truth signal, and the blue signal is the spectrum recovered by UDEQ-iPPG from facial video. We show the recovered signal over three unrolling iterations of UDEQ-iPPG. The blue signal in the top row shows  $\tilde{\mathbf{X}}_t$  from Figure 2, i.e., the signal after the gradient step but before the DEQ denoising step, while the bottom row shows  $\mathbf{X}_t$  after  $\mathcal{R}(\cdot)$  operates upon it. Before the first iteration, the signal extracted from the face is noisy with an incorrect peak pulse-rate. In successive iterations, we notice the DEQ’s ability to denoise the signal while the gradient step ensures consistency with the (noisy) extracted signal. By the end of three unrolling iterations, we successfully recover a cleaner signal.

2) *Statistical Analysis*: To confirm the validity of our method, we perform a modified Bland-Altman analysis (also known as the Tukey Mean-Difference plot) between the estimated pulse rates and the ground-truth rate pulse rates for the UDEQ-iPPG method. In Figure 8, we plot the difference between the predicted pulse rate and ground-truth pulse rate against the ground-truth pulse rate for the MMSE-HR, PURE, and UBFC-rPPG dataset. Each datapoint represents a single 30 s time window for the MMSE-HR dataset, a 10 s time window for the PURE dataset, and a 10 s time window for the UBFC-rPPG dataset. For all three datasets, our mean difference is close to zero, showing our method’s agreement with the ground-truth. Furthermore, the 95% limits of agreement, defined as the mean difference  $\pm 1.96 \times$  the standard deviation, are less than or equal to 4.6 bpm from the ground-truth pulse rate on all datasets. This shows strong performance, particularly over a range of pulse rates from 50 bpm to 130 bpm, confirming the effectiveness of our method.

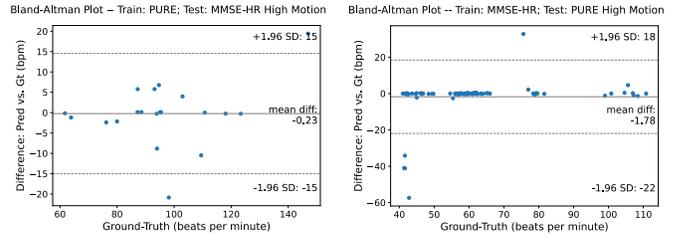


Fig. 9: Bland Altman analysis for cross dataset scenario when testing on “high-motion” videos only. In both scenarios UDEQ-iPPG performs quite well, except for a few outliers. We analyze failure cases in Figure 11.

3) *Cross-Dataset Model Evaluation*: Our methods’ learned denoisers can generalize to unseen data; we show this in Table II, in which we train and test on different datasets. UDEQ-iPPG has an MAE of less than 4 bpm in all scenarios, while Unrolled iPPG does well in most scenarios. We notice that UDEQ-iPPG transfers more effectively to other datasets. We computed the Wilcoxon signed-rank test on the absolute prediction errors of our methods to determine if UDEQ-iPPG statistically outperforms Unrolled iPPG or DE-Prox-iPPG. As shown in Table II, we observed a significant improvement ( $p < 0.05$ ) in all cases except for the cross-domain scenario of training on MMSE-HR and testing on PURE. In this specific case, we failed to reject the null hypothesis ( $p > 0.05$ ), indicating that UDEQ-iPPG performs comparably to the baseline models on the PURE dataset. Our algorithms transfer well to other datasets, though they suffer expected performance decreases when the domain shifts. We hypothesize that this happens because we learn *the prior* associated with each dataset. Since this prior is data-driven, it may not capture the recording scenarios, noise, etc. of another dataset.

To evaluate our algorithm’s robustness to motion and illumination, we identify the test videos in the MMSE-HR and PURE datasets that exhibit the highest motion, and we perform a cross-dataset evaluation on these high-motion subsets (i.e., the *most difficult* videos of each dataset). On both the MMSE-HR and PURE datasets, we define the “high-motion” videos of each subject as the videos of that subject in which the standard deviation of facial landmark positions (averaged over all landmarks and all frames) is highest. In the Appendix Section VII, we explain exactly which videos belong to each set.

In Figure 9, we plot the Bland Altman analysis when training UDEQ-iPPG on PURE and testing on the high-motion test videos of MMSE-HR. We similarly plot the analysis when we train on MMSE-HR and test on the high-motion videos of PURE. In both scenarios, we are performing a cross-dataset evaluation on the *most difficult* test samples. Through the Bland Altman analysis, we notice that most of our pulse rate estimates are close to the ground-truth, with low average error and occasional incorrect estimates. This shows that our method generalizes well to significant changes in motion and illumination, even in the cross-dataset scenario.

That being said, domain shift challenges exist. As enumerated by [82], distributions shift due to changes in *movement*,

cameras, illumination, and skin tones between recording setups. All these must be addressed either explicitly or implicitly to build a generalizable rPPG model. To address this, the work of [82] imposed priors on their model such that the iPPG signal agrees with the priors. Other work [83] optimized iPPG features by aligning the domains using contrastive learning and also provided more coverage over out-of-distribution samples by feature space regularization. Other work [84] viewed domain shifts as “noise” and used adversarial noise generation in the source domain to learn better rPPG features. The key idea behind addressing domain shifts would be to model the distributions of movement, illumination, and skin tones, and then find a way to minimize the “distance” between these distributions. We leave this as future work.

4) *Pulse Rate Variability Analysis*: We perform a Bland-Altman analysis between the predicted LFNU, HFNU, and LF/HF against the ground-truth for both ContrastPhys+ and UDEQ-iPPG as shown in Figure 10. We reduce the number of parameters in ContrastPhys+ to match that of UDEQ-iPPG. We notice that we underestimate the ground-truth low-frequency power by about -0.03 normalized units (and analogously overestimate the high-frequency power by +0.03 normalized units), which shows good agreement with the ground-truth. We see similar results for the LF/HF ratio, albeit with higher error. Over all three metrics, we notice that UDEQ-iPPG performs better than ContrastPhys+ when observing the absolute mean difference. We surmise both methods’ performance on PRV estimation can be improved in two ways: 1) Using more complicated networks with more parameters—particularly for the noise  $\mathcal{Q}(\cdot)$ —could result in better fitting of the iPPG signal in high-noise scenarios, and 2) Designing specific loss functions for PRV estimation could produce better inter-beat interval signals, crucial for PRV estimation. We leave these explorations as future work.

### F. Failure cases and potential for improvement

While our method successfully captures the pulse rate in many scenarios, we do see failure cases as shown in Figure 11. We plot the spectrum for both the measurements/output (blue) and the ground-truth spectrum (orange). On the PURE test dataset, we see that the spectrum of the measurements is broad; therefore, our algorithm struggles to recover the peak frequency and/or harmonics, and outputs a flat spectrum. In the middle and right examples, the measurements contain higher power at an incorrect frequency; the algorithm believes this is the fundamental frequency and boosts these frequencies. A proper way to address these failure cases would be uncertainty quantification, which would highlight instances in which the output of the algorithm is unreliable. A posterior sampling method would be appropriate, which we leave to future work.

### G. Ablation Studies

1) *Model Class: DEQ versus Traditional Network*: Firstly, we explore the effect of implementing the learned denoising operators as DEQs and traditional neural networks in UDEQ-iPPG. We examine the four algorithmic configurations for model type where  $\mathcal{R}(\cdot) \in \{\text{DEQ}, \text{NN}\}$ ,  $\mathcal{Q}(\cdot) \in \{\text{DEQ}, \text{NN}\}$ ,

TABLE III: Varying the type of model used for each denoiser, and the results on three datasets.

Dataset	Method	$\mathcal{R}(\cdot)$	$\mathcal{Q}(\cdot)$	MAE	RMSE	PTE6
				bpm ↓	bpm ↓	% ↑
MMSE-HR	—	DEQ	DEQ	1.21	3.02	93.02
	UDEQ-iPPG	DEQ	NN	<b>0.95</b>	<b>2.53</b>	<b>94.57</b>
	—	NN	DEQ	1.36	3.2	90.70
	Unrolled iPPG	NN	NN	1.12	2.80	93.02
PURE	—	DEQ	DEQ	0.30	0.80	100
	UDEQ-iPPG	DEQ	NN	<b>0.18</b>	<b>0.35</b>	100
	—	NN	DEQ	0.29	0.63	100
	Unrolled iPPG	NN	NN	0.41	0.62	100
UBFC-rPPG	—	DEQ	DEQ	7.10	10.10	63.64
	UDEQ-iPPG	DEQ	NN	<b>0.46</b>	<b>0.82</b>	100
	—	NN	DEQ	6.15	9.45	83.33
	Unrolled iPPG	NN	NN	0.92	1.84	90.91

where “DEQ” is the fixed-point instantiation of the model in Figure 5, while “NN” refers to a single application of the same model. The results are in Table III. The results show that the UDEQ-iPPG configuration, in which the signal  $\mathbf{X}$  is modeled as the output of a fixed-point function ( $\mathcal{R} = \text{DEQ}$ ) and the noise  $\mathbf{E}$  is modeled as unconstrained neural network ( $\mathcal{Q} = \text{NN}$ ), achieves performance consistent with the state-of-the-art. The PURE dataset shows similar results; all model combinations perform well, but again the UDEQ-iPPG configuration performs best. On the UBFC-rPPG dataset, we see a clear distinction: when modeling the noise signal as a DEQ, our model struggles to converge. Both the Unrolled iPPG configuration ( $\mathcal{R} = \text{NN}$ ,  $\mathcal{Q} = \text{NN}$ ) and the UDEQ-iPPG configuration ( $\mathcal{R} = \text{DEQ}$ ,  $\mathcal{Q} = \text{NN}$ ) show good performance. Our intuition is that the pulse signal has fixed-point structure, while the noise does not; therefore, instantiating  $\mathcal{R}(\cdot)$  as a DEQ and instantiating the noise  $\mathcal{Q}$  as a traditional neural network learns best.

2) *Modeling the noise explicitly*: To determine the effect of explicitly modeling the noise, we examine training with  $\mathcal{R}(\cdot)$  only, i.e., solving the optimization problem

$$\min_{\mathbf{X}, \mathbf{E}} \frac{1}{2} \|\mathbf{Z} - \mathbf{F}^{-1}\mathbf{X}\|_F^2 + \lambda \cdot \rho(\mathbf{X}). \quad (14)$$

We show the results in Table IV. It is clear that Unrolled iPPG, DE-Prox-iPPG, and UDEQ-iPPG all benefit from modeling the noise in addition to the signal.

3) *Computational Efficiency*: In addition to parameter efficiency, we measure the number of GFLOPs during the forward pass of our models, and compare it to Physformer [16] and ContrastPhys+ [23], very recent state-of-the-art algorithms for pulse signal recovery. Unrolled iPPG has the lowest number of forward propagation GFLOPs because forward propagation through the model occurs  $T = 3$  times. For both UDEQ-iPPG and DE-Prox-iPPG, we compute the forward propagation until a fixed point is reached for our signal. This results in a higher number of Forward GFLOPs relative to Unrolled iPPG. Compared to ContrastPhys+, Unrolled iPPG has a  $100\times$  reduction in GFLOPs, while UDEQ-iPPG and DE-Prox-iPPG have  $13\times$  and  $15\times$  reduction in GFLOPs, respectively. We achieve these reductions while achieving results comparable and consistent with the state-of-the-art.

4) *Interpretability*: Finally, we seek to understand how our denoisers make their prediction. Inspired by Grad-CAM [85], we examine the last convolution layer of our  $\mathcal{R}(\cdot)$  network

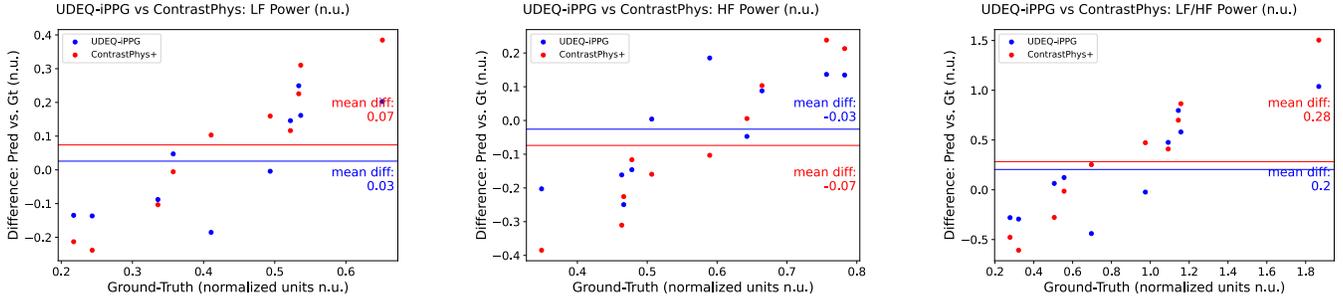


Fig. 10: Comparing the predicted PRV against the ground-truth PRV metrics for both ContrastPhys+ [23] and UDEQ-iPPG. We achieve lower absolute errors across LF, HF, and LF/HF metrics when normalizing the number of learnable parameters across models.

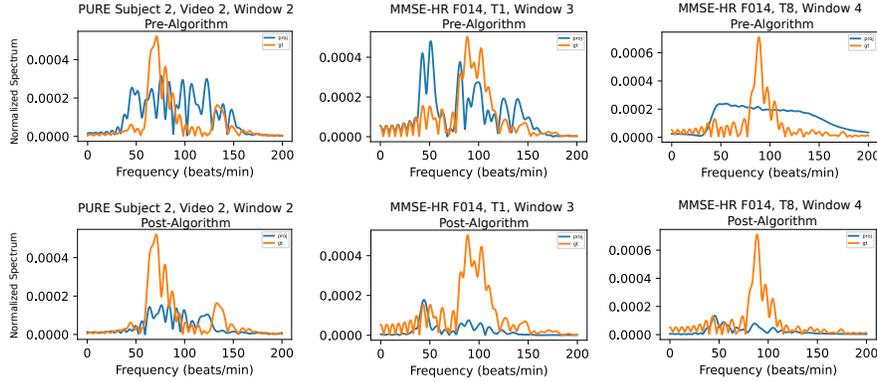


Fig. 11: Failure cases of our algorithm. We plot the spectrum before and after our algorithm; the orange signal is the ground-truth while the blue signal is our signal measurements/output. We notice that when our measurements contain a broad spectrum/low SNR, or when there is a peak frequency at the incorrect frequency, our method can struggle.

TABLE IV: Pulse rate estimation on the MMSE-HR dataset with and without modeling the noise component  $\mathbf{E}$ , where “ $\times$ ” solves Eq. (14) and “ $\checkmark$ ” solves Eq. (2).

Method	$\mathcal{Q}(\cdot)?$	MAE (bpm) $\downarrow$	RMSE (bpm) $\downarrow$	PTE6 (%) $\uparrow$
Unrolled iPPG	$\times$	1.44	3.87	92.25
	$\checkmark$	<b>1.11</b>	<b>2.82</b>	<b>93.02</b>
DE-Prox	$\times$	1.50	3.29	90.70
	$\checkmark$	<b>1.22</b>	<b>2.80</b>	<b>93.80</b>
UDEQ-iPPG	$\times$	1.70	6.68	92.24
	$\checkmark$	<b>0.95</b>	<b>2.53</b>	<b>94.57</b>

TABLE V: Counting the Forward Propagation and Backward Propagation GFLOPs for each method

Type	Forward Propagation GFLOPs
Physformer [16]	237.08
ContrastPhys+ [23]	4594.96
Unrolled iPPG	34.64
UDEQ-iPPG	332.84
DE-Prox-iPPG	304.53

in Unrolled iPPG and examine the gradients of the output with respect to the input. The main idea is that the magnitude of the gradient of the model’s output with respect to the last convolutional layer’s input should have the following two properties: 1) It should be large at the ground-truth frequency

with respect to the other frequencies in the first unrolling iteration, in order to regress to the fundamental features; and 2) It should be more well-distributed across frequencies in the last unrolling iteration, as the network learns the finer details. We visualize this idea in Figure 12, which is analogous to Figure 7 except that we include the gradient signal of the output with respect to the input of the network at each iteration. We notice in the first iteration that the network is primarily looking to capture the fundamental frequency and its first harmonic, generating near-zero gradients at other frequencies. By the third unrolling iteration, the network has identified the fundamental frequency, and generates larger gradients at other frequencies for a better signal fit. Future work could explicitly evaluate the learned regularization prior at different test samples [86] in the domain of the function to understand how the range evolves.

## VI. CONCLUSION

In this paper, we model pulse signal recovery from facial video as an inverse problem with learnable priors and propose solutions using unrolling methods, fixed-point methods, and a combination of unrolling and fixed-point methods. In our first method, we apply successive gradient descent and neural network-based denoising steps on the pulse signal and noise for a *fixed number* of iterations—known commonly as

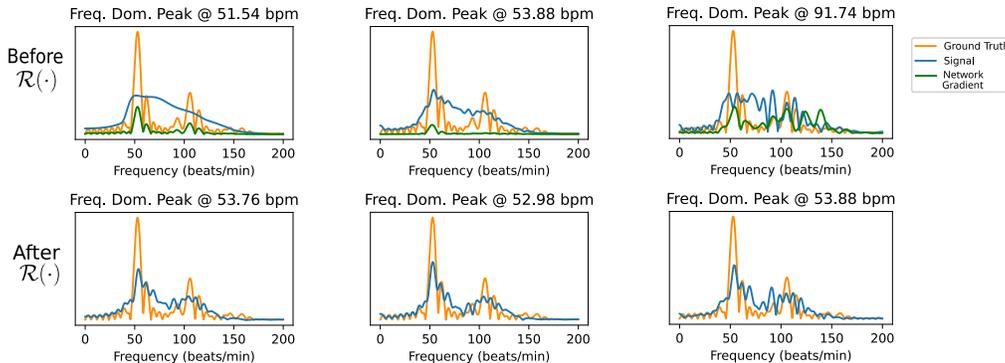


Fig. 12: We plot the gradient of the output of the  $\mathcal{R}(\cdot)$  network with respect to the input in green to demonstrate “where” the Unrolled iPPG network is making the most changes. In the first iterations of the unrolling, the network promotes the fundamental frequency and its first harmonic (i.e., the fundamental components of the signal) while ignoring other frequencies. After obtaining the key components, the network then generates higher gradients for frequencies unrelated to the peak frequency.

“unrolling” proximal gradient descent—and call our method Unrolled iPPG. Instead of applying a *fixed* number of unrolling steps, Deep-Equilibrium-Proximal-iPPG (DE-Prox-iPPG), our second method, iterates both the gradient descent and denoising steps until fixed-point convergence. Finally, we combine unrolling and fixed-point methods in Unrolled-DEQ-iPPG, which applies the neural network-based denoising operator on the pulse signal until fixed-point convergence in each unrolling step. Our results show performance consistent with the state-of-the-art, while all three methods show good performance using significantly fewer parameters than previous methods.

## VII. APPENDIX

### A. Data splits for results in Table I

Below, we enumerate the dataset splits used to train and evaluate our model

- **MMSE-HR [75]:** The dataset contains forty subjects who performed between two and six tasks each. To compare fairly with previous methods, we follow the work of [17] and perform **subject-independent** cross validation i.e. we select thirty-nine subjects for training and one subject for validation and record the estimated and ground-truth pulse rates. We repeat this procedure a total of forty times so that each subject is validated. We then compute the MAE, RMSE, PTE6, and  $\rho$  aggregated over all forty folds of cross-validation.
- **PURE [34]:** We follow the dataset splits proposed by the authors of the dataset, which are repeated below:
  - **Train:** 01-01, 01-02, 01-03, 01-04, 01-05, 01-06, 04-01, 04-02, 04-03, 04-04, 04-05, 04-06, 05-01, 05-02, 05-03, 05-04, 05-05, 05-06, 07-01, 07-02, 07-03, 07-04, 07-05, 07-06, 08-01, 08-02, 08-03, 08-04, 08-05, 08-06.
  - **Validation:** 06-01, 06-02, 06-03, 06-04, 06-05, 06-06, 09-01, 09-02, 09-03, 09-04, 09-05, 09-06.
  - **Test:** 02-01, 02-02, 02-03, 02-04, 02-05, 02-06, 03-01, 03-02, 03-03, 03-04, 03-05, 03-06, 10-01, 10-02, 10-03, 10-04, 10-05, 10-06.

- **UBFC-rPPG:** This dataset contains 42 subjects. We did not use some subjects at training due to corrupted ground-truth. During testing, we used all subjects reported by previous work for a fair comparison. We use the following splits:

- **Train/Validation:** Subject1, Subject3, Subject4, Subject5, Subject6, Subject9, Subject10, Subject12, Subject 13, Subject14, Subject15, Subject16, Subject17, Subject22, Subject23, Subject30, Subject31, Subject32, Subject33, Subject34, Subject35, Subject36, Subject37
- **Test:** Subject38, Subject39, Subject40, Subject41, Subject42, Subject43, Subject44, Subject45, Subject46, Subject47, Subject48, Subject49

- **BP4D+ [75], [76]:** This dataset contains 140 subjects, totaling 548 videos. Previous literature evaluated over 5 folds of cross-validation but did not describe their splits. We perform subject-wise five-fold cross-validation using a fixed random seed, and report the aggregate results.

### B. Data splits for test sets of Figure 9

The high-motion videos for the MMSE-HR and PURE datasets, which are the test sets used to generate Figure 9, are listed below. For the MMSE-HR dataset, we select the video, per subject, in which the standard deviation of landmark positions (averaged over all landmarks and all frames of the video) is highest. For the PURE dataset, videos with the highest standard deviation of landmark positions are the videos labeled “slow translation,” “fast translation,” and “medium rotation”; three videos, per subject, were considered “high motion”.

- **High-Motion (MMSE-HR):** F005-T10, F006-T11, F007-T11, F008-T10, F009-T11, F010-T11, F011-T11, F012-T11, F013-T8, F014-T8, F015-T8, F016-T8, F017-T8, F018-T1, F019-T11, F020-T10, F021-T11, F022-T10, F023-T10, F024-T10, F025-T10, F026-T10, F027-T10, M001-T11, M002-T11, M003-T10, M004-T10, M005-T10, M006-T10, M007-T10, M008-T11,

M009-T10, M010-T11, M011-T8, M012-T11, M013-T11, M014-T10, M015-T10, M016-T11, M017-T10.

- **High-motion (PURE):** Videos labeled as 03-Slow Translation, 04-Fast Translation, and 06-Medium Rotation.

## REFERENCES

- [1] D. McDuff, "Camera measurement of physiological vital signs," *ACM Computer Survey*, vol. 55, no. 9, 2023. [Online]. Available: <https://doi.org/10.1145/3558518>
- [2] X. Liu, G. Narayanswamy, A. Paruchuri, X. Zhang, J. Tang, Y. Zhang, Y. Wang, S. Sengupta, S. Patel, and D. McDuff, "rppg-toolbox: Deep remote ppg toolbox," *arXiv preprint arXiv:2210.00716*, 2022.
- [3] F. Shaffer and J. P. Ginsberg, "An overview of heart rate variability metrics and norms," *Frontiers in public health*, vol. 5, p. 258, 2017.
- [4] L. A. Aarts, V. Jeanne, J. P. Cleary, C. Lieber, J. S. Nelson, S. B. Oetomo, and W. Verkruysse, "Non-contact heart rate monitoring utilizing camera photoplethysmography in the neonatal intensive care unit—a pilot study," *Early human development*, vol. 89, no. 12, 2013.
- [5] V. R. Shenoy, C. Q. Kingston, M. Singh, I. C. Fleming, N. J. Durr, R. Chellappa, and A. M. Giladi, "Perfusion assessment of healthy and injured hands using video-based deep learning models," *Plastic and Reconstructive Surgery*, 2025. [Online]. Available: [https://journals.lww.com/plasreconsurg/fulltext/9900/\\_perfusion\\_assessment\\_of\\_healthy\\_and\\_injured\\_hands.2657.aspx](https://journals.lww.com/plasreconsurg/fulltext/9900/_perfusion_assessment_of_healthy_and_injured_hands.2657.aspx)
- [6] E. M. Nowara, T. K. Marks, H. Mansour, and A. Veeraraghavan, "Near-infrared imaging photoplethysmography during driving," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 4, pp. 3589–3600, 2022.
- [7] S.-Q. Liu, X. Lan, and P. C. Yuen, "Remote photoplethysmography correspondence feature for 3d mask face presentation attack detection," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [8] 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*, vol. 18, no. 10, 2010.
- [9] M. Lewandowska, J. Rumiński, T. Kocejko, and J. Nowak, "Measuring pulse rate with a webcam — a non-contact method for evaluating cardiac activity," in *2011 Federated Conference on Computer Science and Information Systems (FedCSIS)*, 2011.
- [10] G. de Haan and V. Jeanne, "Robust pulse rate from chrominance-based rppg," *IEEE Trans. on Biomedical Eng.*, 2013.
- [11] G. de Haan and A. van Leest, "Improved motion robustness of remote-ppg by using the blood volume pulse signature," *Physiological Measurement*, vol. 35, no. 9, aug 2014. [Online]. Available: <https://dx.doi.org/10.1088/0967-3334/35/9/1913>
- [12] E. Magdalena Nowara, T. K. Marks, H. Mansour, and A. Veeraraghavan, "Sparseppg: Towards driver monitoring using camera-based vital signs estimation in near-infrared," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2018.
- [13] W. Chen and D. McDuff, "Deepphys: Video-based physiological measurement using convolutional attention networks," in *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [14] X. Liu, J. Fromm, S. Patel, and D. McDuff, "Multi-task temporal shift attention networks for on-device contactless vitals measurement," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/e1228be46de6a0234ac22ded31417bc7-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/e1228be46de6a0234ac22ded31417bc7-Paper.pdf)
- [15] Z. Yu, X. Li, and G. Zhao, "Remote photoplethysmograph signal measurement from facial videos using spatio-temporal networks," in *Proc. BMVC*, 2019.
- [16] Z. Yu, Y. Shen, J. Shi, H. Zhao, P. H. Torr, and G. Zhao, "Physformer: Facial video-based physiological measurement with temporal difference transformer," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 4186–4196.
- [17] E. M. Nowara, D. McDuff, and A. Veeraraghavan, "The benefit of distraction: Denoising camera-based physiological measurements using inverse attention," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 4955–4964.
- [18] X. Liu, Z. Jiang, J. Fromm, X. Xu, S. Patel, and D. McDuff, "Metaphys: Few-shot adaptation for non-contact physiological measurement," in *Proc. of the Conf. on Health, Inference, and Learning*, ser. CHIL '21. New York, NY, USA: Association for Computing Machinery, 2021. [Online]. Available: <https://doi.org/10.1145/3450439.3451870>
- [19] E. Lee, E. Chen, and C.-Y. Lee, "Meta-rppg: Remote heart rate estimation using a transductive meta-learner," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVII 16*. Springer, 2020.
- [20] X. Liu, B. Hill, Z. Jiang, S. Patel, and D. McDuff, "Efficientphys: Enabling simple, fast and accurate camera-based cardiac measurement," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2023, pp. 5008–5017.
- [21] X. Liu, M. Zhang, Z. Jiang, S. Patel, and D. McDuff, "Federated remote physiological measurement with imperfect data," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2022, pp. 2155–2164.
- [22] J. Gideon and S. Stent, "The way to my heart is through contrastive learning: Remote photoplethysmography from unlabelled video," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021.
- [23] Z. Sun and X. Li, "Contrast-phys+: Unsupervised and weakly-supervised video-based remote physiological measurement via spatiotemporal contrast," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 8, pp. 5835–5851, 2024.
- [24] Z. Yue, M. Shi, and S. Ding, "Facial video-based remote physiological measurement via self-supervised learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 11, 2023.
- [25] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, 2009. [Online]. Available: <https://doi.org/10.1137/080716542>
- [26] V. Monga, Y. Li, and Y. C. Eldar, "Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing," *IEEE Signal Processing Magazine*, vol. 38, no. 2, 2021.
- [27] S. Bai, J. Z. Kolter, and V. Koltun, "Deep equilibrium models," in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2019/file/01386bd6d8e091c2ab4c7c7de644d37b-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/01386bd6d8e091c2ab4c7c7de644d37b-Paper.pdf)
- [28] D. Gilton, G. Ongie, and R. Willett, "Deep equilibrium architectures for inverse problems in imaging," *IEEE Transactions on Computational Imaging*, vol. 7, 2021.
- [29] W. W. Nichols, M. O'Rourke, E. R. Edelman, and C. Vlachopoulos, *McDonald's blood flow in arteries: theoretical, experimental and clinical principles*. CRC press, 2022.
- [30] V. R. Shenoy, T. K. Marks, H. Mansour, and S. Lohit, "Unrolled ippg: Video heart rate estimation via unrolling proximal gradient descent," in *2023 IEEE International Conference on Image Processing (ICIP)*, 2023, pp. 2715–2719.
- [31] R. Favilla, V. C. Zuccala, and G. Coppini, "Heart rate and heart rate variability from single-channel video and ica integration of multiple signals," *IEEE Journal of Biomedical and Health Informatics*, vol. 23, no. 6, 2018.
- [32] W. Wang, A. C. Den Brinker, S. Stuijk, and G. De Haan, "Algorithmic principles of remote ppg," *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 7, 2016.
- [33] S. Tulyakov, X. Alameda-Pineda, E. Ricci, L. Yin, J. F. Cohn, and N. Sebe, "Self-adaptive matrix completion for heart rate estimation from face videos under realistic conditions," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [34] R. Spetlik, V. Franc, J. Cech, and J. Matas, "Visual heart rate estimation with convolutional neural network," in *British Machine Vision Conference*, 2018. [Online]. Available: <https://api.semanticscholar.org/CorpusID:52219725>
- [35] Z. Yu, X. Li, and G. Zhao, "Remote photoplethysmograph signal measurement from facial videos using spatio-temporal networks," *arXiv preprint arXiv:1905.02419*, 2019.
- [36] S.-Q. Liu and P. C. Yuen, "Robust remote photoplethysmography estimation with environmental noise disentanglement," *IEEE Transactions on Image Processing*, vol. 33, 2024.
- [37] F. Bousefsaf, A. Pruski, and C. Maaoui, "3d convolutional neural networks for remote pulse rate measurement and mapping from facial video," *Applied Sciences*, vol. 9, no. 20, 2019.

- [38] M. Hu, F. Qian, D. Guo, X. Wang, L. He, and F. Ren, "Eta-rppgnet: Effective time-domain attention network for remote heart rate measurement," *IEEE Transactions on Instrumentation and Measurement*, vol. 70, 2021.
- [39] Z. Yu, X. Li, X. Niu, J. Shi, and G. Zhao, "Autohr: A strong end-to-end baseline for remote heart rate measurement with neural searching," *IEEE Signal Processing Letters*, vol. 27, 2020.
- [40] A. Das, H. Lu, H. Han, A. Dantcheva, S. Shan, and X. Chen, "Bypnet: Video-to-byp signal prediction for remote heart rate estimation," in *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*. IEEE, 2021.
- [41] J. Du, S.-Q. Liu, B. Zhang, and P. C. Yuen, "Dual-bridging with adversarial noise generation for domain adaptive rppg estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 10 355–10 364.
- [42] J. Lin, C. Gan, and S. Han, "Tsm: Temporal shift module for efficient video understanding," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [43] X. Liu, Y. Zhang, Z. Yu, H. Lu, H. Yue, and J. Yang, "rppg-mae: Self-supervised pretraining with masked autoencoders for remote physiological measurements," *IEEE Transactions on Multimedia*, 2024.
- [44] E. Mejía-Mejía, J. M. May, R. Torres, and P. A. Kyriacou, "Pulse rate variability in cardiovascular health: A review on its applications and relationship with heart rate variability," *Physiological Measurement*, vol. 41, no. 7, p. 07TR01, 2020.
- [45] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," in *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., vol. 31. Curran Associates, Inc., 2018. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2018/file/69386f6bb1dfed68692a24c8686939b9-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2018/file/69386f6bb1dfed68692a24c8686939b9-Paper.pdf)
- [46] C. G. Broyden, "A class of methods for solving nonlinear simultaneous equations," *Mathematics of Computation*, 1965.
- [47] D. W. Peaceman and H. H. Rachford, Jr., "The numerical solution of parabolic and elliptic differential equations," *Journal of the Society for Industrial and Applied Mathematics*, vol. 3, no. 1, 1955. [Online]. Available: <https://doi.org/10.1137/0103003>
- [48] D. G. Anderson, "Iterative procedures for nonlinear integral equations," *J. ACM*, vol. 12, no. 4, oct 1965. [Online]. Available: <https://doi.org/10.1145/321296.321305>
- [49] S. G. Krantz and H. R. Parks, *The Implicit Function Theorem: History, Theory, and Applications*. Birkhäuser Boston, MA, 2012.
- [50] S. Bai, V. Koltun, and J. Z. Kolter, "Multiscale deep equilibrium models," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 5238–5250. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/3812f9a59b634c2a9c574610eba5bed-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/3812f9a59b634c2a9c574610eba5bed-Paper.pdf)
- [51] S. Gurumurthy, S. Bai, Z. Manchester, and J. Z. Kolter, "Joint inference and input optimization in equilibrium networks," in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34. Curran Associates, Inc., 2021, pp. 16 818–16 832. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2021/file/8c3c27ac7d298331a1bdf0a5e8703d3-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/8c3c27ac7d298331a1bdf0a5e8703d3-Paper.pdf)
- [52] S. Bai, Z. Geng, Y. Savani, and J. Z. Kolter, "Deep equilibrium optical flow estimation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 620–630.
- [53] A. Pokle, Z. Geng, and J. Z. Kolter, "Deep equilibrium approaches to diffusion models," in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35. Curran Associates, Inc., 2022, pp. 37 975–37 990. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/f7f47a73d631c0410cbe2748a8015241-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/f7f47a73d631c0410cbe2748a8015241-Paper-Conference.pdf)
- [54] P. Micaelli, A. Vahdat, H. Yin, J. Kautz, and P. Molchanov, "Recurrence without recurrence: Stable video landmark detection with deep equilibrium models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 22 814–22 825.
- [55] A. Gkillas, D. Ampeliotis, and K. Berberidis, "A highly interpretable deep equilibrium network for hyperspectral image deconvolution," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [56] Q. Chen, Y. Wang, Z. Geng, Y. Wang, J. Yang, and Z. Lin, "Equilibrium image denoising with implicit differentiation," *IEEE Transactions on Image Processing*, vol. 32, 2023.
- [57] J. Liu, X. Xu, W. Gan, s. shoushtari, and U. Kamilov, "Online deep equilibrium learning for regularization by denoising," in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35. Curran Associates, Inc., 2022. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/a2440e23f6a8c037eff1dc4f1156aa35-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/a2440e23f6a8c037eff1dc4f1156aa35-Paper-Conference.pdf)
- [58] J. Hu, S. Shoushtari, Z. Zou, J. Liu, Z. Sun, and U. S. Kamilov, "Robustness of deep equilibrium architectures to changes in the measurement model," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [59] Z. Zou, J. Liu, B. Wohlberg, and U. S. Kamilov, "Deep equilibrium learning of explicit regularizers for imaging inverse problems," *IEEE Open J. Signal Process*, 2023, doi:10.1109/OJSP.2023.3296036.
- [60] Z. Geng, X.-Y. Zhang, S. Bai, Y. Wang, and Z. Lin, "On training implicit models," in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34. Curran Associates, Inc., 2021, pp. 24 247–24 260. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2021/file/cb8da6767461f2812ae4290eac7cbe42-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/cb8da6767461f2812ae4290eac7cbe42-Paper.pdf)
- [61] S. W. Fung, H. Heaton, Q. Li, D. Mckenzie, S. Osher, and W. Yin, "Jfb: Jacobian-free backpropagation for implicit networks," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 6, pp. 6648–6656, Jun. 2022. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/20619>
- [62] S. Bai, V. Koltun, and Z. Kolter, "Stabilizing equilibrium models by jacobian regularization," in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 554–565. [Online]. Available: <https://proceedings.mlr.press/v139/bai21b.html>
- [63] C. Wei and J. Z. Kolter, "Certified robustness for deep equilibrium models via interval bound propagation," in *International Conference on Learning Representations*, 2022.
- [64] M. Li, Y. Wang, and Z. Lin, "CerDEQ: Certifiable deep equilibrium model," in *Proceedings of the 39th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, Eds., vol. 162. PMLR, 17–23 Jul 2022, pp. 12 998–13 013. [Online]. Available: <https://proceedings.mlr.press/v162/li22t.html>
- [65] y. yang, J. Sun, H. Li, and Z. Xu, "Deep admm-net for compressive sensing mri," in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2016/file/1679091c5a880faf6fb5e6087eb1b2dc-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2016/file/1679091c5a880faf6fb5e6087eb1b2dc-Paper.pdf)
- [66] J. Zhang and B. Ghanem, "Ista-net: Interpretable optimization-inspired deep network for image compressive sensing," in *CVPR*, 2018, pp. 1828–1837.
- [67] S. Lohit, D. Liu, H. Mansour, and P. T. Boufounos, "Unrolled projected gradient descent for multi-spectral image fusion," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 7725–7729.
- [68] Y. Li, Y. Zou, L. Chen, and Y. Fu, "Physical degradation model-guided interferometric hyperspectral reconstruction with unfolding transformer," 2025. [Online]. Available: <https://arxiv.org/abs/2506.21880>
- [69] N. Parikh and S. Boyd, "Proximal algorithms," *Found. Trends Optim.*, vol. 1, no. 3, jan 2014. [Online]. Available: <https://doi.org/10.1561/2400000003>
- [70] M. Hutchinson, "A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines," *Communications in Statistics - Simulation and Computation*, vol. 18, no. 3, 1989. [Online]. Available: <https://doi.org/10.1080/03610918908812806>
- [71] X. Niu, H. Han, S. Shan, and X. Chen, "Synrhythm: Learning a deep heart rate estimator from general to specific," in *2018 24th International Conference on Pattern Recognition (ICPR)*, 2018.
- [72] X. Niu, Z. Yu, H. Han, X. Li, S. Shan, and G. Zhao, "Video-based remote physiological measurement via cross-verified feature disentangling," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*. Springer, 2020.
- [73] R. Song, H. Chen, J. Cheng, C. Li, Y. Liu, and X. Chen, "PulseGAN: Learning to generate realistic pulse waveforms in remote photoplethysmography," *IEEE Journal of Biomedical and Health Informatics*, vol. 25, no. 5, 2021.
- [74] H. Lu, H. Han, and S. K. Zhou, "Dual-gan: Joint bvp and noise modeling for remote physiological measurement," in *Proceedings of the IEEE/CVF*

*Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 12404–12413.

- [75] Z. Zhang, J. M. Girard, Y. Wu, X. Zhang, P. Liu, U. Ciftci, S. Canavan, M. Reale, A. Horowitz, H. Yang, J. F. Cohn, Q. Ji, and L. Yin, “Multimodal spontaneous emotion corpus for human behavior analysis,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 3438–3446.
- [76] I. O. Ertugrul, J. F. Cohn, L. A. Jeni, Z. Zhang, L. Yin, and Q. Ji, “Cross-domain au detection: Domains, learning approaches, and measures,” in *2019 14th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2019)*, 2019, pp. 1–8.
- [77] A. Comas, T. K. Marks, H. Mansour, S. Lohit, Y. Ma, and X. Liu, “Turnip: Time-series u-net with recurrence for nir imaging ppg,” in *2021 IEEE International Conference on Image Processing (ICIP)*, 2021, pp. 309–313.
- [78] R. Stricker, S. Müller, and H.-M. Groß, “Non-contact video-based pulse rate measurement on a mobile service robot,” *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, 2014. [Online]. Available: <https://api.semanticscholar.org/CorpusID:8529212>
- [79] S. Bobbia, R. Macwan, Y. Benezeth, A. Mansouri, and J. Dubois, “Un-supervised skin tissue segmentation for remote photoplethysmography,” *Pattern Recognition Letters*, vol. 124, 2019, award Winning Papers from the 23rd Inter. Conf. on Pattern Recognition (ICPR). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167865517303860>
- [80] S. Zhang, X. Zhu, Z. Lei, H. Shi, X. Wang, and S. Z. Li, “Faceboxes: A cpu real-time face detector with high accuracy,” 2017. [Online]. Available: <https://arxiv.org/abs/1708.05234>
- [81] A. Kumar, T. K. Marks, W. Mou, Y. Wang, M. Jones, A. Cherian, T. Koike-Akino, X. Liu, and C. Feng, “Luvli face alignment: Estimating landmarks’ location, uncertainty, and visibility likelihood,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [82] Y. Zhang, H. Lu, X. Liu, Y. Chen, and K. Wu, “Advancing generalizable remote physiological measurement through the integration of explicit and implicit prior knowledge,” *IEEE Transactions on Image Processing*, vol. 34, pp. 3764–3778, 2025.
- [83] H. Lu, Z. Yu, X. Niu, and Y. Chen, “Neuron structure modeling for generalizable remote physiological measurement,” in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 18 589–18 599.
- [84] J. Du, S.-Q. Liu, B. Zhang, and P. C. Yuen, “Dual-bridging with adversarial noise generation for domain adaptive rppg estimation,” in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 10 355–10 364.
- [85] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-cam: Visual explanations from deep networks via gradient-based localization,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 618–626.
- [86] Z. Fang, S. Buchanan, and J. Sulam, “What’s in a prior? learned proximal networks for inverse problems,” in *The Twelfth International Conference on Learning Representations*, 2024.



**Vineet R. Shenoy** is a PhD student in the Department of Electrical and Computer Engineering at Johns Hopkins University. He is broadly interested in Computer Vision and Machine Learning, with a current focus on signal recovery algorithms. Previously, he earned his B.S. in Electrical and Computer Engineering and Computer Science at Rutgers University New-Brunswick, where he received the James Leroy Potter Award for Original Investigation as well as the Chancellor’s Leadership Award. He is a member of the Phi Beta Kappa and Tau Beta Pi honor societies.



**Suhās Lohit** (S’15, M’20) received the M.S degree in Computer Engineering in 2015 and the Ph.D. degree in Electrical Engineering in 2019 from Arizona State University. He is currently a Principal Research Scientist at Mitsubishi Electric Research Laboratories in Cambridge, MA, USA, where he has worked since 2020. His research interests lie in machine learning, computer vision and computational imaging, and artificial intelligence.



**Hassan Mansour** (S’99, M’09, SM’17) received the B.E. degree in computer and communications engineering from the American University of Beirut, Beirut, Lebanon, in 2003, and the M.A.Sc. degree in electrical and computer engineering and the Ph.D. degree in electrical and computer engineering from The University of British Columbia, Vancouver, BC, Canada, in 2005 and 2009, respectively. Between January 2010 and January 2013, he was a Postdoctoral Research Fellow with the Department of Computer Science, the Mathematics Department, and the Department of Earth, Ocean, and Atmospheric Sciences, The University of British Columbia. He is currently a Senior Principal Research Scientist with Mitsubishi Electric Research Laboratories, Cambridge, MA, USA. His research interests are in inverse problems, compressed sensing, sparse signal reconstruction, image enhancement, and scalable video compression and transmission. His current research is focused on the design of efficient acquisition schemes and reconstruction algorithms for natural images, radar sensing, video analytics, and inverse scattering problems. Dr. Mansour is a member of the IEEE Computational Imaging Technical Committee and the IEEE Sensor Array and Multichannel Technical Committee. He is also an Associate Editor for the IEEE Transactions on Signal Processing.



**Rama Chellappa** (Fellow, IEEE) is a Bloomberg Distinguished Professor with the Department of Electrical and Computer Engineering (Whiting School of Engineering) and Biomedical Engineering (School of Medicine) with a secondary appointment with the Department of Computer Science, Johns Hopkins University (JHU). At JHU, he is affiliated with the Center for Imaging Sciences, the Center for Language and Speech Processing, the Institute for Assured Autonomy, the Data Science and Artificial Intelligence Institute, and the Mathematical Institute for Data Science. He holds the position as College Park Professor with the University of Maryland. He holds nine patents. His current research interests are computer vision, pattern recognition, machine intelligence, and artificial intelligence. He has received many prestigious awards from the IEEE, the University of Maryland, the University of Southern California, and the International Association of Pattern Recognition. He has been recognized as an Outstanding ECE Alumnus by Purdue University and as a Distinguished Alumnus by the Indian Institute of Science, India. He is a Fellow of Optica, AAAI, AAAS, ACM, AIMBE, IAPR, NAI, and OSA, a member of the National Academy of Engineering and a Foreign Fellow of the Indian National Academy of Engineering.



**Tim K. Marks** (Member, IEEE) received the AB degree in Physics from Harvard University in 1991, the MS degree in Cognitive Science from the University of California, San Diego in 2001, and the PhD degree in Cognitive Science from the University of California, San Diego (UCSD) in 2006. Dr. Marks is a Senior Principal Research Scientist and the Leader of the Computer Vision Team at the Mitsubishi Electric Research Laboratories (MERL), in Cambridge, MA, where he has worked since 2008. From 2006–2008, he was a postdoctoral researcher in robotics at the UCSD Department of Computer Science, in collaboration with the computer vision group at the NASA/Caltech Jet Propulsion Laboratory. Before attending graduate school, he was a high school math and physics teacher and then a mathematics textbook editor for Houghton Mifflin. His current research interests span many topics in computer vision and machine learning. He has published over 60 papers in journals and conferences (including the IEEE Transactions on Pattern Analysis and Machine Intelligence, CVPR, ICCV, ECCV, IROS, and NeurIPS), as well as a college Number Theory textbook. He has been granted over 20 US patents.