MetaPhys: Few-Shot Adaptation for Non-Contact Physiological Measurement

Xin Liu
University of Washington
Seattle, WA, USA
xliu0@cs.washington.edu

Ziheng Jiang
University of Washington & OctoML
Seattle, WA, USA
ziheng@cs.washington.edu

Josh Fromm
OctoML
Seattle, WA, USA
jwfromm@octoml.ai

Xuhai Xu
University of Washington
Seattle, WA, USA
xuhaixu@cs.washington.edu

Shwetak Patel
University of Washington
Seattle, WA, USA
shwetak@cs.washington.edu

Daniel McDuff
Microsoft Research AI
Redmond, WA, USA
damcduff@microsoft.com

ABSTRACT

There are large individual differences in physiological processes, making designing personalized health sensing algorithms challenging. Existing machine learning systems struggle to generalize well to unseen subjects or contexts and can often contain problematic biases. Video-based physiological measurement is no exception. Therefore, learning personalized or customized models from a small number of unlabeled samples is very attractive as it would allow fast calibrations to improve generalization and help correct biases. In this paper, we present a novel meta-learning approach called MetaPhys for personalized video-based cardiac measurement for non-contact pulse and heart rate monitoring. Our method uses only 18-seconds of video for customization and works effectively in both supervised and unsupervised manners. We evaluate our approach on two benchmark datasets and demonstrate superior performance in cross-dataset evaluation with substantial reductions (42% to 44%) in errors compared with state-of-the-art approaches. We also find that our method leads to large reductions in bias due to skin type.

CCS CONCEPTS

• Computing methodologies → Computer vision.

KEYWORDS

remote physiological sensing, rPPG, mobile health, healthcare, computer vision, meta learning, few-shot learning

ACM Reference Format:

Xin Liu, Ziheng Jiang, Josh Fromm, Xuhai Xu, Shwetak Patel, and Daniel McDuff. 2021. MetaPhys: Few-Shot Adaptation for Non-Contact Physiological Measurement. In *ACM Conference on Health, Inference, and Learning (ACM CHIL '21)*, April 8–10, 2021, Virtual Event, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3450439.3451870>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM CHIL '21, April 8–10, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8359-2/21/04...\$15.00

<https://doi.org/10.1145/3450439.3451870>

1 INTRODUCTION

The importance of scalable health sensing has been acutely highlighted during the SARS-CoV-2 (COVID-19) pandemic. The virus has been linked to increased risk of myocarditis and other serious cardiac (heart) conditions [35]. Contact sensors (electrocardiograms, oximeters) are the current gold-standard for measurement of heart function. However, these devices are still not ubiquitously available, especially in low-resource settings. The development of video-based non-contact sensing of vital signs presents an opportunity for highly scalable physiological monitoring. Furthermore, in clinical settings non-contact sensing could reduce the risk of infection for vulnerable patients (e.g., infants and elderly) and the discomfort caused by obtrusive leads and electrodes [40].

While there are compelling advantages of camera-based sensing, the approach also presents unsolved challenges. The use of ambient illumination means camera-based measurement is sensitive to **environmental differences** in the intensity and composition of the incident light. Camera **sensor differences** mean that hardware can differ in sensitivity across the frequency spectrum and auto adjustments (e.g., white balancing) and video compression codecs can further impact pixel values [29]. People (the subjects) exhibit large **individual differences** in appearance (e.g., skin type, facial hair) and physiology (e.g, pulse dynamics). Finally, **contextual differences** mean that motions in a video at test time might be different from those seen in the training data. One specific example of challenges to generalization is biases in performance across skin types [30]. This problem is not isolated to physiological measurement as studies have found systematic biases in facial gender classification, with error rates up to 7x higher on women than men and poorer performance on people with darker skin types [6].

Collecting large corpora of high-quality physiological data presents challenges: 1) recruiting and instrumenting participants is often expensive and requires advanced technical expertise, 2) the data can reveal the identity of the subjects and/or sensitive health information meaning it is difficult for researchers to share such datasets. Therefore, training supervised models that generalize well across environments and subjects is often difficult. For these reasons, we observe that performance on cross-dataset evaluation is significantly worse than within-dataset evaluation using current state-of-the-art methods [7, 24].

Calibration of consumer health sensors is often performed in a clinic, where a doctor or nurse will collect readings from a high-end sensor to calibrate a consumer-level device the patient owns. The reason for this is partly due to the variability within readings from consumer devices across different individuals. Ideally, we would be able to train a personalized model for each individual; however, standard supervised learning training schemes require large amounts of labeled data. Getting enough physiological training data of each individual is difficult because it requires using medical-grade devices to provide reliable labels. Being able to generate a personalized model from a small amount of training samples would enable customization based on a few seconds or minutes of video captured while visiting a clinic where people have access to a gold-standard device. Furthermore, if this process could be achieved without the need for these devices (i.e., in an unsupervised manner), that would have even greater impact. Integrating remote physiological measurement into telehealth systems could provide patients' vital signs for clinicians during remote diagnosis. Given that requests for telehealth appointments have increased more than 10x during COVID-19, and that this is expected to continue into the future [36], robust personalized models are of growing importance.

Meta-learning, or learning to learn, has been extensively studied in the past few years [18]. Instead of learning a specific generalized mapping, the goal of meta-learning is to design a model that can adapt to a new task or context with a small amount of data. Due to the inherent ability for fast adaption, meta-learning is a good candidate strategy for building personalized models (e.g., personalization in dialogue and video retargeting [21, 25].) However, we argue that meta-learning is underused in healthcare. The goal of this work is to develop a meta-learning based personalization framework in remote physiological measurement whereby we can use a limited amount of data from a previously unseen individual (task) to mimic how a clinician might manually calibrate sensor readings for a specific patient. When meta-learning is applied to remote physiological measurement, there are two kinds of scenarios: 1) supervised adaptation with a few samples of labeled data from a clinical grade sensor and 2) unsupervised adaptation with unlabeled data. We hypothesize that supervised adaptation is more likely to yield a robust personalized model with only a few labels, while unsupervised adaptation may personalize the model less effectively but require much lower effort and complexity in practice.

In this paper, we propose a novel meta-learning approach to address the aforementioned challenges called MetaPhys. Our contributions are: 1) A meta-learning based deep neural framework, supporting both **supervised and unsupervised** few-shot adaptation, for camera-based vital sign measurement; 2) A systematic cross-dataset evaluation showing that our system considerably outperforms the state-of-the-art (42% to 52% reduction in heart rate error); 3) To perform an ablation experiment, freezing weights in the temporal and appearance branches to test sensitivity during adaptation; 4) To analyze performance for subjects with different skin types. To our best knowledge, MetaPhys is the first work that leverages pseudo labels in training a physiological sensing model

and the first unsupervised deep learning method in remote physiological measurement. Our code, example models, and video results can be found on our project page.¹

2 RELATED WORK

2.1 Video-Based Physiological Measurement

Video-based physiological measurement is a growing interdisciplinary domain that leverages ubiquitous imaging devices (e.g., webcams, smartphones' cameras) to measure vital signs and other physiological processes. Early work established that changes in light reflected from the body could be used to capture subtle variations blood volume and motion related to the photoplethysmogram (PPG) [38, 39] and ballistocardiogram (BCG) [2], respectively. Video analysis enables non-contact, spatial and temporal measurement of arterial and peripheral pulsations and allows for magnification of these signals [8, 43], which may help with examination (e.g., [1]). Based on the PPG and BCG signal, heart rate can be extracted [2, 34]. Subsequent research has shown HRV and waveform morphology can also be obtained from facial imaging data and achieved high-precision in the measurement of inter-beat-interval [27].

However, the relationship between pixels and underlying physiological changes in a video is complex and neural models have shown strong performance compared to source separation techniques [7, 24, 44, 45]. Conventional supervised learning requires a large amount of training data to produce a generalized model. However, obtaining synchronized physiological and facial data is complicated and expensive. Current public datasets have a limited number of subjects and diversity in regards of appearance (including skin type), camera sensors, environmental conditions and subject motions. Therefore, if the subject of interest is not in the training data or the video is otherwise different, performance can be considerably degraded, a result that is not acceptable for a physiological sensor.

Lee et al. [20] recognized the potential for meta-learning applied to imaging-based cardiac pulse measurement. Their method (Meta-rPPG) focuses on using transductive inference based meta-learning and a LSTM encoder-decoder architecture which to our knowledge was not validated in previous work. Instead, our proposed meta-learning framework is built on top of a state-of-the-art on-device network [24] and aims to explore the potential of both supervised and unsupervised on-device personalized meta-learning. More specifically, Meta-rPPG uses a synthetic gradient generator and a prototypical distance minimizer to perform transductive inference to enable self-supervised meta-learning. This learning mechanism requires a number of rather complex steps with transductive inference. We propose a somewhat simpler mechanism that is physiologically and optically grounded [24, 42] and achieves greater accuracy.

2.2 Meta-Learning and Person Specific Models

The ability to learn from a small number of samples or observations is often used as an example of the unique capabilities of human intelligence. However, machine learning systems are often brittle

¹<https://github.com/xliucs/MetaPhys>

in a similar context. Meta-learning approaches tackle this problem by creating a general learner that is able to adapt to a new task with a small number of training samples, inspired by how humans can often master a new skill without many observations [18]. However, most of the previous work in meta-learning focuses on supervised vision problems [37, 47] and in the computer vision literature has mainly been applied to image analysis [22, 41]. Supervised regression in video settings has received less attention. One of few examples is object or face tracking [9, 31]. In these tasks, the learner needs to adapt to the individual differences in appearance of the target and then track it across frames, even if the appearance changes considerably over time in the video. Choi et al. [9] present a matching network architecture providing the meta-learner with information in the form of loss gradients obtained using the training samples.

The property of fast adaptation makes meta-learning a good candidate for personalizing models, it has been used in various applications such as dialogue agents [25], gaze estimation [17], sleep stage classification [3], activity recognition [15], and video retargeting [21]. For example, Banluesombatkul et al. proposed a MAML-based meta-learning system to perform fast adaptation of a sleep stage classification model using biosignals [3]. More recently, MetaPix [21] leveraged a meta-learning training schema with a small amount of video to adapt a universal generator to a particular background and human in the problem of video retargeting. Similarly, our proposed meta-learning framework is also capable of personalizing a universal remote physiological model to a new person or an environmental setting.

3 METHOD

3.1 Physiological Meta-Learning

In camera-based cardiac measurement, the goal is to separate pixel changes due to pulsatile variations in blood volume and body motions from other changes that are not related to the pulse signal. Examples of “noise” in this context that might impact the performance on the task include: changes in the environment (illumination) and changes in appearance of the subject and motions (e.g., facial expressions, rigid head motions). A model trained within a traditional supervised learning regime might perform well if illumination, non-pulsatile motions, and facial appearances in the test set are similar to those in the training set. However, empirical evidence shows that performance usually significantly degrades from one dataset to another, suggesting that traditional training is likely to overfit to the training set to some extent [7]. Therefore, to achieve state-of-the-art performance in remote physiological measurement on cross-dataset evaluation, the system should have: 1) a good initial representation of the mapping from the raw video data to the pulse signal, and 2) a strategy for adapting to unseen individuals and environments.

To achieve this, we propose a system called MetaPhys (Fig. 1), an adaptable meta-learning based on-device framework aimed at efficient and personalized remote physiological sensing. MetaPhys uses a pretrained convolutional attention network as the backbone (described below) and leverages a novel personalized meta-learning schema to overcome the aforementioned limitations. We adopt Model-Agnostic Meta-Learning (MAML) [13] as our personalized

parameter update schema. MAML produces a general initialization as the starting point for fast adaptation to a diverse set of unseen tasks with only a few training samples. However, applying MAML to the task of camera-based physiological measurement has differences to many previously explored meta-learning problems. Existing meta-learning approaches are often evaluated on classification or some toy regression tasks due to the lack of regression benchmark datasets [18]. Our problem is a non-trivial vision-based regression task due to the subtle nature of the underlying physiological signal. Algorithm 1 outlines the training process for MetaPhys, we first pretrain the backbone network to get an initial spatial-temporal representation. Then we treat each individual as a task τ_i . At training time, we split the data into a support set (K video frames) and a query set (K' video frames) for each individual (task). The support set is used to update the task’s parameters and yield a personalized model θ_i . The query set is used to assess the effectiveness of the personalized model and further update the global initialization θ to make future adaptation better. A robust personalized model θ_i aims to provide a more accurate attention mask to the corresponding motion branch and to preform precise physiological measurement for the target individual as well as the target’s environment. During the testing stage, MetaPhys has the updated global initialization $\hat{\theta}$, and can generate $\hat{\theta}_i$ for each test individual (task) by optimizing the test support set as $\hat{\theta}_{\tau_i} \leftarrow \hat{\theta} - \alpha \nabla_{\hat{\theta}} \mathcal{L}_{\tau_i} f(\hat{\theta})$. With this training and testing schema, the robust global initialization $\hat{\theta}$ generated from MetaPhys not only leverages the pretrained representation but also learns how to adapt to new individuals and environmental noise quickly.

3.2 Spatial and Temporal Model Architecture Backbone

Our ultimate goal is a computationally efficient on-device meta-learning framework that offers inference at 150 fps. Therefore, we adopt the state-of-the-art architecture (TS-CAN) [24] for remote cardiopulmonary monitoring. TS-CAN is an end-to-end neural architecture with appearance and motion branches. The input is a series of video frames and the output is the first-derivative of the pulse estimate at the corresponding time points. Tensor shifting modules (TSM) [23] are used that shift frames along the temporal axis allowing for information exchange across time. This helps capture temporal dependencies beyond consecutive frames. The appearance branch and attention mechanism help guide the motion branch to focus on regions with high pulsatile signal (e.g., skin) instead of others (e.g., clothes, hair) (see Fig. 1). However, we discover empirically that this network does not necessarily generalize well across datasets with differences in subjects, lighting, backgrounds and motions (see Table 1). One of the main challenges when employing TS-CAN is that the appearance branch may not generate an accurate mask while testing on unseen subjects or environments because of the differences in appearance of skin pixels. Without a good attention mask, motions from other sources are likely to be given more weight, thus damaging the quality of our physiological estimate.

Figure 1: We present MetaPhys, an approach for few-shot unsupervised adaptation for personalized camera-based physiological measurement models.

3.3 Supervised or Unsupervised Learning

We explore both supervised and unsupervised training regimes for MetaPhys. Supervised personalization may be suitable in clinical settings that require highly precise adaptation and where there is access to reference devices. Unsupervised personalization may be preferable for consumer measurement when convenience and scalability is of a greater priority and calibration with a clinical grade device might be difficult.

For the supervised version of MetaPhys we use the gold standard reference signal from a fingertip PPG or blood pressure wave (BPW) to train the meta-learner and perform few-shot adaptation when testing. In contrast to the supervised version, in the unsupervised case we use pseudo labels during the training of the MetaPhys meta-learner and parameter updates rather than the ground-truth signal from the medical-grade devices. We use a physiologically-based unsupervised remote physiological measurement model to generate pseudo pulse signal estimates without relying on gold standard measurements. More specifically, we leverage the Plane-Orthogonal-to-Skin (POS) [42] method, which is the current state-of-the-art for demixing in this context. POS calculates a projection plane orthogonal to the skin-tone, derived based on optical and physiological principles, that is then used for pulse extraction. The POS method can be summarized as follows: 1) spatially averaging pixel values for each frame with the region-of-interest, 2) temporally normalizing the resulting signals within a certain window size calculated relative to the frame rate, 3) applying a fixed matrix projection to offset the specular reflections and other noise, 4) band-pass filtering of the resulting pulse waveform.

We observe that even though our unsupervised model uses the POS signal for meta-training, MetaPhys’s performance significantly outperforms POS once trained. As Algorithm 1 illustrates, the pseudo label generator G produces pseudo labels for both K support frames and K' query frames for adaptation and parameter updates. We used pseudo labels for the query set (K') at training time, as we observed similar empirical results in preliminary testing whether we used pseudo labels or ground-truth labels.

4 EXPERIMENTS

4.1 Datasets

AFRL [11]: 300 videos of 25 participants (17 males) were recorded at 658x492 resolution and 30 fps. Pulse measurements were recorded via a contact reflectance PPG sensor and used for training. Electrocardiograms (ECG) were recorded for evaluating performance. Each participant was recorded six times with increasing head motion in each task (10 degrees/second, 20 degrees/second, 30 degrees/second). The participants were asked to sit still for the first two tasks and perform three motion tasks rotating their head about the vertical axis. In last set of tasks, participants were asked to orient their head randomly once every second to one of nine pre-defined locations. The six recordings were repeated twice in front of two backgrounds.

UBFC [4]: 42 videos of 42 participants were recorded at 640x480 resolution and 30 fps in uncompressed 8-bit RGB format. A CMS50E transmissive pulse oximeter was used to obtain the ground truth

Table 1: Pulse Measurement (Heart Rate) on the MMSE-HR and UBFC datasets.

Method	Train / Test Datasets	MAE	RMSE	SNR	ρ
Pretrain + Unsupervised MetaPhys	(AFRL & UBFC) / MMSE	1.87	3.12	5.04	0.89
Pretrain + Supervised MetaPhys	(AFRL & UBFC) / MMSE	2.98	4.86	3.81	0.72
MetaPhys (No pretrain)	(AFRL & UBFC) / MMSE	3.67	5.50	2.41	0.70
Supervised Pretrain + FT on Test Support Set	(AFRL & UBFC) / MMSE	4.05	5.68	2.76	0.76
Supervised Pretrain [24]	(AFRL & UBFC) / MMSE	3.78	5.75	2.67	0.77
(Unsupervised) CHROM [10]	None / MMSE	3.2	5.71	5.42	0.75
(Unsupervised) POS [42]	None / MMSE	3.98	6.66	5.74	0.67
(Unsupervised) ICA [33]	None / MMSE	4.12	6.46	6.09	0.67

Method	Train / Test Datasets	MAE	RMSE	SNR	ρ
Pretrain + Unsupervised MetaPhys	(AFRL & MMSE) / UBFC	2.46	3.12	4.28	0.96
Pretrain + Supervised MetaPhys	(AFRL & MMSE) / UBFC	1.90	2.62	3.84	0.96
MetaPhys (No pretrain)	(AFRL & MMSE) / UBFC	3.80	5.32	0.13	0.84
Supervised Pretrain + FT on Test Support Set	(AFRL & MMSE) / UBFC	6.26	7.37	-0.23	0.72
(Unsupervised) Meta-rPPG [20]	Self-Collected / UBFC	5.97	7.42	-	0.53
Supervised Pretrain [24]	(AFRL & MMSE) / UBFC	4.42	6.13	1.87	0.79
(Unsupervised) POS [42]	None / UBFC	6.44	9.48	0.55	0.66
(Unsupervised) CHROM [10]	None / UBFC	7.31	9.85	0.93	0.57
(Unsupervised) ICA [33]	None / UBFC	10.2	14.4	-0.19	0.50

MAE = Mean Absolute Error, RMSE = Root Mean Squared Error, ρ = Pearson Correlation, SNR = BVP Signal-to-Noise Ratio.

Algorithm 1 MetaPhys: Meta-learning for physiological signal personalization

Require: S : Subject-wise video data
Require: A batch of personalized tasks τ where each task τ_i contains N data points from S_i
Require: A pseudo label generator G for unsupervised meta-learning

- 1: $\theta \leftarrow$ **Pre-training** TS-CAN on AFRL dataset
- 2: **for each** $\tau_i \in \tau$ **do**
- 3: **if** Supervised **then**
- 4: $K \leftarrow$ Sample K support frames from videos of τ_i with ground truth labels
- 5: $K' \leftarrow$ Sample K' query frames from videos of τ_i with ground truth labels
- 6: **else**
- 7: $K \leftarrow$ Sample K support frames from videos of τ_i with pseudo labels from G
- 8: $K' \leftarrow$ Sample K' query frames from videos of τ_i with pseudo labels from G
- 9: **end if**
- 10: $\theta_{\tau_i} \leftarrow \theta - \alpha \nabla_{\theta} \mathcal{L}_{\tau_i} f(K, \theta)$, Update the personalized params. based on indiv. support loss
- 11: **end for**
- 12: $\hat{\theta} \leftarrow \theta - \beta \nabla_{\theta} \sum_{\tau_i} \mathcal{L}_{\tau_i} f(K'_{\tau_i}, \theta_{\tau_i})$, Update the global params. based on individuals' query loss

PPG data. All the experiments were conducted indoors with different sunlight and indoor illumination. Participants were also asked to play time sensitive mathematical games to augment the heart rate during the data collection.

MMSE [46]: 102 videos of 40 participants were recorded at 1040x1392 resolution and 25 fps. A blood pressure wave signal was measured at 1000 fps as the gold standard. The blood pressure wave was used as the training signal for this data as a PPG signal was not available. The distribution of skin types based on the Fitzpatrick scale [14] is: II=8, III=11, IV=17, V+VI=4.

4.2 Implementation Details

MetaPhys was implemented in PyTorch [32], and all the experiments were conducted on a Nvidia 2080Ti GPU. We first implemented the backbone network (TS-CAN) and modified it to use a window size of 20 frames (rather than 10) because we empirically observed a larger window size led to better overall performance. Then, we implemented MetaPhys based on a gradient computation framework called higher [16]. Compared with most previous meta-learning studies that were trained and evaluated on a single dataset (e.g., miniimagenet [41]), we used three datasets to perform pre-training and cross-dataset training and evaluation. Our backbone was pretrained on the AFRL dataset, and the training (described in Algorithm 1) and evaluation of our meta-learner were performed with the UBFC and MMSE datasets. We picked the size of the support set (K) for personalization to be 540 video frames for each individual. For a 30 fps video recording this equates to an 18-second recording which is a reasonably short calibration period. During the meta training and adaptation, we used an Adam optimizer [19] with an outer learning rate (β) of 0.001 and a stochastic gradient descent (SGD) optimizer with an inner learning rate (θ) of 0.005. We trained the meta-learner for 10 epochs, and performed one step adaptation (i.e., gradient descent).

As baselines, we implemented traditional supervised training (TS-CAN) on AFRL and evaluated on MMSE and UBFC. Fine tuning with the support set and testing on the query set was implemented as our adaptation baseline. To assure a fair comparison across all experiments, we forced the test data (test query set) to remain the same within each task. We also implemented three established unsupervised algorithms (CHROM, POS, ICA) using iPhys-Toolbox [26]. We applied post-processing to the outputs of all the methods in the same way. We first divided the remainder of the recordings for each participant into 360-frame windows (approximately 12 seconds), with no overlap, and applied a 2nd-order butterworth filter with a cutoff frequency of 0.75 and 2.5 Hz (these represent a realistic range of heart rates we would expect for adults). We then computed four metrics for each window: mean absolute error (MAE), root mean squared error (RMSE), signal-to-noise ratio (SNR) and correlation (ρ) in heart-rate estimations. Unlike most prior

work which evaluated performance on whole videos (often 30 or 60 seconds worth of data), we perform evaluation on 12 second sequences which is considerably more challenging as the model has much less information for inference.

5 RESULTS AND DISCUSSION

5.1 Comparison with the State-of-the-Art:

In this section, we compare the performance of MetaPhys with other state-of-the-art approaches using Mean Absolute Error(MAE) of heart rate. Details of the evaluation metrics are in the Appendix section. For the MMSE dataset, our proposed supervised and unsupervised MetaPhys with pretraining outperformed the state-of-the-art results by 7% and 42% in MAE, respectively (see Table 1). On the UBFC dataset, supervised and unsupervised MetaPhys with pretraining showed even greater benefits reducing error by 57% and 44% compared to the previous state-of-the-art, respectively. Meta-learning alone is not as effective as meta-learning using weights initialized in a pretraining stage (19% and 50% improvements in MMSE and UBFC). We also compared our methods against the only other meta-learning based method (Meta-rPPG) where we reduced the MAE by 68%. Furthermore, we compared MetaPhys against a traditional personalization method (fine-tuning), and our approach gained 54% and a 61% improvements in terms of MAE on MMSE and UBFC, respectively. We also evaluated performance using different time windows in the support set (6s, 12s and 18s), and the results showed training with 18s (RMSE: 3.12) outperformed 6s (RMSE: 5.43) and 12s (RMSE: 5.53) on the MMSE dataset. A similar trend was also observed on the UBFC dataset (RMSE of 18s: 3.12, RMSE of 12s: 4.48, RMSE of 6s: 3.46). Fig. 3 compares the heart rate estimates from MetaPhys and the-state-of-art algorithm (TS-CAN [24]) (y-axes) to the gold-standard measurements (x-axes). The strong correlation coefficients and bias/limits in the Bland-Altman plots presented support that MetaPhys can produce more accurate and reliable estimates. Our results from unsupervised MetaPhys indicate that pseudo labels provided by a relatively simple de-mixing approach (POS) can be used in a meta-learning context to obtain strong results. The training with these pseudo labels, sometimes with noisy labels, well outperforms the de-mixing approach itself.

5.2 Unsupervised vs. Supervised Adaptation

Next, we examine the difference between a supervised and unsupervised training regime in MetaPhys. For UBFC, the *supervised* model (MAE=1.90 BPM), outperformed the *unsupervised* model (MAE=2.46 BPM). Whereas, for the MMSE dataset the *unsupervised* model (MAE=1.87 BPM) outperformed the *supervised* model (MAE=2.98 BPM). The fact that the unsupervised model achieves broadly comparable results to the supervised model is surprising and encouraging because there are many applications where unsupervised adaptation would be more convenient and efficient (e.g., calibrating a heart rate measurement app on a smartphone without needing a reference device). We also observe that the unsupervised model, even though it used the POS signal as training input, significantly outperforms POS on both datasets, suggesting MetaPhys is able to form a better representation.

5.3 Visualizing Adaption

To help us understand why MetaPhys outperforms the state-of-the-art models we visualized the attention masks for different subjects. In Fig. 4-A, we compare the attention masks from the appearance branch of TS-CAN based on four training schemes which are: 1) supervised training with TS-CAN, 2) pretraining TS-CAN on AFRL and then fine tuning TS-CAN on the support set used for the meta-learning experiments, 3) pretraining on AFRL and supervised MetaPhys training, 4) pretraining on AFRL and unsupervised MetaPhys training. The differences are subtle, but on inspection we can notice that MetaPhys leads to masks that put higher weight on regions with greater pulsatile signal (e.g., forehead and cheeks) and less weight on less important regions (e.g., clothing - see P5 as an example). In Fig. 4-B, we visualize the progression of learning for the four different methods. Again the changes during learning are subtle, but the traditional supervised methods seem more likely to overfit even over a relatively small number of epochs meaning that the attention to important regions of the face is not as high as with the meta-learning approach, presumably because the traditional supervised learning has to capture a more generic model which is not well adapted to any one specific individual.

5.4 Freezing Appearance vs. Motion

We questioned whether the adaptation of the appearance mask was the main, or sole reason for the improvements provided by MetaPhys. To test this, we froze the weights in the motion branch of the TS-CAN during the meta-training stage and only updated weights in the appearance branch. From the results of these experiments, we observe that there is a 20% increase in MAE, indicating that MetaPhys not only noticeably improves the quality of the attention mask, but also learns additional temporal dynamics specific to an individual's pulse waveform.

5.5 Robustness to Skin Type

Our motivation for adopting a meta-learning approach is to improve generalization. One challenge with camera-based PPG methods, and PPG methods in general, is their sensitivity to skin type. A larger melanin concentration in people with darker skin leads to higher light absorption compared to lighter skin types [30], thus reducing the reflectance signal-to-noise ratio. Fig. 2 shows a bar plot of the MAE in heart rate estimates by skin type (we group types I+II and V+VI as there were relatively few subjects in these categories). Both the AFRL and UBFC datasets are heavily skewed towards lighter Caucasian skin type categories. Therefore supervised methods trained on these datasets (e.g., TS-CAN) tend to overfit and not perform well on other skin types. Entirely unsupervised baselines do not perform any better, possibly because they were mostly designed and validated with lighter skin type data as well. While the highest errors for unsupervised MetaPhys still come in the darkest skin type categories, the reduction in error for types V+VI is considerable (68% compared to POS, 50% compared to TS-CAN). We are encouraged that these results are a step towards more consistent performance across people of different appearances.

Figure 2: Left) MAE in HR estimates (12-second windows) for the UBFC and MMSE-HR datasets. Right) MAE in HR estimates by skin type on the MMSE-HR dataset. Standard error bars shown.

Figure 3: Left) Estimated HR and gold-standard HR reference measurements in MMSE and UBFC datasets and the corresponding Bland-Altman plots from TS-CAN [24]. Right) Estimated HR and gold-standard HR reference measurements in the MMSE and UBFC datasets and the corresponding Bland-Altman plots from MetaPhys.

5.6 Future Work

5.6.1 Federated Meta-Learning: Federated learning (FL) refers to a training paradigm in which a global model is trained across multiple users by exchanging weights/gradients via a central server. In federated learning, a user's data will be stored locally reducing

access to private information. For these reasons FL is attractive in the health domain [5, 12]. MetaPhys could be deployed in a federated learning system as personalized weights and gradients of each individual could easily be uploaded to the central server without sharing images or videos of the subject. By doing so, a more robust central model could be used as the backbone to further improve

Figure 4: (A) An illustration comparing the attention masks of five subjects. The masks were generated in using four training schemes: 1) traditional supervised training (TS-CAN), 2) TS-CAN with fine tuning, 3) supervised MetaPhys and 4) unsupervised MetaPhys. (B) An illustration comparing the attention masks in the learning progress from four training schemes.

the generalizability and performance. In the future, we will explore how to extend MetaPhys in federated settings.

5.6.2 Explainable MetaPhys: Although Figure 4 provides certain degree of insight as to why MetaPhys outperforms traditional fine tuning, it focuses on interpreting the spacial dimension and does not shed light on how well MetaPhys captures temporal information. PPG is a time-varying signal, therefore extracting unique temporal information in each individual is key. Towards more explainable remote physiological measurement, we plan to investigate why and how MetaPhys adapts to different faces and environments.

5.7 Limitations

There is a trend towards inferior performance when the skin type of the subject is darker. We acknowledge this limitation and plan to use resampling to help address this bias in future. Synthetic data may also be an answer to this problem and has shown promising early results [28]. Both the MMSE and UBFC datasets have somewhat limited head motion and future work will also investigate whether meta-learning can help with generalization to other motion conditions. In this paper, our proposed method is based on MAML [13], however, it is also worth to explore other other meta-learning algorithms. Moreover, we acknowledge that the datasets we evaluated on are non-clinical datasets. We plan to validate the robustness, feasibility and safety of our system in clinical adoptions.

6 CONCLUSIONS

In this paper, We present a novel few-shot adaptation framework for non-contact physiological measurement called MetaPhys. Our proposed method leverages an optically-grounded unsupervised learning technique and a modern meta-learning framework to jointly train a remote physiological network. MetaPhys is also the first work that uses pseudo labels in training a physiological sensing

model and the first unsupervised deep learning method in remote physiological measurement. Our results not only substantially improves on the state-of-the-art and the performance on various skin types, but also reveal why and how our method achieved such improvement.

7 APPENDIX

7.1 Evaluation Metrics

Mean Absolute Error (MAE): The MAE between our model estimates and the gold-standard heart rates from the contact sensor measurements were calculated as follows for each 12-second time window:

$$MAE = \frac{1}{T} \sum_{i=1}^T |HR_i - HR'_i| \quad (1)$$

Root Mean Squared Error (RMSE): The RMSE between our model estimates and the gold-standard heart rate from the contact sensor measurements were calculated as follows for each 12-second time window:

$$RMSE = \sqrt{\frac{1}{T} \sum_{i=1}^T (HR_i - HR'_i)^2} \quad (2)$$

In both cases, HR is the gold-standard heart rate and HR' is the estimated heart rate from the video respectively. The gold-standard HR frequency was determined from the calculated from gold-standard PPG signal (UBFC dataset) or blood pressure wave (MMSE dataset).

We also compute the Pearson correlation between the estimated heart rates and the gold-standard heart rates from the contact sensor measurements across all the subjects.

Signal-to-Noise Ratios (SNR): We calculate blood volume pulse signal-to-noise ratios (SNR) [10]. This captures the signal quality of the recovered pulse estimates without penalizing heart rate estimates that are slightly inaccurate. The gold-standard HR frequency was determined from the gold-standard PPG waveform (UBFC dataset) or blood pressure wave (MMSE dataset).

$$SNR = 10 \log_{10} \left(\frac{\sum_{f=30}^{240} ((U_t(f) \hat{S}(f))^2)}{\sum_{f=30}^{240} (1 - U_t(f) \hat{S}(f))^2} \right) \quad (3)$$

where $\hat{S}$ is the power spectrum of the BVP signal (S), f is the frequency (in BPM), HR is the heart rate computed from the gold-standard device and $U_t(f)$ is a binary template that is one for the heart rate region from HR-6 BPM to HR+6BPM and its first harmonic region from $2^*HR-12BPM$ to $2^*HR+12BPM$, and 0 elsewhere.

7.2 Code and Additional Results:

Available at: <https://github.com/xliucs/MetaPhys>.

REFERENCES

- [1] Freddy Abnoui, Guson Kang, John Giacomini, Alan Yeung, Shirin Zarafshar, Nicholas Vesom, Euan Ashley, Robert Harrington, and Celina Yong. 2019. A novel noninvasive method for remote heart failure monitoring: the EuleriAn video Magnification apPLications In heart Failure study (AMPLIFY). *NPJ digital medicine* 2, 1 (2019), 1–6.
- [2] Guha Balakrishnan, Fredo Durand, and John Guttag. 2013. Detecting pulse from head motions in video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3430–3437.
- [3] Nannapas Banluesombatkul, Pichayoot Ouappaphan, Pitshaporn Leelaarporn, Payongkit Lakhan, Busarakum Chaitusaney, Nattapong Jaimchariyatam, Ekapol Chuangsuwanich, Wei Chen, Huy Phan, Nat Dilokthanakul, et al. [n.d.]. MetaSleepLearner: A Pilot Study on Fast Adaptation of Bio-signals-Based Sleep Stage Classifier to New Individual Subject Using Meta-Learning. ([n.d.]).
- [4] Serge Bobbia, Richard Macwan, Yannick Benezeth, Alamin Mansouri, and Julien Dubois. 2019. Unsupervised skin tissue segmentation for remote photoplethysmography. *Pattern Recognition Letters* 124 (2019), 82–90.
- [5] Theodora S Brisimi, Ruidi Chen, Theofanie Mela, Alex Olshevsky, Ioannis Ch Paschalidis, and Wei Shi. 2018. Federated learning of predictive models from federated electronic health records. *International journal of medical informatics* 112 (2018), 59–67.
- [6] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. 77–91.
- [7] Weixuan Chen and Daniel McDuff. 2018. Deepphys: Video-based physiological measurement using convolutional attention networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 349–365.
- [8] Weixuan Chen and Daniel McDuff. 2020. DeepMag: Source-Specific Change Magnification Using Gradient Ascent. *ACM Transactions on Graphics (TOG)* 40, 1 (2020), 1–14.
- [9] Janghoon Choi, Junseok Kwon, and Kyoung Mu Lee. 2019. Deep meta learning for real-time target-aware visual tracking. In *Proceedings of the IEEE International Conference on Computer Vision*. 911–920.
- [10] Gerard De Haan and Vincent Jeanne. 2013. Robust pulse rate from chrominance-based rPPG. *IEEE Transactions on Biomedical Engineering* 60, 10 (2013), 2878–2886.
- [11] Justin R Estepp, Ethan B Blackford, and Christopher M Meier. 2014. Recovering pulse rate during motion artifact with a multi-imager array for non-contact imaging photoplethysmography. In *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 1462–1469.
- [12] Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. 2020. Personalized federated learning: A meta-learning approach. *arXiv preprint arXiv:2002.07948* (2020).
- [13] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. 1126–1135.
- [14] Thomas B Fitzpatrick. 1988. The validity and practicality of sun-reactive skin types I through VI. *Archives of dermatology* 124, 6 (1988), 869–871.
- [15] Taesik Gong, Yeonsu Kim, Jinwoo Shin, and Sung-Ju Lee. 2019. MetaSense: few-shot adaptation to untrained conditions in deep mobile sensing. In *Proceedings of the 17th Conference on Embedded Networked Sensor Systems*. 110–123.
- [16] Edward Grefenstette, Brandon Amos, Denis Yarats, Phu Mon Htut, Artem Molchanov, Franziska Meier, Douwe Kiela, Kyunghyun Cho, and Soumith Chintala. 2019. Generalized inner loop meta-learning. *arXiv preprint arXiv:1910.01727* (2019).
- [17] Junfeng He, Khoi Pham, Nachiappan Valliappan, Pingmei Xu, Chase Roberts, Dmitry Lagun, and Vidhya Navalpakkam. 2019. On-device few-shot personalization for real-time gaze estimation. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*. 0–0.
- [18] Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. 2020. Meta-learning in neural networks: A survey. *arXiv preprint arXiv:2004.05439* (2020).
- [19] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [20] Eugene Lee, Evan Chen, and Chen-Yi Lee. 2020. Meta-rPPG: Remote Heart Rate Estimation Using a Transductive Meta-Learner. *Proceedings of the European Conference on Computer Vision (ECCV)* (2020).
- [21] Jessica Lee, Deva Ramanan, and Rohit Girdhar. 2019. Metapix: Few-shot video retargeting. *arXiv preprint arXiv:1910.04742* (2019).
- [22] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. 2017. Meta-sgd: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835* (2017).
- [23] Ji Lin, Chuang Gan, and Song Han. 2019. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE International Conference on Computer Vision*. 7083–7093.
- [24] Xin Liu, Josh Fromm, Shwetak Patel, and Daniel McDuff. 2020. Multi-Task Temporal Shift Attention Networks for On-Device Contactless Vitals Measurement. *arXiv preprint arXiv:2006.03790* (2020).
- [25] Andrea Madotto, Zhaojiang Lin, Chien-Sheng Wu, and Pascale Fung. 2019. Personalizing dialogue agents via meta-learning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 5454–5459.
- [26] Daniel McDuff and Ethan Blackford. 2019. iphys: An open non-contact imaging-based physiological measurement toolbox. In *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 6521–6524.
- [27] Daniel McDuff, Sarah Gontarek, and Rosalind W Picard. 2014. Remote detection of photoplethysmographic systolic and diastolic peaks using a digital camera. *IEEE Transactions on Biomedical Engineering* 61, 12 (2014), 2948–2954.
- [28] Daniel McDuff, Javier Hernandez, Erroll Wood, Xin Liu, and Tadas Baltrusaitis. 2020. Advancing Non-Contact Vital Sign Measurement using Synthetic Avatars. *arXiv preprint arXiv:2010.12949* (2020).
- [29] Daniel J McDuff, Ethan B Blackford, and Justin R Estepp. 2017. The impact of video compression on remote cardiac pulse measurement using imaging photoplethysmography. In *2017 12th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017)*. IEEE, 63–70.
- [30] Ewa M Nowara, Daniel McDuff, and Ashok Veeraraghavan. 2020. A Meta-Analysis of the Impact of Skin Tone and Gender on Non-Contact Photoplethysmography Measurements. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 284–285.
- [31] Eunbyung Park and Alexander C Berg. 2018. Meta-tracker: Fast and robust online adaptation for visual object trackers. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 569–585.
- [32] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*. 8026–8037.
- [33] Ming-Zher Poh, Daniel J McDuff, and Rosalind W Picard. 2010. Advancements in noncontact, multiparameter physiological measurements using a webcam. *IEEE transactions on biomedical engineering* 58, 1 (2010), 7–11.
- [34] Ming-Zher Poh, Daniel J McDuff, and Rosalind W Picard. 2010. Non-contact, automated cardiac pulse measurements using video imaging and blind source separation. *Optics express* 18, 10 (2010), 10762–10774.
- [35] Valentina O Puntmann, M Ludovica Carerj, Imke Wieters, Masia Fahim, Christophe Arendt, Jędrzej Hoffmann, Anastasia Schendrygina, Felicitas Escher, Mariuca Vasa-Nicotera, Andreas M Zeiher, et al. 2020. Outcomes of cardiovascular magnetic resonance imaging in patients recently recovered from coronavirus disease 2019 (COVID-19). *JAMA cardiology* (2020).
- [36] Anthony C Smith, Emma Thomas, Cainte L Snoswell, Helen Haydon, Ateev Mehrotra, Jane Clemensen, and Liam J Caffery. 2020. Telehealth for global emergencies: Implications for coronavirus disease 2019 (COVID-19). *Journal of telemedicine and telecare* (2020), 1357633X20916567.
- [37] Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*. 4077–4087.
- [38] Chihiro Takano and Yuji Ohta. 2007. Heart rate measurement based on a time-lapse image. *Medical engineering & physics* 29, 8 (2007), 853–857.
- [39] Wim Verkruyse, Lars O Svaasand, and J Stuart Nelson. 2008. Remote plethysmographic imaging using ambient light. *Optics express* 16, 26 (2008), 21434–21445.
- [40] Mauricio Villarroel, Sitthichok Chaichulee, João Jorge, Sara Davis, Gabrielle Green, Carlos Arteta, Andrew Zisserman, Kenny McCormick, Peter Watkinson,

- and Lionel Tarassenko. 2019. Non-contact physiological monitoring of preterm infants in the Neonatal Intensive Care Unit. *npj Digital Medicine* 2, 1 (2019), 1–18.
- [41] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. 2016. Matching networks for one shot learning. In *Advances in neural information processing systems*. 3630–3638.
- [42] Wenjin Wang, Albertus C den Brinker, Sander Stuijk, and Gerard de Haan. 2016. Algorithmic principles of remote PPG. *IEEE Transactions on Biomedical Engineering* 64, 7 (2016), 1479–1491.
- [43] Hao-Yu Wu, Michael Rubinstein, Eugene Shih, John Guttag, Frédo Durand, and William Freeman. 2012. Eulerian video magnification for revealing subtle changes in the world. *ACM transactions on graphics (TOG)* 31, 4 (2012), 1–8.
- [44] Zitong Yu, Xiaobai Li, and Guoying Zhao. 2019. Remote photoplethysmograph signal measurement from facial videos using spatio-temporal networks. In *Proc. BMVC*. 1–12.
- [45] Qi Zhan, Wenjin Wang, and Gerard de Haan. 2020. Analysis of CNN-based remote-PPG to understand limitations and sensitivities. *Biomedical Optics Express* 11, 3 (2020), 1268–1283.
- [46] Zheng Zhang, Jeff M Girard, Yue Wu, Xing Zhang, Peng Liu, Umur Ciftci, Shaun Canavan, Michael Reale, Andy Horowitz, Huiyuan Yang, et al. 2016. Multimodal spontaneous emotion corpus for human behavior analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3438–3446.
- [47] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. 2018. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 8697–8710.