<u>Viewpoint</u>

# Personalization of AI Using Personal Foundation Models Can Lead to More Precise Digital Therapeutics

Peter Washington, PhD

Department of Medicine, Division of Clinical Informatics and Digital Transformation, University of California, San Francisco, San Francisco, CA, United States

**Corresponding Author:**
Peter Washington, PhD
Department of Medicine
Division of Clinical Informatics and Digital Transformation
University of California, San Francisco
10 Koret Way
San Francisco, CA, 94143
United States
Phone: 1 (415) 353 2067
Email: peter.washington@ucsf.edu

## Abstract

Digital health interventions often use machine learning (ML) models to make predictions of repeated adverse health events. For example, models may be used to analyze patient data to identify patterns that can anticipate the likelihood of disease exacerbations, enabling timely interventions and personalized treatment plans. However, many digital health applications require the prediction of highly heterogeneous and nuanced health events. The cross-subject variability of these events makes traditional ML approaches, where a single generalized model is trained to classify a particular condition, unlikely to generalize to patients outside of the training set. A natural solution is to train a separate model for each individual or subgroup, essentially overfitting the model to the unique characteristics of the individual without negatively overfitting in terms of the desired prediction task. Such an approach has traditionally required extensive data labels from each individual, a reality that has rendered personalized ML infeasible for precision health care. The recent popularization of self-supervised learning, however, provides a solution to this issue: by pretraining deep learning models on the vast array of unlabeled data streams arising from patient-generated health data, personalized models can be fine-tuned to predict the health outcome of interest with fewer labels than purely supervised approaches, making personalization of deep learning models much more achievable from a practical perspective. This perspective describes the current state-of-the-art in both self-supervised learning and ML personalization for health care as well as growing efforts to combine these two ideas by conducting self-supervised pretraining on an individual's data. However, there are practical challenges that must be addressed in order to fully realize this potential, such as human-computer interaction innovations to ensure consistent labeling practices within a single participant.

## Introduction

In recent years, the intersection of consumer digital health and machine learning (ML) has emerged to enable ML-powered digital therapeutics, which have been developed in areas such as interventions for substance use [1-4]; technologies for managing mental health conditions such as anxiety, stress, and depression [5-8]; and autism therapeutics using Google Glass [9,10]. The models powering these digital therapies typically analyze large streams of an individual patient's data in order to anticipate adverse health events or actionable patient-reported outcomes. However, a significant computational challenge arises when dealing with the prediction of nuanced and subjective health events that are typically self-reported by participants in the form of patient-reported outcomes, such as mental health states like stress and anxiety. For such prediction targets, the cross-subject variability poses an obstacle for traditional ML approaches, as one participant's label of "moderately stressed" might be another participant's "lightly stressed".

Conventional ML methodologies typically involve training a single generalized model to classify a specific condition [11], such as for diagnostic or screening purposes. However, attempting to apply a universal model often leads to poor generalization to individuals and health systems that were not represented in the training data. An alternative solution involves training separate models for each individual or subgroup, tailoring the model to the unique characteristics of the patient. However, this approach would traditionally demand extensive labeled data from each participant, a requirement that has historically hindered the feasibility of personalized ML applications in precision health care.

The relatively recent advent of self-supervised learning (SSL), made popular in the context of pretraining large language models like ChatGPT (OpenAI), has enabled a transformative solution to address the challenges associated with personalized ML in health care [12-14]. SSL is a machine learning paradigm in which a model is trained to understand and represent the underlying structure of its input data without relying on externally provided labels. By pretraining deep learning models on vast amounts of unlabeled data streams derived from patient-generated health data to understand the baseline temporal dynamics of the data stream without a single label, SSL provides a means to fine-tune personalized models with significantly fewer labeled data points than when using traditionally supervised learning. This relatively new paradigm opens new avenues for making ML personalization in health care more practical, thereby overcoming one of the major hurdles that has historically impeded progress in this area.

This perspective explores the integration of SSL and personalization in scenarios where there are large unlabeled data streams generated per patient, focusing in particular on the potential of personalized SSL to improve the performance of digital therapeutics that provide some sort of digital therapy or digital intervention when a prediction about the participant in question is made by an ML model.

## Personalized Models in Health Care

Traditional ML methodologies, which often rely on a one-size-fits-all model, face substantial challenges when confronted with the diverse and nuanced nature of health outcomes. The need for personalized models that cater to individual characteristics has led to a paradigm where a single ML model is trained on data streams coming from a single user and evaluated on future data coming from that same user (Figure 1).

Several examples of personalized ML models for health care have been published in the past decade. Zhang et al [15] developed Patient2Vec, a representation learning approach for longitudinal electronic health record data used to predict clinical events into the future. Luu et al [16] trained a generalized model that was then fine-tuned to predict step count in a personalized manner, achieving 98%-99% accuracy in the personalized case and 96%-99% accuracy with the generalized models. Li et al [17] compared a personalized model for stress prediction against 2 baselines, subject-inclusive and subject-exclusive generalized models, finding that the personalized models significantly outperformed both sets of generalized models. This finding indicates that personalization using only an individual's data outperforms personalization when combining the personal data with data from other users, at least for highly heterogeneous outcomes such as affective computing.

Federated learning, where distributed local models are trained and sent to a central global server for weight aggregation, is naturally connected to the idea of personalized ML. Each "local" model is, by definition, a personalized model. Federated learning has been successfully applied to certain health care settings. For example, Rudovic et al [18] developed a personalized federated learning approach for pain estimation from face images where clients train models using local data, aggregate the model weights in a central server, and then send the global model back to the clients for fine-tuning. This federated learning approach enables the classification of a traditionally difficult classification task due to its inherent subjectivity and heterogeneity between individuals, namely, pain estimation using computer vision.

Traditional applications of personalized ML apply to scenarios where there are vast amounts of data labels per patient. Unfortunately, this situation is often unattainable. In contexts where the data labels pertain to patient-generated health data, it is especially infeasible to collect many labels. To address this practical issue with traditional personalized ML, this perspective explores the idea of performing SSL on an individual's unlabeled data streams to create a personalized foundation model.

**Figure 1.** In many biomedical domains, there exist massive unlabeled data streams with sparse annotations of the health event of interest. In personalized self-supervised learning, we can pretrain on data coming earlier from the participant and then fine-tune on an ideally small number of patient-provided labels. Evaluation then occurs later temporally. HR: heart rate; SpO2: oxygen saturation.
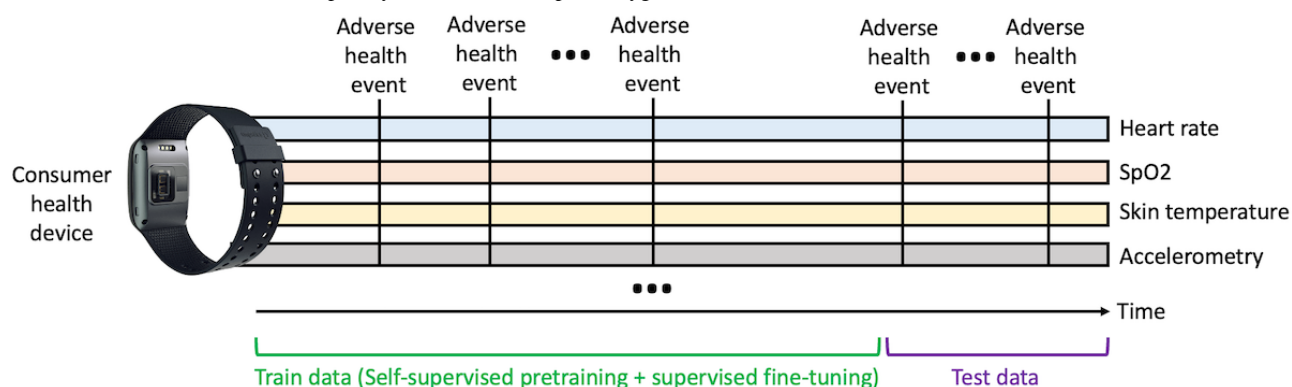
## Personalized Foundation Models: Combining Personalization With Self-Supervised Learning

SSL holds great promise to improve the performance of ML models in health care [19], broadly speaking. SSL involves leveraging the inherent information within the data itself to create supervisory signals for training. SSL has been traditionally applied to large datasets containing data from a broad array of patients. In passive data generation contexts, however, such as when patients wear a monitor that continuously collects biosignals, it can be productive to run SSL separately for each patient, as each patient has a large amount of data sampled several times a second. These separate pretraining procedures per patient can result in a "personal foundation model." Because foundation models can learn using much less data than would have been required if no SSL took place, the personal foundation models can enable learning of complex health outcomes where the supervisory signal drastically varies across patients.

SSL for personalization of longitudinal time series data for health care can be achieved through a variety of adaptations of popular SSL pretraining strategies (Figure 2). An inherently multimodal approach is to predict the missing portion of a signal given the values of signals from separate data modalities (Figure 2A) [20], treating the prediction as a multiple-output regression task [21]. Another approach is to perform contrastive learning algorithms such as SimCLR [22] on the signals to maximize representational similarity between augmented versions of the same time period while minimizing similarity between 2 distinct time windows (Figure 2B) [23,24]. More sophisticated generative approaches, such as masked autoencoders [25] and latent masking [26], can also be used to predict masked portions of input signals (Figure 2C), including in a multimodal manner [27].

Personalized modeling combined with SSL has recently enabled the successful prediction of traditionally heterogeneous and subjective health outcomes. For example, Li and Sano [28] used unsupervised representation learning to predict outcomes related to wellbeing, such as mood and stress. Li et al [29] computed personalized brain function networks from functional magnetic resonance imaging using SSL. Spathis et al [30] used SSL to learn user-specific representations of wearable data streams and demonstrated that these personalized representations can be fine-tuned to a variety of downstream tasks.

One important consideration is that increases in model performance might be due to either the personalization aspect or the SSL aspect. SSL without personalization has been repeatedly documented to improve ML model performance [31-34]. Thus, it is important to systematically isolate both conditions in isolation as baselines to determine the true contribution of each component.

Another caveat to personalized SSL is that within-subject consistency in labeling is crucial, and initial studies have found that improvement gains observed using personalized SSL require consistency in data labeling within a user. For example, Islam and Washington [35,36] applied personalized multimodal SSL to the Wearable Stress and Affect Detection dataset [37], observing significant improvements in model performance when compared to a baseline model using identical data without self-supervised pretraining. By contrast, Eom et al [38] evaluated a multimodal dataset collected by Hosseini et al [39] consisting of wearable biosensors measured from nurses working during the COVID-19 outbreak. Eom et al [38] did not observe increased performance on average when using personalized models pretrained on each individual's data compared to baseline models, likely due to particularly noisy and irregular data collection procedures arising from nurses providing data during a stressful event. This highlights the importance of using datasets that have consistent labeling within a participant in order to make personalized SSL actually work.

**Figure 2.** Examples of self-supervised learning approaches for longitudinal time series data. (A) An inherently multimodal approach is to predict the missing portion of a signal given the values of signals from separate data modalities. (B) Another approach is to perform contrastive learning on the signals by training a network to maximize similarity between a data point and an augmented version of that data point while minimizing similarity between that data point and a separate data point. (C) A third possible strategy is to predict the missing portion of a signal using a masked autoencoder or similar model.
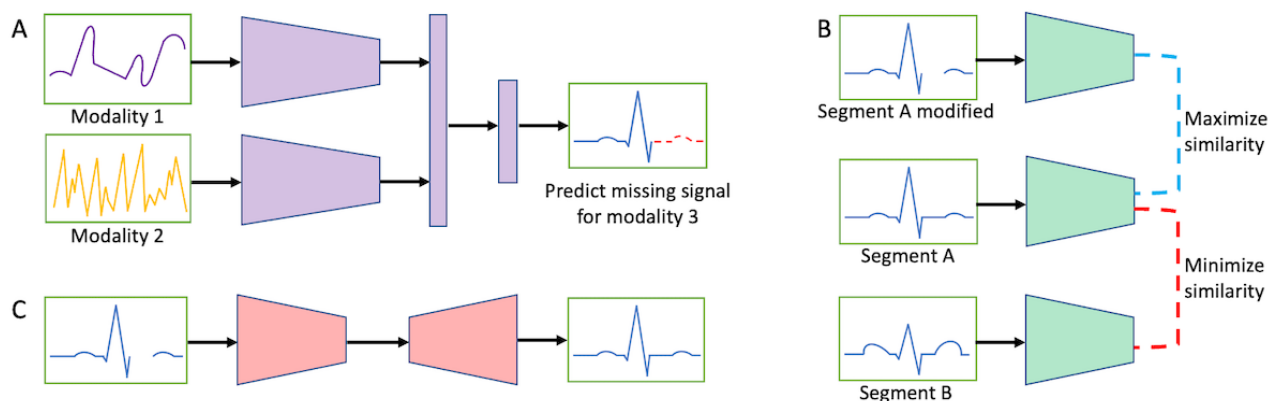
## Future Opportunities

Applications of personalized SSL to recurrent health predictions have been successful thus far under clean data scenarios. By harnessing the power of SSL, these applications have demonstrated the ability to glean intricate patterns and dependencies within longitudinal health data. As advancements continue in this burgeoning field, the promise of enhanced precision, early intervention, and improved overall health outcomes appear increasingly attainable for health domains and datasets that are traditionally "challenging" due to their inherent subjectivity, heterogeneity, and complexity.

Despite the initial successes described here, there are likely myriad digital health applications that have yet to be realized because they were not previously feasible prior to the advent of SSL. For example, recent advances in personalized SSL for emotion recognition [40] have the potential to improve the efficiency of the personalization of artificial intelligence–powered digital therapeutics for children with autism [41,42]. While the state-of-the-art of emotion recognition models hovers around 70% accuracy [43], previous emotion personalization efforts without self-supervision were able to achieve strong performances [44]. It is likely that further improvements with fewer labels will be possible with personalized SSL. This approach has yet to be applied to digital therapeutics more broadly, and this gap suggests the possibility of more precise digital therapeutics in the coming years.

## Ongoing Challenges

Personalized SSL studies can often be framed as several independent N=1 studies, where each study and corresponding model consists of training, validation, and testing data that all come from a single user. Such studies must be careful about overfitting across 2 dimensions: within subjects and between subjects. While between-subject overfitting, or overfitting to some patients while failing to generalize to other patients, is often discussed, discussions and evaluations of overfitting within a subject appear relatively sparse in the literature. Future work should explore overfitting in this temporal dimension more thoroughly.

Another ill-studied area is the intersection of performance discrepancies and personalization. Personalization of models should, in theory, lead to a reduction in ML performance discrepancies across groups. The capability of model personalization to reduce these discrepancies has yet to be thoroughly studied. However, it is plausible that personalized models could still propagate existing performance gaps across

groups if the underlying data remains skewed or if the personalization process disproportionately benefits certain groups [45]. A thorough understanding of this will require rigorous evaluation across a wide range of populations.

Another key challenge of personalized foundation models is that individuals change over time. As an extreme example to illustrate the point, a personalized model that was trained on an individual during their youth may be irrelevant during their 30s. The paradigm of continual (or online) learning, or the continual retraining of models as new data become available, can offer a solution. By allowing models to adapt incrementally, continuous learning can ensure that they evolve alongside the user, capturing shifts in behavior, preferences, and needs over time. Possible approaches can include incremental fine-tuning [46-48], where the model is periodically retrained on newly available data while retaining previously learned weights; experience replay [49,50], where a subset of past data is stored and combined with new data during model updates; and meta-learning [51,52], where the model learns how to quickly adapt to new data by leveraging prior knowledge, making it efficient in learning new tasks from fewer examples.

A final critical challenge is addressing human factors that influence the quality, consistency, and usability of patient-generated data in personalized SSL pipelines. As Slade et al [53,54] highlight, participants often encounter both technical and behavioral barriers during data collection, including device discomfort, app usability issues, and low perceived relevance of labeling tasks. These factors can lead to sporadic participant engagement, mislabeled or missing data, and dropout, ultimately undermining the effectiveness of models that rely on temporal consistency and high-volume personal data streams. Designing for human factors through mechanisms such as clearer feedback loops, improved incentives, and user-centered data collection interfaces will be essential to support robust protocol adherence leading to successful personalization.

## Conclusion

The training of personalized foundation models by learning from the vast unlabeled time series data that are often generated from patients can lead to ML applications in health care that expand beyond the traditional realm of diagnostics, such as adaptive and customized digital therapeutics. This area of research is relatively understudied in comparison to other aspects of ML-powered digital health, though it is likely that the advent and increasingly widespread application of SSL will lead to a proliferation of such applications.

XSL·FO

RenderX

## Authors' Contributions

Conceptualization: PW
Writing – original draft: PW
Writing – review and editing.: PW
Funding acquisition: PW

## Conflicts of Interest

None declared.

## References

1.  Beaulieu T, Knight R, Nolan S, Quick O, Ti L. Artificial intelligence interventions focused on opioid use disorders: a review of the gray literature. Am J Drug Alcohol Abuse. Jan 02, 2021;47(1):26-42. [doi: 10.1080/00952990.2020.1817466] [Medline: 33006905]
2.  Carreiro S, Chai PR, Carey J, Lai J, Smelson D, Boyer EW. mHealth for the detection and intervention in adolescent and young adult substance use disorder. Curr Addict Rep. Jun 2018;5(2):110-119. [FREE Full text] [doi: 10.1007/s40429-018-0192-0] [Medline: 30148037]
3.  Hsu M, Ahern DK, Suzuki J. Digital phenotyping to enhance substance use treatment during the COVID-19 pandemic. JMIR Ment Health. Oct 26, 2020;7(10):e21814. [FREE Full text] [doi: 10.2196/21814] [Medline: 33031044]
4.  Sun Y, Kargarandehkordi A, Slade C, Jaiswal A, Busch G, Guerrero A, et al. Personalized deep learning for substance use in Hawaii: protocol for a passive sensing and ecological momentary assessment study. JMIR Res Protoc. Feb 07, 2024;13:e46493. [FREE Full text] [doi: 10.2196/46493] [Medline: 38324375]
5.  Kargarandehkordi A, Slade C, Washington P. Personalized AI-driven real-time models to predict stress-induced blood pressure spikes using wearable devices: proposal for a prospective cohort study. JMIR Res Protoc. Mar 25, 2024;13:e55615. [FREE Full text] [doi: 10.2196/55615] [Medline: 38526539]
6.  Lee S, Kim H, Park MJ, Jeon HJ. Current advances in wearable devices and their sensors in patients with depression. Front Psychiatry. 2021;12:672347. [FREE Full text] [doi: 10.3389/fpsyt.2021.672347] [Medline: 34220580]
7.  Lee Y, Pham V, Zhang J, Chung TM. A digital therapeutics system for the diagnosis and management of depression: work in progress. 2023. Presented at: International Conference on Future Data and Security Engineering; November 22-24, 2023; Da Nang, Vietnam. [doi: 10.1007/978-981-99-8296-7_27]
8.  Pavlopoulos A, Rachiotis T, Maglogiannis I. An overview of tools and technologies for anxiety and depression management using AI. Appl Sci. Oct 08, 2024;14(19):9068. [doi: 10.3390/app14199068]
9.  Daniels J, Schwartz JN, Voss C, Haber N, Fazel A, Kline A, et al. Exploratory study examining the at-home feasibility of a wearable tool for social-affective learning in children with autism. NPJ Digit Med. 2018;1:32. [FREE Full text] [doi: 10.1038/s41746-018-0035-3] [Medline: 31304314]
10. Voss C, Washington P, Haber N, Kline A, Daniels J, Fazel A, et al. Superpower glass: delivering unobtrusive real-time social cues in wearable systems.? 2016. Presented at: 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing; September 12-16, 2016; Heidelberg, Germany. [doi: 10.1145/2968219.2968310]
11. Habehh H, Gohel S. Machine learning in healthcare. Curr Genomics. Dec 16, 2021;22(4):291-300. [FREE Full text] [doi: 10.2174/1389202922666210705124359] [Medline: 35273459]
12. Chowdhury A, Rosenthal J, Waring J, Umeton R. Applying self-supervised learning to medicine: review of the state of the art and medical implementations. Informatics. Sep 10, 2021;8(3):59. [doi: 10.3390/informatics8030059]
13. Rani V, Nabi ST, Kumar M, Mittal A, Kumar K. Self-supervised learning: a succinct review. Arch Comput Methods Eng. 2023;30(4):2761-2775. [FREE Full text] [doi: 10.1007/s11831-023-09884-2] [Medline: 36713767]
14. Spathis D, Perez-Pozuelo I, Marques-Fernandez L, Mascolo C. Breaking away from labels: the promise of self-supervised machine learning in intelligent health. Patterns. Feb 11, 2022;3(2):100410. [FREE Full text] [doi: 10.1016/j.patter.2021.100410] [Medline: 35199063]
15. Zhang J, Kowsari K, Harrison JH, Lobo JM, Barnes LE. Patient2Vec: a personalized interpretable deep representation of the longitudinal electronic health record. IEEE Access. 2018;6:65333-65346. [doi: 10.1109/access.2018.2875677]
16. Luu L, Pillai A, Lea H, Buendia R, Khan FM, Dennis G. Accurate step count with generalized and personalized deep learning on accelerometer data. Sensors (Basel). May 24, 2022;22(11):3989. [FREE Full text] [doi: 10.3390/s22113989] [Medline: 35684609]
17. Li J, Washington P. A comparison of personalized and generalized approaches to emotion recognition using consumer wearable devices: machine learning study. JMIR AI. May 10, 2024;3:e52171. [FREE Full text] [doi: 10.2196/52171] [Medline: 38875573]
18. Rudovic O, Tobis N, Kaltwang S, Schuller B, Rueckert D, Cohn JF, et al. Personalized federated deep learning for pain estimation from face images. ArXiv. Preprint posted online on January 12, 2021. [FREE Full text]
19. Krishnan R, Rajpurkar P, Topol EJ. Self-supervised learning in medicine and healthcare. Nat Biomed Eng. Dec 2022;6(12):1346-1352. [doi: 10.1038/s41551-022-00914-1] [Medline: 35953649]

20. Wu Y, Daoudi M, Amad A. Transformer-based self-supervised multimodal representation learning for wearable emotion recognition. IEEE Trans Affective Comput. Jan 2024;15(1):157-172. [doi: 10.1109/taffc.2023.3263907]

21. Weng D, Cheng M, Liu Z, Liu Q, Chen E. Diffusion auto-regressive transformer for effective self-supervised time series forecasting. ArXiv. Preprint posted online on October 8, 2024. [doi: 10.48550/arXiv.2410.05711]

22. Chen T, Kornblith S, Norouzi M, Hinton G. A simple framework for contrastive learning of visual representations. ArXiv. Preprint posted online on Febraury 13, 2020. [FREE Full text]

23. Liu Z, Alavi A, Li M, Zhang X. Self-supervised contrastive learning for medical time series: a systematic review. Sensors (Basel). Apr 23, 2023;23(9):4221. [FREE Full text] [doi: 10.3390/s23094221] [Medline: 37177423]

24. Raghu A, Chandak P, Alam R, Guttag J, Stultz CM. Sequential multi-dimensional self-supervised learning for clinical time series. 2023. Presented at: International Conference on Machine Learning; July 23-29, 2023; Honolulu, Hawaii.

25. He K, Chen X, Xie S, Li Y, Dollar P, Girshick R. ?Masked autoencoders are scalable vision learners. 2022. Presented at: IEEE/CVF Conference on Computer Vision and Pattern Recognition; June 18-24, 2022; New Orleans, LA. [doi: 10.1109/cvpr52688.2022.01553]

26. Deldari S, Spathis D, Malekzadeh M, Kawsar F, Salim FD, Mathur A. Crossl: cross-modal self-supervised learning for time-series through latent masking. 2024. Presented at: 17th ACM International Conference on Web Search and Data Mining; March 4-8, 2024; Merida, Mexico. [doi: 10.1145/3616855.3635795]

27. Tang P, Zhang X. Mtsmae: masked autoencoders for multivariate time-series forecasting. 2022. Presented at: 34th International Conference on Tools with Artificial Intelligence; October 31-November 2, 2022; Macao, China. [doi: 10.1109/ictai56018.2022.00150]

28. Li B, Sano A. Extraction and interpretation of deep autoencoder-based temporal features from wearables for forecasting personalized mood, health, and stress. Proc ACM Interact Mob Wearable Ubiquitous Technol. Jun 15, 2020;4(2):1-26. [doi: 10.1145/3397318]

29. Li H, Srinivasan D, Zhuo C, Cui Z, Gur RE, Gur RC, et al. Computing personalized brain functional networks from fMRI using self-supervised deep learning. Med Image Anal. Apr 2023;85:102756. [FREE Full text] [doi: 10.1016/j.media.2023.102756] [Medline: 36706636]

30. Spathis D, Perez-Pozuelo I, Brage S, Wareham NJ, Mascolo C. Self-supervised transfer learning of physiological representations from free-living wearable data. 2021. Presented at: Conference on Health, Inference, and Learning; April 8-10, 2021; Online. [doi: 10.1145/3450439.3451863]

31. Chen Y, Lo Y, Lai F, Huang C. Disease concept-embedding based on the self-supervised method for medical information extraction from electronic health records and disease retrieval: algorithm development and validation study. J Med Internet Res. Jan 27, 2021;23(1):e25113. [FREE Full text] [doi: 10.2196/25113] [Medline: 33502324]

32. Shurrab S, Duwairi R. Self-supervised learning methods and applications in medical imaging analysis: a survey. PeerJ Comput Sci. 2022;8:e1045. [FREE Full text] [doi: 10.7717/peerj-cs.1045] [Medline: 36091989]

33. Spathis D, Perez-Pozuelo I, Marques-Fernandez L, Mascolo C. Breaking away from labels: the promise of self-supervised machine learning in intelligent health. Patterns. Feb 11, 2022;3(2):100410. [FREE Full text] [doi: 10.1016/j.patter.2021.100410] [Medline: 35199063]

34. Zhao Q, Liu Z, Adeli E, Pohl KM. Longitudinal self-supervised learning. Med Image Anal. Jul 2021;71:102051. [FREE Full text] [doi: 10.1016/j.media.2021.102051] [Medline: 33882336]

35. Islam T, Washington P. Individualized stress mobile sensing using self-supervised pre-training. Appl Sci (Basel). Nov 2023;13(21):12035. [doi: 10.3390/app132112035] [Medline: 39507765]

36. Islam T, Peter W. Personalized prediction of recurrent stress events using self-supervised learning on multimodal time-series data. 2023. Presented at: International Conference on Machine Learning 2023 Workshop on Artificial Intelligence & Human Computer Interaction; July 23-29, 2023; Honolulu, HI.

37. Schmidt P, Reiss A, Duerichen R, Marberger C, Van Laerhoven K. Introducing WESAD, a multimodal dataset for wearable stress and affect detection. 2018. Presented at: 20th ACM International Conference on Multimodal Interaction; October 16-20, 2018; Boulder, CO. [doi: 10.1145/3242969.3242985]

38. Eom S, Eom S, Washington P. SIM-CNN: self-supervised individualized multimodal learning for stress prediction on nurses using biosignals. MedRXiv. Preprint posted online on August 28, 2023. [FREE Full text] [doi: 10.1101/2023.08.25.23294640]

39. Hosseini S, Gottumukkala R, Katragadda S, Bhupatiraju RT, Ashkar Z, Borst CW, et al. A multimodal sensor dataset for continuous stress detection of nurses in a hospital. Sci Data. Jun 01, 2022;9(1):255. [FREE Full text] [doi: 10.1038/s41597-022-01361-y] [Medline: 35650267]

40. Nimitsurachat P, Washington P. Audio-based emotion recognition using self-supervised learning on an engineered feature space. AI (Basel). Mar 2024;5(1):195-207. [FREE Full text] [doi: 10.3390/ai5010011] [Medline: 38715564]

41. Penev Y, Dunlap K, Husic A, Hou C, Washington P, Leblanc E, et al. A mobile game platform for improving social communication in children with autism: a feasibility study. Appl Clin Inform. Oct 2021;12(5):1030-1040. [FREE Full text] [doi: 10.1055/s-0041-1736626] [Medline: 34788890]

42. Voss C, Schwartz J, Daniels J, Kline A, Haber N, Washington P, et al. Effect of wearable digital intervention for improving socialization in children with autism spectrum disorder: a randomized clinical trial. JAMA Pediatr. May 01, 2019;173(5):446-454. [FREE Full text] [doi: 10.1001/jamapediatrics.2019.0285] [Medline: 30907929]

43. Washington P, Kalantarian H, Kent J, Husic A, Kline A, Leblanc E, et al. Improved digital therapy for developmental pediatrics using domain-specific artificial intelligence: machine learning study. JMIR Pediatr Parent. Apr 08, 2022;5(2):e26760. [FREE Full text] [doi: 10.2196/26760] [Medline: 35394438]

44. Kline A, Voss C, Washington P, Haber N, Schwartz H, Tariq Q, et al. Superpower glass. GetMobile Mobile Comp and Comm. Nov 14, 2019;23(2):35-38. [doi: 10.1145/3372300.3372308]

45. Warr M. Beat Bias? Personalization, bias, and generative AI. 2024. Presented at: Society for Information Technology & Teacher Education International Conference; Mar 25, 2024; Los Angelos, NV.

46. Rebuffi SA, Kolesnikov A, Sperl G, Lampert CH. icarl: incremental classifier and representation learning. 2017. Presented at: IEEE Conference on Computer Vision and Pattern Recognition; July 21-26, 2017; Honolulu, HI. [doi: 10.1109/cvpr.2017.587]

47. Rosenfeld A, Tsotsos JK. Incremental learning through deep adaptation. IEEE Trans Pattern Anal Mach Intell. Mar 1, 2020;42(3):651-663. [doi: 10.1109/tpami.2018.2884462]

48. Zhou Z, ShinShin J, Zhang L, Gurudu S, Gotway M, Liang J. Fine-tuning convolutional neural networks for biomedical image analysis: actively and incrementally. 2017. Presented at: IEEE Conference on Computer Vision and Pattern Recognition; July 21-26, 2017; Honolulu, HI. [doi: 10.1109/cvpr.2017.506]

49. Buzzega P, Boschini M, Porrello A, Calderara S. Rethinking experience replay: a bag of tricks for continual learning. 2021. Presented at: 25th International Conference on Pattern Recognition; January 10-15, 2021; Milan, Italy. [doi: 10.1109/icpr48806.2021.9412614]

50. Rolnick D, Ahuja A, Schwarz J, Lillicrap T, Wayne G. Experience replay for continual learning. 2019. Presented at: NeurIPS 2019; December 8-14, 2019; Vancouver, Canada.

51. Javed M, White M. Meta-learning representations for continual learning. 2019. Presented at: NeurIPS 2019; December 8-14, 2019; Vancouver, Canada.

52. Son J, Lee S, Kim G. When meta-learning meets online and continual learning: a survey. IEEE Trans Pattern Anal Mach Intell. Jan 2025;47(1):413-432. [doi: 10.1109/tpami.2024.3463709]

53. Slade C, Sun Y, Chao WC, Chen CC, Benzo RM, Washington P. Current challenges and opportunities in active and passive data collection for mobile health sensing: a scoping review. JAMIA Open. Aug 2025;8(4):ooaf025. [doi: 10.1093/jamiaopen/ooaf025] [Medline: 40688708]

54. Slade C, Benzo RM, Washington P. Design guidelines for improving mobile sensing data collection: prospective mixed methods study. J Med Internet Res. Nov 18, 2024;26:e55694. [FREE Full text] [doi: 10.2196/55694] [Medline: 39556828]

## Abbreviations

**ML:** machine learning
**SSL:** self-supervised learning