Towards unsupervised wearable representations for longitudinal cardio-fitness prediction

Dimitris Spathis, Ignacio Perez-Pozuelo, Soren Brage,
Nicholas J. Wareham, Cecilia Mascolo
University of Cambridge

Abstract

Introduction. Wearables and smartphones are becoming increasingly popular tools for objectively monitoring physical activity in the real world [1]. To date, research has primarily focused on the purely supervised task of human activity recognition, demonstrating limited success in inferring high-level health outcomes from low-level signals. In this abstract, we present our recent work in task-agnostic physiological representations as well as models tailored to fitness prediction.

Task-agnostic physiological representations with self-supervision. We first design a novel self-supervised representation learning method using wearable activity and heart rate (HR) signals without semantic labels. With a deep neural network, we set the HR responses as the *supervisory signal* for the activity data, leveraging their underlying physiological relationship. In addition, we propose a custom quantile loss function that accounts for the long-tailed HR distribution present in the general population. We evaluate our model in the *Fenland Study*, the largest free-living combined-sensing dataset (comprising >280k hours of wrist accelerometer & wearable ECG data). We find that the pre-training task creates a model that can accurately forecast HR based only on cheap activity sensors and that we can leverage the information captured through this task through downstream transfer learning with linear classifiers, capturing physiologically meaningful, personalized information. For instance, the resulting embeddings can predict outcomes associated with individuals' health, fitness and demographic characteristics (AUC >70), outperforming unsupervised autoencoders and common bio-markers [2].

Improving cardio-fitness prediction with free-living wearables. Second, acknowledging fitness as a strong predictor of overall health and mortality [3], which, however, can only be measured with expensive instruments (e.g. VO_2max test), we further develop supervised models for accurate prediction of fine-grained fitness levels. Although fitness can be approximated using resting heart rate and self-reported exercise habits, their accuracy is low compared to estimates based on dynamic data. Modern wearables capture dynamic sensor data which could improve fitness prediction. In this work, we analyze ECG-measured movement and heart rate signals in free-living conditions from a bigger cohort of the *Fenland Study*, who also underwent a standard exercise test (n=*11,059*). Our neural network leverages time-series features and traditional metadata to predict VO_2max, yielding a high correlation (r = 0.82) in a holdout sample. This model outperforms conventional non-exercise fitness models and traditional bio-markers, using measurements of normal daily living without the need to undertake a specific exercise test. Additionally, we show the adaptability of this approach for predicting fitness change over time in a longitudinal subsample (n = *2,675*) who repeated measurements after almost a decade. The latent representations that arise from such a model pave the way for fitness-aware population subtyping and interventions at scale.

Conclusion. Models which can aggregate high-resolution sensor data and map them to meaningful outcomes will become even more important due to recent improvements in on-device computation and the accumulation of larger datasets. Overall, here we discuss new unsupervised and supervised models for behavioral and physiological data with implications for large-scale health and lifestyle monitoring.

Keywords— deep learning, self-supervision, wearable sensor time-series, longitudinal data

References

[1] Althoff, T., Hicks, J. L., King, A. C., Delp, S. L., & Leskovec, J. (2017). Large-scale physical activity data reveal worldwide activity inequality. *Nature*, 547 (7663), 336-339.

[2] Spathis, D., Perez-Pozuelo, I., Brage, S., Wareham, N. J., & Mascolo, C. (2021). Self-supervised transfer learning of physiological representations from free-living wearable data. *In ACM Proceedings of the Conference on Health, Inference, and Learning (CHIL).*

[3] Mandsager, K., Harb, S., Cremer, P., Phelan, D., Nissen, S. E., Jaber, W. (2018). Association of cardiorespiratory fitness with long-term mortality among adults undergoing exercise treadmill testing. *JAMA network open,* 1(6), e183605-e183605.