

Unsupervised Adversarial Domain Adaptation in Wearable Physiological Sensing for Construction Workers' Health Monitoring Using Photoplethysmography

Yogesh Gautam, S.M.ASCE¹; Yizhi Liu, Ph.D., A.M.ASCE²;
and Houtan Jebelli, Ph.D., A.M.ASCE³

¹Ph.D. Student, Dept. of Civil and Environmental Engineering, Univ. of Illinois Urbana-Champaign, Champaign, IL. Email: ygautam2@illinois.edu

²Assistant Professor, Dept. of Civil and Environmental Engineering, Syracuse Univ., Syracuse, NY. Email: yk15505@psu.edu

³Assistant Professor, Dept. of Civil and Environmental Engineering, Univ. of Illinois Urbana-Champaign, Champaign, IL (corresponding author). Email: hjebelli@illinois.edu

ABSTRACT

Recent advancements in wearable physiological sensing and artificial intelligence have made some remarkable progress in workers' health monitoring in construction sites. However, the scalable application is still challenging. One of the major complications for deployment has been the distribution shift observed in the physiological data obtained through sensors. This study develops a deep adversarial domain adaptation framework to adapt to out-of-distribution data (ODD) in the wearable physiological device based on photoplethysmography (PPG). The domain adaptation framework is developed and validated with reference to the heart rate predictor based on PPG. A heart rate predictor module comprising feature generating encoder and predictor is initially trained with data from a given training domain. An unsupervised adversarial domain adaptation method is then implemented for the test domain. In the domain adaptation process, the encoder network is adapted to generate domain invariant features for the test domain using discriminator-based adversarial optimization. The results demonstrate that this approach can effectively accomplish domain adaptation, as evidenced by a 27.68% reduction in heart rate prediction error for the test domain. The proposed framework offers potential for scaled adaptation in the jobsite by addressing the ODD problem.

INTRODUCTION

The digital healthcare sector has experienced a rapid expansion, with its current market size reaching 211 billion USD ("Digital Health Market Size, Share & Trends Report, 2030"). One of the factors driving this growth is the increasing popularity of wearable health monitoring devices. These devices have become particularly appealing in industries like construction, where worker health has been a major concern (Ojha et al. 2023). Wearable health monitoring systems offer several advantages in this context, including continuous tracking of physical activity, non-invasive detection of health conditions, and monitoring of vital signs.

Various wearable health monitoring devices have been studied for the construction sites like electrocardiogram (ECG), electromyography (EMG), and photoplethysmography (PPG) (Jebelli et al. 2018; Ojha et al. 2023; Shakerian et al. 2021). Among these, PPG has gained considerable importance in recent years due to its ability to measure multiple physiological parameters, including heart rate, blood pressure, and oxygen saturation levels (Temko 2017), and detect cardio-pulmonary problems like atrial fibrillation, sleep apnea, hypertension etc. (Elgendi et al.

2019). While past investigations have displayed the feasibility of PPG for construction workers' health monitoring (Jebelli et al. 2018), wide-scale adaptation in the construction industry is still lagging.

One of the major barriers to a wide-range adaptation of wearable health monitoring in the construction industry has been the wide range of distribution shifts observed in the physiological signals across various times, locations, and populations. It occurs when the distribution of the input data changes between the training and testing phases, leading to a mismatch between the two distributions. This mismatch between the marginal distribution can result in poor generalization and reduced accuracy of the model. Distribution shifts in physiological signals can arise due to various reasons, such as changes in the environment, changes in the population demographics being studied, changes in the data collection process, or changes in time (Shimodaira 2000). Construction sites by nature are very dynamic having wide variations in the working environment, workers' demography, and physical and mental demands of the tasks. To enable wearable health monitoring adaptation on a larger scale in the construction industry the distribution shift problem must be solved.

The study designs, develops, and validates an unsupervised adversarial domain adaptation-based neural network framework for adapting ODD to the training distribution for PPG signals. The framework uses a convolution-based heart rate predictor to extract features from the PPG signals. Adversarial domain adaptation is used to adapt the distribution of these features for the test domain. The framework is validated by leveraging the adapted features to predict heart rate. The results indicate that the approach can provide consistent accuracy for the ODD, with over 27.68% reduction in error for heart rate prediction on the test domain.

This study contributes to the existing field of knowledge by providing a holistic framework for domain adaptation to solve the ODD problem for PPG signals. The study is the first of its type to introduce this concept of adversarial domain adaptation which is popular in computer vision for physiological signals based on photoplethysmography.

WEARABLE BIOSENSORS AND MACHINE LEARNING FOR WORKER'S HEALTH MONITORING IN THE FIELD

The use of wearable biosensors and machine learning algorithms has shown tremendous potential in monitoring the health of workers in the construction industry. In the field, workers are exposed to various environmental hazards, including high temperatures, noise, dust, and chemicals (Mo et al. 2022). These hazards can have a significant impact on the workers' health, both physical and mental, causing various illnesses and injuries. Monitoring the workers' health in real-time can help detect early signs of health issues and prevent more severe consequences. Moreover, a healthy workforce can directly contribute to increasing productivity.

Wearable biosensors are small devices that can be attached to the workers' body and can collect various physiological data, such as heart rate, blood pressure, temperature, and oxygen saturation. Various sensors based on ECG, EMG, and PPG signals have been examined in the literature for health monitoring in the construction industry (Hwang and Lee 2017; Ojha et al. 2023).

PPG devices, because of their small size, easy usage, and access to important cardio-pulmonary information have been a popular choice for health monitoring in the construction industry as well as commercial usage. PPG signals detect changes in blood volume and oxygen saturation levels, making them a useful tool for monitoring cardiovascular health (Temko 2017).

Various research has been done using PPG to estimate a variety of cardiovascular activities, like heart rate, oxygen saturation levels, heart rate variability, etc., and multiple diagnoses like peripheral artery diseases, sleep apnea, arterial stiffness, hypertension, etc. (Elgendi et al. 2019). In the author's prior study, the feasibility of utilizing PPG to assess the physical and mental well-being of construction workers was examined and confirmed (Jebelli et al. 2018). Other investigations have explored further usage of PPG in construction jobsites. For instance, Shakerian et al. developed a machine learning-based heat stress assessment system for construction jobsites (Shakerian et al. 2021); Ojha et al. performed statistical analysis to analyze PPG response over varying levels of dehydration for construction workers (Ojha et al. 2020).

Though machine learning algorithms can provide robust data-driven solutions for various health monitoring applications, they may suffer from distribution shift problems if there is a difference in marginal distribution between the training and testing set which may usually occur in a real-life scenario (Lipton et al. 2018). Inference based on the test set with different distribution than the marginal distribution can provide inaccurate detections.

DISTRIBUTION SHIFTS IN THE DATA: ONE MAJOR CHALLENGE OF ML ALGORITHMS FOR HEALTH MONITORING IN THE FIELD

There have been several investigations in the literature to solve the distribution shift problem. Some common statistical methods are importance weighting (Shimodaira 2000), and multi-model method (Jp et al. 2007). Transfer learning has been a popular method for domain adaptation for neural network-based architectures. Some recent investigations have explored transfer learning-based solutions for PPG and ECG signals which have demonstrated improved accuracy (Li et al. 2021; Radha et al. 2021). While transfer learning may be favorable in the presence of labeled data, labeled data may not be necessarily available in real-world scenarios like say construction settings where environment, working conditions, and workers' demographics change frequently. Also, while statistical methods have displayed higher performance for domain adaptation, it requires detailed knowledge of data distribution and subject variation which may not be tedious in a real-life application like a construction environment. This demands unsupervised methods which can perform in the absence of labeled data as well as deep theoretical knowledge of the data distributions. Adversarial domain adaptation has developed as an effective deep-learning framework for unsupervised domain adaptation in the last few years.

Adversarial domain adaptation is influenced by Generative Adversarial Network (GANs) that has emerged as a powerful tool in the field of deep learning, with diverse applications including image and video generation, text-to-image synthesis, and domain adaptation (Goodfellow et al. 2020). Similar to GAN, in an adversarial domain adaptation model, the generator learns to generate domain invariant features for a target domain, while the discriminator attempts to distinguish between the features from the source and target domains. Both the generator and the discriminator are trained iteratively in a game-like manner, where the generator tries to produce more realistic domain invariant features to fool the discriminator, and the discriminator tries to differentiate between the domains. Adversarial domain adaptation aims to train a model on a source domain and then adapt it to a target domain by introducing an adversarial loss that forces the feature representations of the source and target domains to be indistinguishable. Various architectures and training strategies have been proposed for adversarial domain adaptation (Volpi et al. 2017). While these techniques have achieved high performance for domain adaptation for 2D images their validation for 1D physiological signals is very limited.

To tackle the distribution shift problem observed in physiological signals using an unsupervised method, the study designs, develops, and validates an adversarial domain adaptation framework based on the PPG signal. The framework attempts to reduce the distribution shift for a test domain. The experiments and validation for this study have been performed in a construction setting. By developing this framework, this study paves the way for the scalable deployment and more advanced applications of PPG in health monitoring for the construction industry.

ADVERSARIAL DOMAIN ADAPTATION FOR SOLVING THE DISTRIBUTION SHIFT IN PPG SIGNALS

The developed framework, illustrated in Figure 1, comprises of two phases. In the first phase, a convolution-based encoder-predictor network is trained to extract features in the latent space and consecutively predict heart rate. The encoder network uses blood volume pulse (BVP) obtained through photoplethysmography as input and generates an 8-dimensional feature vector as output. The 8-dimensional feature vector is then used as input for the predictor network to predict heart rate. The features generated by the encoder network are used as the reference for the domain adaptation method in this study. In the second phase, an encoder-discriminator-based min-max optimization is performed for adversarial domain adaptation. The encoder in the second phase uses the BVP as input and learns to produce features with a distribution similar to the features of the training domain. The discriminator is trained to differentiate between the distribution difference of the feature vectors existing between the training domain and the test domain generated by the respective encoders. This constant adversarial optimization between the encoder and discriminator network results in domain adapted encoder for the test domain. The domain-adapted encoder from the second phase is then used with the predictor from the first phase to predict the heart rate for the new domain. The heart rate prediction task acts as the validator for the domain adaptation method. Each component will be explained in detail in the following methodology sections, which include (1) Encoder-predictor network, and (2) Adversarial domain adaptation.




Figure 1: Overview of the developed framework for adversarial domain adaptation.

Encoder-predictor module. This subsection presents the encoder-predictor network, designed to predict heart rate from the acquired blood volume pulse (BVP) data through photoplethysmography. The module comprises two subnetworks, a multi-stage convolutional

encoder network, and a multi-layer perceptron-based predictor network. The encoder network (E) uses BVP data as input to extract an 8-dimensional feature vector. The predictor network (P) uses the feature generated by the encoder network as input to estimate heart rate. The architecture of the encoder predictor network is presented in Figure 2.

Figure 2 subnetwork A depicts the architecture of the proposed encoder network developed for feature extraction. The network leverages BVP signals from photoplethysmography as input to generate 8D features $f = \{f_1, f_2, \dots, f_8\}$, that can estimate heart rate of the given subject. The encoder network is constructed of two convolutional blocks followed by two dense layers, to extract features. Each convolutional block consists of 1D convolution followed by batch normalization and a one-dimensional max pool layer. For the first convolution block (C_1), kernel size of 5×1 is used to capture a larger temporal range whereas for the second convolution block (C_2), a kernel size of 3×1 is used. The predictor represented by Figure 2 subnetwork B comprises of two hidden dense layers of 4 and single units respectively. The final dense unit predicts the heart rate H_r .




Figure 2: Convolution-based encoder-predictor network for heart rate prediction based on BVP signals acquired from photoplethysmography.

To train the encoder-predictor network mean absolute error (MAE) is used as the loss function defined by Equation 1.

$$\mathcal{L}(\theta) = \frac{1}{n} * \sum_{i=1}^N |H_r^*(i) - H_r(i)| \quad (1)$$

In this equation, $H_r(i)$ indicates the network estimated heart rate and $H_r^*(i)$ indicates ground truth heart rate. For the training, the Adam optimizer is deployed with a batch size of 64 and a learning rate of 0.001. As the study intends to use a limited amount of data to mimic real-life scenarios, data from three subjects are only deployed during training with 80% of the data being used as the training set and 20% of the set used as the validation set. Using a smaller number of subjects during training facilitates a large number of subject choices in the testing set. Domain adaptation is implemented for the test subject with the highest heart rate estimation error. The results for the training and error for the test subjects are presented in the results section.

Adversarial domain adaptation. In the second phase of the framework, adversarial min-max optimization is performed to adapt the encoder to generate features having similar distribution to that of the training domain. The encoder from the first phase (E) is used to initialize the encoder of the second phase (E). E uses same architecture as that of E described in

the previous sub-section and portrayed in Figure 2 subnetwork A. Thus, E has BVP data points as the input and 8D domain adapted feature vector $f' = \{f'_1, f'_2, \dots, f'_8\}$ as the output. The purpose of the discriminator (D) is to distinguish between the features of the training domain and testing domain. The input to the discriminator is the concatenated vector of the unprocessed input BVP and the feature vector f' generated by the encoder E . The extra concatenation is performed to facilitate the discriminator with the information about the original signal and the corresponding generated features. As the discriminator performs binary classification tasks, the output of the discriminator is a binary value representing whether the input features are from a source or a target domain. The features generated by the encoder E based on the training cohort are used as the source feature and the feature generated by the adapted encoder E based on the test cohort is represented as target features during the training process of the discriminator.

The discriminator network is constructed of three convolutional blocks followed by a single unit of dense layers representing the binary output to distinguish between the features of the training and testing domain. Each convolutional block consists of 1D convolution followed by batch normalization and a one-dimensional max pool layer of kernel size two. The output from the last convolutional block is flattened followed by a single dense unit.

During training, for each epoch, first the discriminator is trained, and then the encoder (E) is trained freezing the discriminator. For both instances, the discriminator performs the binary classification task of distinguishing between features from the given domain so binary cross entropy is used as the loss function as given by equation 2. In equation 2, x_1 represents the training domain and x_2 is the test domain for which the encoder (E) is adapted through min-max optimization of the encoder and discriminator. The results from the adapted domain are presented in the Results section.

$$\mathcal{L}(E, D): \min(E), \max(D) = \mathbb{E}_{x_1 \in \text{train}} \log(D(x_1)) + \mathbb{E}_{x_2 \in \text{test}} \log(1 - D(E(x_2))) \quad (2)$$

EXPERIMENT AND DATASET DESCRIPTION

The proposed framework has been trained and tested with data acquired in a construction setting. An off-the-shelf wristband biosensor was used to acquire blood volume pulse signals using photoplethysmography from the non-dominant wrist of the subjects. The study involves 7 participants, with three of them used as training data and the remaining used as test data selected randomly.

The construction tasks involve three different activities as represented in Figure 3. The first task requires the subjects to carry a 2.3 Kg (5 lbs.) bag between a distance of 6 meters. This low-intensity transportation task is carried out for a total of 480 seconds, with a break of 120 seconds after the initial 180 seconds. The second high-intensity transportation is similar to the first but with a load of 13.6 Kg (30 lbs.) instead. In the third task, the subjects are required to tie up rebars, each measuring 0.5 inches in diameter and with a length of 4 feet. The rebars are arranged in a cross position within an area measuring 1.21m by 0.6m. This task is carried out for a total of 480 seconds, with a break of 120 seconds after the initial 180 seconds. During this task, the subjects tie up the rebars using tie wire and a manual wire twister. A public dataset has also been leveraged for further testing to test the framework against a more severe distribution shift as the public dataset has been acquired in a different location, and environmental conditions with different population demography (Reiss et al. 2019). Three subjects were chosen from the public

dataset, one randomly and two having the highest MAE for heart rate prediction which have been included in the test set as subjects 8, 9, and 10 respectively.




Figure 3: Material handling and rebar tying experimental setup; 3A: An overview of the experimental location; 3B: Subject performing rebar tying task with wristband highlighted in a white outline.

RESULTS AND DISCUSSION

This section presents all the results from the encoder-predictor network and adversarial domain adaptation. For the encoder-predictor network, during the training, mean absolute error (MAE) is used as the loss function, and Adam Optimizer is used along with training for 100 epochs with a batch size of 64. Twenty percent of the training data is used as the validation set. A callback function was deployed during training to save the model with the lowest validation. The lowest validation error of 7.5 BPM (beats per minute) was achieved with a training error of 7.14 BPM. The training and validation loss as a function of BPM is presented in Figure 4B.




Figure 4: Prediction of the encoder-predictor module; 4A: Sample prediction on subject 8 with MAE of 9.1 BPM; 4B: Training-validation loss for the encoder-predictor network.

The trained model when tested on the remaining subjects, the highest MAE (30.7 BPM) is observed for subject 9 followed by subject 10 (MAE: 26.4 BPM) as represented in Table 1 along

with mean squared error (MSE). A sample prediction for subject 8 is represented in Figure 4A using the initial encoder-predictor network.

Table 1: Heart rate prediction error for test subjects using the initial encoder-predictor network.

	Train	Val	S4	S5	S6	S7	S8	S9	S10
MAE	7.14	7.5	8.9	7.7	9.9	13.8	9.1	30.7	26.4
MSE	132	141	153	149	235	358	142	1490	1231

The subjects for the domain adaptation have been chosen based on the error observed from Table 1 and the results are presented in Table 2. Two subjects with the highest MAE and one subject with a low MAE, i.e., subject 9, subject 10, and subject 6 have been chosen for domain adaptation. For all the subjects, a separate domain adaptation model is trained. The resulting encoder network is validated by processing the output features of the adapted encoder through the predictor network. For the subject with the highest MAE a reduction of 27.68% is observed while for the subject with lower MAE the improvement is negligible. One of the possible reasons for this could be that as the subject already has MAE error close to the validation error, much improvement could not be achieved.

Table 2: Improvement in heart rate prediction after performing adversarial domain adaptation.

	MAE heartbeat rate		
	Before domain adaptation	After domain adaptation	% Reduction
Subject-6	9.9	9.82	0.8%
Subject-9	30.7	22.2	27.68%
Subject -10	26.4	20.1	23.86%

CONCLUSION

The study introduces an adversarial domain adaptation framework aimed at adapting out-of-distribution data based on photoplethysmography. In this study, the authors first developed a feature-extracting encoder network based on heart rate prediction and then adapted the encoder for the test domain. The performance of the framework was evaluated against heart rate prediction for the adapted encoder. Results indicate that the designed framework based on adversarial optimization can enhance domain adaptation for PPG signals. This research contributes to the scalability of PPG-based health monitoring for construction workers. Further extended examinations can be performed by using other forms of validators like blood pressure, respiration rate, etc. Additionally, more robust domain adaptation frameworks could be developed by acquiring and utilizing larger datasets and using more advanced architecture and adversarial training frameworks for better feature representation and adaptation. Research can also extend this study for testing the framework in more selective subjects with controlled environmental parameters like temperature, humidity, and subject demographics for studying

adaptability against population, and environmental variations. Overall, the proposed framework addresses the ODD problem and offers the potential for scaled adaptation in job sites.

ACKNOWLEDGEMENT

The work presented in this paper was supported financially by a National Science Foundation award (No. 2222654, ‘uture of onstruction Workplace Health Monitoring’). ny opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- Elgendi, M., R. Fletcher, Y. Liang, N. Howard, N. H. Lovell, D. Abbott, K. Lim, and R. Ward. 2019. “The use of photoplethysmography for assessing hypertension.” *npj Digital Medicine* 2019 2:1, 2 (1): 1–11. Nature Publishing Group.
- Goodfellow, I., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. 2020. “Generative adversarial networks.” *Commun ACM*, 63 (11): 139–144. ACM PUB27 New York, NY, USA.
- Hwang, S., and S. H. Lee. 2017. “Wristband-type wearable health devices to measure construction workers’ physical demands.” *Autom Constr*, 83: 330–340. Elsevier.
- Jebelli, H., B. Choi, H. Kim, and S. Lee. 2018. “Feasibility study of a wristband-type wearable sensor to understand construction workers’ physical and mental status.” *Construction Research Congress 2018: Construction Information Technology - Selected Papers from the Construction Research Congress 2018*, 2018-April: 367–377. American Society of Civil Engineers (ASCE).
- Jp, S. T. A., M. Krauledat, and K. F. De. 2007. “Covariate Shift Adaptation by Importance Weighted Cross Validation.” *The Journal of Machine Learning Research*, 8: 985–1005. JMLR.org PUB6573.
- Li, Q., Q. Li, A. S. Cakmak, G. Da Poian, D. L. Bliwise, V. Vaccarino, A. J. Shah, and G. D. Clifford. 2021. “Transfer learning from ECG to PPG for improved sleep staging from wrist-worn wearables.” *Physiol Meas*, 42 (4): 044004. IOP Publishing.
- Lipton, Z. C., Y.-X. Wang, and A. J. Smola. 2018. *Detecting and Correcting for Label Shift with Black Box Predictors*. PMLR.
- Mo, S., Y. Wang, and F. Xiong. 2022. “Identification and prioritization of key health hazards to workers in roadway construction.” *Transportation Safety and Environment*, 4 (2). Oxford Academic.
- Ojha, A., S. Shakerian, M. Habibnezhad, and H. Jebelli. 2023. “Feasibility Verification of Multimodal Wearable Sensing System for Holistic Health Monitoring of Construction Workers.” *Lecture Notes in Civil Engineering*, 239: 283–294. Springer Science and Business Media Deutschland GmbH.
- Ojha, A., S. Shakerian, M. Habibnezhad, H. Jebelli, S. Lee, and M. S. Fardhosseini. 2020. *Feasibility of Using Physiological Signals from a Wearable Biosensor to Monitor Dehydration of Construction Workers*. 4.
- Radha, M., P. Fonseca, A. Moreau, M. Ross, A. Cerny, P. Anderer, X. Long, and R. M. Aarts. 2021. “A deep transfer learning approach for wearable sleep stage classification with photoplethysmography.” *npj Digital Medicine* 2021 4:1, 4 (1): 1–11. Nature Publishing Group.

- Reiss, A., I. Indlekofer, P. Schmidt, and K. Van Laerhoven. 2019. "Deep PPG: Large-Scale Heart Rate Estimation with Convolutional Neural Networks." *Sensors* 2019, Vol. 19, Page 3079, 19 (14): 3079. Multidisciplinary Digital Publishing Institute.
- Shakerian, S., M. Habibnezhad, A. Ojha, G. Lee, Y. Liu, H. Jebelli, and S. H. Lee. 2021. "Assessing occupational risk of heat stress at construction: A worker-centric wearable sensor-based approach." *Saf Sci*, 142: 105395. Elsevier.
- Shimodaira, H. 2000. "Improving predictive inference under covariate shift by weighting the log-likelihood function." *J Stat Plan Inference*, 90 (2): 227–244. North-Holland.
- Temko, A. 2017. "Accurate Heart Rate Monitoring during Physical Exercises Using PPG." *IEEE Trans Biomed Eng*, 64 (9): 2016–2024. IEEE Computer Society.
- Volpi, R., P. Morerio, S. Savarese, and V. Murino. 2017. "Adversarial Feature Augmentation for Unsupervised Domain Adaptation." *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 5495–5504. IEEE Computer Society.