
Mixed Effects Random Forests for Personalised Predictions of Clinical Depression Severity

Rob Lewis¹ Asma Ghandeharioun^{*1} Szymon Fedor^{*1} Paola Pedrelli² Rosalind Picard¹ David Mischoulon²

Abstract

This work demonstrates how mixed effects random forests enable accurate predictions of depression severity using multimodal physiological and digital activity data collected from an 8-week study involving 31 patients with major depressive disorder. We show that mixed effects random forests outperform standard random forests and personal average baselines when predicting clinical Hamilton Depression Rating Scale scores (HDRS₁₇). Compared to the latter baseline, accuracy is significantly improved for each patient by an average of 0.199-0.276 in terms of mean absolute error ($p \ll 0.05$). This is noteworthy as these simple baselines frequently outperform machine learning methods in mental health prediction tasks. We suggest that this improved performance results from the ability of the mixed effects random forest to personalise model parameters to individuals in the dataset. However, we find that these improvements pertain exclusively to scenarios where labelled patient data are available to the model at training time. Investigating methods that improve accuracy when generalising to new patients is left as important future work.

1. Introduction

In recent years, smartphones and wearable technologies have become increasingly ubiquitous, while data processing and machine learning (ML) capabilities have matured. As a result, excitement has risen about a new digital era for mental health diagnosis, treatment, and prevention, where technology will be used to augment clinical workflows and to improve direct-to-consumer products such as digital therapeutics (Doraiswamy et al., 2019; Hsin et al., 2018).

Digital phenotyping refers to the collection of continuous,

^{*}Equal contribution ¹MIT Media Lab, Massachusetts Institute of Technology, Cambridge, MA, USA. ²The Depression Clinical and Research Program, Massachusetts General Hospital, Boston, MA, USA. Correspondence to: Robert Lewis <rob-lewis@media.mit.edu>.

multimodal, and in-situ data from patients using smartphone and wearable devices (Torous et al., 2016; Onnela & Rauch, 2016). It is considered central to digital mental health, as one can use ML on these data to enable the accurate forecasting and detailed understanding of mental health psychopathology (e.g., by discovering phenotypes that associate with mental health states). To give an example, a regressor could be learned to map clinical depression scores, such as the Hamilton Depression Rating Scale, HDRS (Hamilton, 1960), onto these digital data streams. Given such data can be collected passively and continuously, more frequent predictions of depression severity (via the HDRS score) can be made. This can provide substantial improvements to the patient’s care: for example, by enabling the early detection of treatment response or acute phase onset (e.g., relapse), care pathways can be adjusted expeditiously to improve patient outcomes (Huckvale et al., 2019).

However, enabling this vision is not trivial due to a multitude of factors, including the high degrees of heterogeneity both in the presentation of mental health conditions and in patients’ digital data streams. In such data scenarios, where multiple observations exist for each patient (i.e., *repeated measures* data) and these observations are not *independent and identically distributed* across patients (i.e., the data are *non-IID*), it becomes increasingly difficult to learn a single predictor that is accurate for all the patients in the cohort. Indeed, this performance degradation can be so severe that it has been observed that simple personal baselines – such as calculating the mean or median of that target variable over observations in the training set – often outperform state-of-the-art ML models (Demasi et al., 2017; Pedrelli et al., 2020). Such limitations are a significant barrier to using ML to augment mental health care and understanding.

In this work, we directly address the question of the utility of ML for HDRS depression severity prediction by comparing mixed effects random forests to simpler baselines. Mixed effects methods extend standard ML models to handle heterogeneous data. They do so by learning a subset of parameters for each individual: as such, they can be considered a form of model *personalisation*. Our contributions are twofold: first, we outline how we apply this technique to a *digital phenotyping* dataset that was collected in a clinical context and contains features related to known biomarkers of depression (e.g., quality of sleep). Second, we report

empirical results that suggest mixed effects methods significantly improve depression severity predictions over standard random forests and average personal baselines.

2. Related Work

2.1. Machine Learning for Mental Health Prediction

There is promising evidence on the feasibility of using ML to predict mental health states (e.g., clinical scores of depression) from physiological and behavioural data. Features have been engineered from various modalities, including self-report surveys, location patterns, smartphone usage (e.g., social media), electrodermal activity, accelerometer, and heart rate data (Grünerbl et al., 2014; De Choudhury et al., 2014; Saeb et al., 2015; Canzian & Musolesi, 2015; Suhara et al., 2017; Ghandeharioun et al., 2017; Pedrelli et al., 2020). However, much work remains. For example, more longitudinal studies are required to collect multimodal data streams and annotate them with clinical scores so high-quality data is available to train ML models, and more quantitative contributions are required to identify ML methods that reliably outperform personal baselines.

2.2. Mixed Effects Machine Learning

Mixed effects models are used in statistics and econometrics for longitudinal data, where *repeated measures* are collected from individuals in the system (Wu & Zhang, 2006; Fitzmaurice et al., 2012). They incorporate *random effect* parameters into models in addition to the *fixed effect* terms, which adjusts the model’s assumptions to account for heterogeneous data with multiple sources of random variability (e.g., both intra- and inter-individual). As a result, mixed effects methods allow stronger statistical conclusions to be made about the factors that correlate with the observed variance.

More recently, mixed effects methods have been considered in the context of ML (Hajjem et al., 2011; Sela & Simonoff, 2012; Ngufor et al., 2019), where they are included as a way to improve predictive accuracy. By granting the model the flexibility to learn some *random effect* parameters for each individual, predictions are *personalised* and, thus, accuracy is increased. This improvement has been noted in several contexts though, to the best of our knowledge, prior work has not yet studied if mixed effects methods can improve accuracy when predicting mental health severity using multimodal digital phenotyping data.

3. Methods

In this work we empirically assess the performance of a mixed effects random forest (MERF) method¹ on repeated

¹We explicitly acknowledge the authors of the MERF theory (Hajjem et al., 2014) and of the opensource Python implementation

measures HDRS scores. The method is referred to as *mixed effects* as it contains both *fixed effect* parameters – i.e., those that are shared by all *clusters*² in the dataset – and *random effect* parameters – i.e., those that are unique for each *cluster*. Beyond the random effect parameters, we are interested in the random forest component of this method (the *fixed effect*), given the random forest’s ability to maintain performance when there are many more features than observations (i.e., in the $p \gg n$ context) (Chen & Ishwaran, 2012). The MERF model (Hajjem et al., 2014) is defined by:

$$\mathbf{Y}_i = f(\mathbf{X}_i) + \mathbf{Z}_i \mathbf{b}_i + \boldsymbol{\epsilon}_i \quad (1)$$

$$\mathbf{b}_i \sim N(0, \mathbf{D}) \quad (2)$$

$$\boldsymbol{\epsilon}_i \sim N(0, \mathbf{R}_i) \quad (3)$$

$$\mathbf{V}_i = \text{Cov}(\mathbf{Y}_i) = \mathbf{Z}_i \mathbf{D} \mathbf{Z}_i^T + \mathbf{R}_i \quad (4)$$

$$\mathbf{R}_i = \boldsymbol{\sigma}^2 \mathbf{I}_{n_i} \quad (5)$$

Where $i=1, \dots, m$ are *clusters* (i.e., patients) with n_i observations each ($j=1, \dots, n_i$); \mathbf{Y}_i is the regression target variable ($n_i \times 1$); \mathbf{X}_i is a design matrix of input features ($n_i \times p$) and $f(\mathbf{X}_i)$ is the *fixed effect* random forest estimator; \mathbf{Z}_i is also a design matrix ($n_i \times q$), that usually contains a subset of features from \mathbf{X}_i ; \mathbf{b}_i are *random effect* parameters ($q \times 1$) for each i , and $\mathbf{Z}_i \mathbf{b}_i$ is assumed to be linear; $\boldsymbol{\epsilon}_i$ is the measurement error for each i ; and, \mathbf{D} , \mathbf{R}_i and \mathbf{V}_i are covariance matrices with assumptions that i) \mathbf{b}_i and $\boldsymbol{\epsilon}_i$ are independent of each other, and ii) between-cluster variation is the only source of correlation between the repeated measurements \mathbf{Y}_i , such that intra-subject measurement errors $\boldsymbol{\epsilon}_i$ are independent and \mathbf{R}_i is diagonal per (5).

While \mathbf{Z}_i may include many features, only a random intercept is used in the experiments of this paper. Thus, \mathbf{Z}_i becomes a ($n_i \times 1$) vector of ones, and so (1-5) can be expressed for each observation ij as:

$$Y_{ij} = f(X_{ij}) + b_i + \epsilon_{ij} \quad (6)$$

$$E(Y_{ij}|b_i) = f(X_{ij}) + b_i \quad (7)$$

$$E(Y_{ij}) = f(X_{ij}) \quad (8)$$

$E(Y_{ij}|b_i)$ is the *conditional expectation* of the model, and can be computed for clusters i that are known when the model is fit. $E(Y_{ij})$ is the *unconditional expectation* and it represents how the model predicts for a cluster i that is unknown at training time. These expectations are useful for interpreting model prediction errors for existing versus new users. The model parameters are fit using an expectation maximisation (EM) procedure with convergence monitored by a generalised log likelihood objective function. Further details of this training algorithm can be found in Supplementary Material (SM) Section A.

(MERF) whose work we build upon in this empirical paper.

²In this context, clusters refer to the individuals in the system that generate repeated measures data (e.g., patients) and they are specified *a priori* (i.e., it is assumed that the data from each patient forms a cluster and this does not change as the model is fit).

4. Experimental Work

4.1. Data

The dataset consists of 1,643 days of data collected from 31 patients with major depressive disorder (MDD), where MDD is categorised per the DSM-IV³ and patients have a score of ≥ 19 on the 28-item Hamilton Depression Rating Scale, HDRS₂₈ (Hamilton, 1960). Patients were monitored for 8 weeks and several categories of data were collected, with multiple observations per participant (i.e., the data has repeated measures). First, clinical assessments were performed by clinicians during 6 visits (once during screening followed by 5 bi-weekly visits during the 8-week monitoring period). These clinical scores include the HDRS₂₈, as well as a shorter 17-item HDRS scale, HDRS₁₇, which is commonly used to measure depressive symptom severity in clinical trials. Second, multimodal data was collected passively and continuously using mobile phones (MovisensXS, 2012) and physiological-sensor wristbands (Empatica E4, 2019). This data includes modalities known to associate with depressive symptomatology, including electrodermal activity (EDA) and heart rate variability (HRV), sleep characteristics, physical activity, digital activity (e.g., location and smartphone usage), and weather information. Features are created at various levels of temporal aggregation, including hourly and daily. Given the limited space, the feature design is discussed at length in Table 2 of SM Section B, as well as in our previous work (Pedrelli et al., 2020).

For the experiments, the dataset contains 2,820 features, and its rows are filtered to only include data captured on days with clinical scores, resulting in 149 observations in total⁴. Within observations, subsets of features are sometimes missing as a result of various events in the study period (e.g., sensor not worn for part of the day). These missing values are set to -1 before the model is trained.

4.2. Evaluation Scenarios

Three evaluation scenarios are considered. First, a *Random Split* setting with a train:test split of 70:30. Second, a *Time Split* setting, where the initial 3 observations per patient are used for training, and the remaining observations are used for testing. In practice, this scenario reflects making further predictions for existing patients (i.e., those that have already received clinical assessments). Finally, a *User Split* setting is considered, where observations for one patient are held out from the model as a testing set, while the observations

³The Diagnostic and Statistical Manual of Mental Disorders 4th edition (DSM-IV) provides a taxonomy for the classification of mental disorders (Bell, 1994). It is published by the American Psychiatric Association and is used in clinical practice in the USA.

⁴While 5 clinical scores are expected per participant during the observation period (excluding the initial screening visit), several participants had missing observations (6 are missing in total).

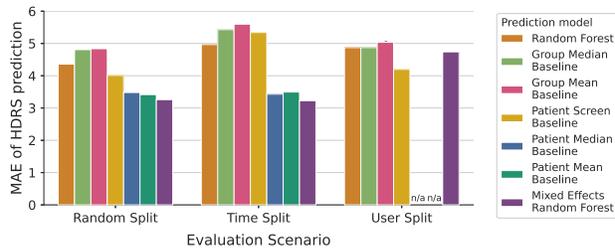


Figure 1. Mean absolute error (MAE) of testing set HDRS₁₇ prediction by evaluation scenario and model type. The MAE is reported here at the *group-level*: i.e., it is the average error across all observations in the testing set. The scenarios and models are described in Sections 4.2 & 4.3, respectively.

for the remaining patients are used in the training set. The sampling is then repeated 31 times so all patients are in the testing set exactly once. In practice, this scenario reflects making predictions for new patients.

4.3. Models Assessed

Several simple baselines are implemented: the median and the mean HDRS₁₇ score of patients in the training set (*Group Median* and *Group Mean*, respectively); the HDRS₁₇ score of the patient at the initial screening visit (*Patient Screen*); and, the median and mean HDRS₁₇ score of the patient from the training set observations (*Patient Median* and *Patient Mean*, respectively). Furthermore, to compare the mixed effects model to a standard ML approach, a random forest regressor without mixed effects is implemented (*Random Forest*)⁵. Finally, the mixed effects random forest regressor is implemented per (1-5) with a random effect term for the intercept (*Mixed Effects Random Forest*; *MERF*). In all cases, the raw HDRS₁₇ score is predicted (range 0-52) without any pre-processing adjustments.

4.4. Evaluation Metrics and Settings

The mean absolute error (MAE) is used to evaluate accuracy. It is calculated at two levels. First, in Figure 1 a single MAE error is calculated across all observations at the *group-level*. Second, in Table 1 a MAE error is calculated for each participant (i.e., the mean error using only predictions from that participant). This second level deepens the assessment of MERF versus the personal baselines. Informed by suggestions from the literature (Demasi et al., 2017), the MAE at the *participant-level* are used to derive a *user lift* metric, which represents the improvement of the MERF model over the baseline (e.g., if the baseline MAE is 4 and the MERF MAE is 3, then the *user lift* is 1). A corollary of this *participant-level* approach is that one can formally test if, on average, the *user lift* is significantly greater than zero. To

⁵The accuracies of other standard ML models are presented in SM Section D. These perform worse than the random forest and are excluded from the main body in the interest of brevity.

Table 1. *Participant-level* errors by scenario. The average of the mean absolute errors (MAE) for each participant is reported (Avg. Err.). The *user lift* metric represents the average performance increase of the *Mixed Effects Random Forest* (MERF) compared to the most accurate personal baseline (PBL). PBL is: the *Patient Mean* if *Random* scenario; the *Patient Median* if *Time Split* scenario; else, the *Patient Screen* score if *User Split* scenario. A one-sample one-tailed nonparametric permutation test ($p \leq 0.05$) is used to test if the *user lift* value is significantly greater than zero. The p-values from this test are aggregated over independent experiments (with $N=10$ random seed repeats) using Fisher’s method. The worst-case error metric (WC Err.) representing the worst error for any participant is also reported (see Section 4.4). **NB:** the permutation test is not calculated in the *User Split* scenario as the *user lift* is clearly less than zero.

SCENARIO	N. SEEDS	AVG. PBL ERR.	AVG. MERF ERR.	AVG. USER LIFT	USER LIFT P-VALUE	WC PBL ERR.	WC MERF ERR.
RANDOM SPLIT	10	3.349	3.165	0.199 (↑)	0.000	8.167	7.933
TIME SPLIT	10	3.450	3.174	0.276 (↑)	0.004	8.500	7.657
USER SPLIT	10	4.198	4.739	-0.541 (↓)	N/A	11.000	12.417

do so a one-sample one-tailed nonparametric permutation test is performed, from which the p-values are reported in Table 1. Finally, the *worst-case error* is also reported, which represents the worst MAE for any participant, and is thus a measure of model robustness in this respect. Experimental settings are discussed further in SM Section C.

4.5. Results

Figure 1 shows the MAE at the *group-level*. In two scenarios – the *Random Split* and *Time Split* – the mixed effects random forest shows an improvement over both the standard random forest and, more importantly, the patient median and mean baselines. However, it is also noteworthy that the mixed effects approach provides less of an improvement in the *User Split* scenario, and indeed using the patient HDRS₁₇ score at screening is a far better predictor in this setting.

Table 1 shows *participant-level* errors and *user lift* metrics. Improvement over the personal baselines is reflected at this level in the *Random Split* and *Time Split* scenarios, with the permutation tests suggesting the improvement is significant ($p \leq 0.05$). Moreover, the worst-case error is also slightly improved in these scenarios by MERF, suggesting it also helps to improve prediction robustness across individuals.

5. Discussion

These results suggest that a mixed effects approach allows random forests to significantly outperform baselines in HDRS₁₇ predictions when patients in the testing set have also contributed some data to the training set. Compared to the standard random forest, this benefit likely stems from the ability of the mixed effects model to fit a random effect intercept parameter for each patient (b_i). Such flexibility ensures the predictions of the learned model are not centred around a group average HDRS₁₇ score, but rather are adjusted by the average observed scores for each patient (cf. the conditional expectation (7)). That said, given the significant lift of MERF over personal average baselines, it is clear that the model learns more than just a *participant-level* intercept. Indeed, it is probable that this additional lift is due to the random forest fixed effect terms, $f(\mathbf{X}_i)$, that –

when estimated using the EM procedure (cf. SM Section A) – learn relations between the multimodal input features and the HDRS₁₇ scores (net of patient specific random effect values) that can be shared across patients.

Nevertheless, this method has limitations. Most notably, it does not provide a lift in the *User Split* scenario. In contrast to the *Time Split*, in the *User Split* scenario no random intercept parameter can be learned for the target user, and thus the model can only predict using the fixed effect parameters learned from the patients in the training set (cf. the unconditional expectation (8)). As such, its accuracy is similar to that of a standard random forest in this scenario.

As future work we will first introduce additional random effect parameters to assess if this further improves accuracy in the *Random Split* and *Time Split* scenarios. Second, we intend to further analyse the model’s errors on a patient-level, assessing if patient characteristics (e.g., the properties of their training data distribution) correlate with them. Third, we intend to repeat these analyses with alternative target variables and datasets (e.g., including individuals without a current MDD diagnosis) to further understand how this method generalises to related mental health prediction tasks. Fourth, as we do not currently use all days of data collected (i.e., days without clinical scores are excluded), we will consider ways to incorporate this information using time-lagged approaches. Fifth, we will attempt to improve model generalisation to new patients (i.e., the *User Split* scenario), e.g., by using patient characteristics to compute initial random effect parameter values for the new patients.

Finally, it is important to comment on the clinical significance of these results. While a lower MAE would be required to enable precision interventions, e.g., just-in-time-adaptive interventions (Nahum-Shani et al., 2018), an MAE ≈ 3.2 across patients is still useful for patient monitoring. For example, per the HDRS₁₇, if a patient’s score worsens from 0 (a level indicating recovery) to ≥ 15 (indicating relapse), then even with an MAE ≈ 3.2 a relapse prediction can still be made with reasonable confidence and used to alert the patient’s care team.

6. Conclusion

This work has shown that extending random forests with random effect intercept parameters significantly improves accuracy over personal baselines when predicting clinical HDRS₁₇ depression scores. These findings only apply to scenarios where patients have labelled data in the training set, but nevertheless are a noteworthy contribution to improving the utility of ML methods for digital mental health.

References

- Bell, C. C. DSM-IV: Diagnostic and Statistical Manual of Mental Disorders. *JAMA*, 272(10):828–829, 09 1994. ISSN 0098-7484. doi: 10.1001/jama.1994.03520100096046. URL <https://doi.org/10.1001/jama.1994.03520100096046>.
- Canzian, L. and Musolesi, M. Trajectories of depression: unobtrusive monitoring of depressive states by means of smartphone mobility traces analysis. In *Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing*, pp. 1293–1304, 2015.
- Chen, X. and Ishwaran, H. Random forests for genomic data analysis. *Genomics*, 99(6):323–329, 2012. ISSN 0888-7543. doi: <https://doi.org/10.1016/j.ygeno.2012.04.003>. URL <https://www.sciencedirect.com/science/article/pii/S0888754312000626>.
- DarkSky API. The DarkSky API. URL <https://darksky.net/about>. (Accessed: May 24, 2021).
- De Choudhury, M., Counts, S., Horvitz, E. J., and Hoff, A. Characterizing and predicting postpartum depression from shared facebook data. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work and Social Computing, CSCW '14*, pp. 626–638, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450325400. doi: 10.1145/2531602.2531675. URL <https://doi.org/10.1145/2531602.2531675>.
- Demasi, O., Kording, K., and Recht, B. Meaningless comparisons lead to false optimism in medical machine learning. *PLoS ONE*, 12, 2017.
- Doraiswamy, P., London, E., Varnum, P., Harvey, B., Saxena, S., Tottman, S., Campbell, S., Ibanez, A., Manji, H., Olama, M., Chou, I.-H., Herrman, H., Jeong, S.-j., Le, T., Montojo, C., Reve, B., Rommelfanger, K., Stix, C., Thakor, N., and Candeias, V. Empowering 8 billion minds: Enabling better mental health for all via the ethical adoption of technologies. *NAM Perspectives*, 10 2019. doi: 10.31478/201910b.
- Empatica E4. E4 wristband technical specifications, 2019. URL <https://support.empatica.com/hc/en-us/articles/202581999-E4-wristband-technical-specifications>. (Accessed: May 24, 2021).
- Fitzmaurice, G. M., Laird, N. M., and Ware, J. H. *Applied longitudinal analysis*, volume 998. John Wiley & Sons, 2012.
- Ghandeharioun, A., Fedor, S., Sangermano, L., Ionescu, D., Alpert, J., Dale, C., Sontag, D., and Picard, R. Objective assessment of depressive symptoms with machine learning and wearable sensors data. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, pp. 325–332, 2017. doi: 10.1109/ACII.2017.8273620.
- Grünerbl, A., Muaremi, A., Osmani, V., Bahle, G., Oehler, S., Tröster, G., Mayora, O., Haring, C., and Lukowicz, P. Smartphone-based recognition of states and state changes in bipolar disorder patients. *IEEE Journal of Biomedical and Health Informatics*, 19(1):140–148, 2014.
- Hajjem, A., Bellavance, F., and Larocque, D. Mixed effects regression trees for clustered data. *Statistics Probability Letters*, 81(4):451–459, 2011. ISSN 0167-7152. doi: <https://doi.org/10.1016/j.spl.2010.12.003>. URL <https://www.sciencedirect.com/science/article/pii/S0167715210003433>.
- Hajjem, A., Bellavance, F., and Larocque, D. Mixed-effects random forest for clustered data. *Journal of Statistical Computation and Simulation*, 84(6):1313–1328, 2014. doi: 10.1080/00949655.2012.741599. URL <https://doi.org/10.1080/00949655.2012.741599>.
- Hamilton, M. A rating scale for depression. *Journal of neurology, neurosurgery, and psychiatry*, 23(1):56, 1960.
- Hsin, H., Fromer, M., Peterson, B., Walter, C., Fleck, M., Campbell, A., Varghese, P., and Califf, R. Transforming Psychiatry into Data-Driven Medicine with Digital Measurement Tools. *npj Digital Medicine*, 1(1):1–4, 2018. ISSN 2398-6352. doi: 10.1038/s41746-018-0046-0. URL <http://dx.doi.org/10.1038/s41746-018-0046-0>.
- Huckvale, K., Venkatesh, S., and Christensen, H. Toward clinical digital phenotyping: a timely opportunity to consider purpose, quality, and safety. *npj Digital Medicine*, 2(1), 2019. ISSN 2398-6352. doi: 10.1038/s41746-019-0166-1. URL <http://dx.doi.org/10.1038/s41746-019-0166-1>.
- Koch, C., Wilhelm, M., Salzmann, S., Rief, W., and Euteneuer, F. A meta-analysis of heart rate variability in major depression. *Psychological Medicine*, 49(12): 1948–1957, 2019. doi: 10.1017/S0033291719001351.

- MERF. Mixed effects random forest (python), Manifold AI. URL <https://github.com/manifoldai/merf>. (Accessed: May 24, 2021).
- MovisensXS. Experience sampling for android, 2012. URL <https://xs.movisens.com>. (Accessed: May 24, 2021).
- Nahum-Shani, I., Smith, S. N., Spring, B., Collins, L., Witkiewitz, K., Tewari, A., and Murphy, S. Just-in-Time Adaptive Interventions (JITAI) in mobile health: Key components and design principles for ongoing health behavior support. *Annals of Behavioral Medicine: A Publication of the Society of Behavioral Medicine*, 52:446–462, 2018.
- Ngufor, C., Van Houten, H., Caffo, B. S., Shah, N. D., and McCoy, R. G. Mixed effect machine learning: A framework for predicting longitudinal change in hemoglobin a1c. *Journal of Biomedical Informatics*, 89:56–67, 2019. ISSN 1532-0464. doi: <https://doi.org/10.1016/j.jbi.2018.09.001>. URL <https://www.sciencedirect.com/science/article/pii/S1532046418301758>.
- Nutt, D., Wilson, S., and Paterson, L. Sleep disorders as core symptoms of depression. *Dialogues in clinical neuroscience*, 10:329–36, 02 2008.
- Onnela, J. P. and Rauch, S. L. Harnessing Smartphone-Based Digital Phenotyping to Enhance Behavioral and Mental Health. *Neuropsychopharmacology*, 41(7):1691–1696, 2016. ISSN 1740634X. doi: 10.1038/npp.2016.7.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Pedrelli, P., Fedor, S., Ghandeharioun, A., Howe, E., Ionescu, D. F., Bhathena, D., Fisher, L. B., Cusin, C., Nyer, M., Yeung, A., Sangermano, L., Mischoulon, D., Alpert, J. E., and Picard, R. W. Monitoring changes in depression severity using wearable and mobile sensors. *Frontiers in Psychiatry*, 11:1413, 2020. ISSN 1664-0640. doi: 10.3389/fpsy.2020.584711. URL <https://www.frontiersin.org/article/10.3389/fpsy.2020.584711>.
- Raschka, S. MLxtend: Providing machine learning and data science utilities and extensions to python’s scientific computing stack. *The Journal of Open Source Software*, 3(24), April 2018. doi: 10.21105/joss.00638. URL <http://joss.theoj.org/papers/10.21105/joss.00638>.
- Saeb, S., Zhang, M., Karr, C. J., Schueller, S. M., Corden, M. E., Kording, K. P., and Mohr, D. C. Mobile phone sensor correlates of depressive symptom severity in daily-life behavior: An exploratory study. *J Med Internet Res*, 17(7):e175, Jul 2015. ISSN 1438-8871. doi: 10.2196/jmir.4273. URL <http://www.jmir.org/2015/7/e175/>.
- Sano, A., Taylor, S., McHill, A. W., Phillips, A. J., Barger, L. K., Klerman, E., and Picard, R. Identifying objective physiological markers and modifiable behaviors for self-reported stress and mental health status using wearable sensors and mobile phones: Observational study. *J Med Internet Res*, 20(6):e210, Jun 2018. ISSN 1438-8871. doi: 10.2196/jmir.9410. URL <http://www.jmir.org/2018/6/e210/>.
- Sela, R. and Simonoff, J. RE-EM trees: A data mining approach for longitudinal and clustered data. *Machine Learning*, 86:169–207, 02 2012. doi: 10.1007/s10994-011-5258-3.
- Suhara, Y., Xu, Y., and Pentland, A. Deepmood: Forecasting depressed mood based on self-reported histories via recurrent neural networks. In *Proceedings of the 26th International Conference on World Wide Web*, pp. 715–724, 2017.
- Torous, J., Kiang, M., Lorme, J., and Onnela, J.-P. New tools for new research in psychiatry: A scalable and customizable platform to empower data driven smartphone research. *JMIR Mental Health*, 3:e16, 05 2016. doi: 10.2196/mental.5165.
- Wu, H. and Zhang, J. *Nonparametric regression methods for longitudinal data analysis*. 2006.

Supplementary Material: Mixed Effects Machine Learning for Personalised Predictions of Clinical Depression Severity

A. Mixed Effects Random Forest Expectation Maximisation Training Procedure

This section resumes the theoretical overview of the mixed effects random forest model introduced in Section 3. To estimate the unknown parameters of $f(\mathbf{X}_i)$ and \mathbf{b}_i in (1-5), maximum likelihood is used. Given a closed-form solution for parameter estimation does not exist, an expectation maximisation (EM) procedure is used to iteratively update the parameters to maximise a generalised log likelihood objective function (GLL). Given this estimation approach is not original to our work, we state the GLL without proof in (9-10). Likewise, we outline the EM procedure with the parameter update equations in Algorithm 1. For a rigorous theoretical overview, we refer the reader to a reference on the theory of mixed effects parameter estimation (Wu & Zhang, 2006) as well as to the work of the originators of the mixed effects random forest (MERF) approach (Hajjem et al., 2011; 2014).

$$\boldsymbol{\epsilon}_i = [\mathbf{Y}_i - f(\mathbf{X}_i) - \mathbf{Z}_i \mathbf{b}_i] \quad (9)$$

$$GLL(f, \mathbf{b}_i | \mathbf{Y}) = \sum_{i=1}^n \{ \boldsymbol{\epsilon}_i^T \mathbf{R}_i^{-1} \boldsymbol{\epsilon}_i + \mathbf{b}_i^T \mathbf{D}^{-1} \mathbf{b}_i + \log |\mathbf{D}| + \log |\mathbf{R}_i| \} \quad (10)$$

Algorithm 1 Fit Mixed Effects Random Forest with EM

System: MERF model defined by (1-5)

for $r = 0$ **to** $N_{iterations}$ **do**

E-step:

(i) Let $\mathbf{Y}_i^* = \mathbf{Y}_i - \mathbf{Z}_i \hat{\mathbf{b}}_{i(r-1)}$

(ii) Fit random forest to \mathbf{Y}_i^* to obtain $\hat{f}(\mathbf{X}_i)_{(r)}$

(iii) Fit $\hat{\mathbf{b}}_{i(r)} = \hat{\mathbf{D}}_{(r-1)} \mathbf{Z}_i^T \hat{\mathbf{V}}_{i(r-1)}^{-1} (\mathbf{Y}_i - \hat{f}(\mathbf{X}_i)_{(r)})$

Where: $\hat{\mathbf{V}}_{i(r-1)}$ is calculated by $\hat{\mathbf{D}}_{(r-1)}$ & $\hat{\mathbf{R}}_{i(r-1)}$ in (4)

M-step:

(i) Update: $\hat{\boldsymbol{\sigma}}_{(r)}^2 = N^{-1} \sum_{i=1}^n \{ \hat{\boldsymbol{\epsilon}}_{i(r)}^T \hat{\boldsymbol{\epsilon}}_{i(r)} + \hat{\boldsymbol{\sigma}}_{(r-1)}^2 [n_i - \hat{\boldsymbol{\sigma}}_{(r-1)}^2 \text{trace}(\hat{\mathbf{V}}_{i(r-1)})] \}$

And: $\hat{\mathbf{D}}_{(r)} = n^{-1} \sum_{i=1}^n \{ \hat{\mathbf{b}}_{i(r)} \hat{\mathbf{b}}_{i(r)}^T + [\hat{\mathbf{D}}_{(r-1)} - \hat{\mathbf{D}}_{(r-1)} \mathbf{Z}_i^T \hat{\mathbf{V}}_{i(r-1)}^{-1} \mathbf{Z}_i \hat{\mathbf{D}}_{(r-1)}] \}$

Where: $\hat{\boldsymbol{\epsilon}}_{i(r)} = \hat{\mathbf{Y}}_i - \hat{f}(\mathbf{X}_i) - \mathbf{Z}_i \hat{\mathbf{b}}_i$ (per (9))

end for

B. Summary of the Multimodal Dataset with Clinical Annotations

Table 2 summarises the data collected during the observational study and how it was engineered into features for the machine learning experiments. The feature designs were informed by prior work identifying biomarkers / correlates of depressive symptomatology. In total, the dataset contains 2,820 input features and a 1-dimensional HDRS₁₇ as the target variable. 1,643 rows of data are collected in total (each representing one day of patient observations). However, in the experiments of this paper, the rows are filtered to only include days where a clinical HDRS₁₇ score was reported. Thus, the total number of observations in the assessed dataset is 149. This corresponds to 5 clinical measures for 31 patients, minus 6 observations which are missing for data quality reasons.

C. Supplementary Information on the Experimental Settings

The experiments for each scenario were repeated 10 times with a different seed used for the random sampling scenario. This reduces the effect of data distribution and model training artefacts on the reported results. The hyperparameters of the standard random forest were tuned using a grid search, where values of the number of features, the max depth of trees, and the number of samples required to form splits and leaves are tuned. However, no hyperparameter tuning was performed on the mixed effects random forest (though this could be pursued as future work).

A one-sample one-tailed nonparametric permutation test was implemented to formally test the significance of the patient-level user lift values (Demasi et al., 2017). A permutation test allows one to calculate a *p-value* for a statistical test without requiring assumptions about the characteristics of the sampling distribution(s). It first calculates the test statistic on the observed sample(s). It then combines the samples, permutes them, and randomly samples from the combined sample. On each of these resamples, the test statistic is recalculated and compared to the observed test statistic (i.e., that from the original samples), allowing one to empirically calculate the probability – i.e. the *p-value* – of observing test statistic values at least as extreme as the observed value. One can then use this *p-value* with a pre-specified alpha (e.g., 0.05) to accept or reject the null

Table 2. Data were collected from several modalities and engineered into features using various aggregation techniques. Each row in the dataset represents the observations for a given day. 1,643 days of data are collected in total. The final model uses 2,820 features as input and a 1-dimensional HDRS score variable as its target.

Modality	Description
HDRS ₁₇ clinical score (target)	The 17-item Hamilton Depression Rating Scale (Hamilton, 1960) is a clinician-administered depression assessment scale. It was administered 6 times during the study: once during the initial screening visit, once at the beginning of the 8-week observation period, and then every other week for the remainder of the study. The regression target variable is created by summing the HDRS ₁₇ items to create a total score with a range of 0 to 52. The mean of the HDRS ₁₇ scores in this dataset is 17.7, the minimum is 5, and the max is 31.
Physiology: Electrodermal Activity (EDA)	Skin conductance level (SCL) and skin conductance response (SCR) are measured on the left and right wrists. Previous work has identified associations between skin conductance and stress / mental health (Sano et al., 2018). Various aggregated statistics are calculated on SCL & SCR (e.g., number of peaks and amplitude) at various temporal aggregations (e.g., hourly, daily, as well as aggregations to night, morning, etc.). Statistics are calculated for each wrist and for the difference between wrists.
Physiology: Heart Rate Variability (HRV)	HRV is measured on the left and right wrists. HRV is often found to associate with Major Depressive Disorder (Koch et al., 2019). Various HRV metrics are calculated in the time and frequency domains (e.g., AVNN, pNN50, rMSSD, SDANN, sDNNIDX, rrSDNN, PSD of the high-, low, very-low frequency signals, HF/LF ratio, etc.) at various temporal aggregations / periods in the day (e.g., hourly, daily, as well as aggregations to night, morning, etc.). Average statistics for heart rate (HR) are also calculated on both wrists at various temporal aggregations in the day.
Sleep	Sleep time is calculated (over 24 hours and during the night). Other sleep characteristics are calculated using actigraphy, such as sleep onset time, wakeups, maximal night uninterrupted sleep, and a sleep regularity index (Ghandeharioun et al., 2017). Sleep disturbance is recognised as a core symptom of depression (Nutt et al., 2008).
Motion / Physical Activity	Features for motion frequency (i.e., fraction of time in motion within a period) and magnitude (i.e., the intensity of the motion) were calculated at various temporal aggregations using accelerometer data collected from the left and right wrists.
Digital Activity (smartphone)	An Android app (MovisensXS, 2012) was used to collect smartphone activity data, including streams for location, call and messaging (sms) activity, and app usage and screen on / off time. Average statistics (mean, counts, sums, variance, etc.) were calculated for these streams at various temporal aggregations.
Environment (weather)	Features related to the weather were also calculated (e.g., temperature, precipitation, humidity, UV index, etc.). The DarkSky API (DarkSky API) was used to obtain historical weather information for each participant by using their location collected through the MovisensXS app.

hypothesis. To aggregate the p -values in Table 1 across the independent experiment repeats (i.e., different seeds) Fisher’s method was used.

It is also worth noting why the one-tailed permutation test is not performed in the *User Split* scenario in Table 1. As the test is one-tailed, it only assesses if the *user lift* is significantly greater than zero. For the *User Split* scenario, the

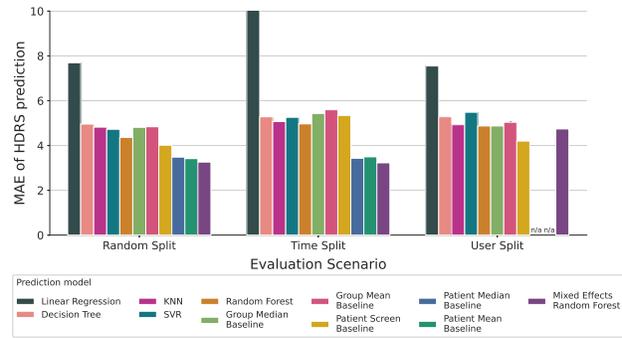


Figure 2. Extended results set including standard machine learning models as additional baseline comparisons. Mean absolute error (MAE) of testing set HDRS₁₇ prediction by evaluation scenario and model type. The MAE is reported here at the *group-level*: i.e., it is the average error across all observations in the testing set. The scenarios are described in Sections 4.2, and the extended set of models are described in Section D.

MERF error is clearly worse than the *Patient Screen* baseline and thus the *user lift* is greater than zero. The one-tailed permutation test is thus not performed, as it does not test this side of the test statistic distribution, and thus reporting a non-significant p -value here may seem incongruous with the magnitude of the negative user lift value.

The system is implemented in Python 3. Scikit-learn (Pedregosa et al., 2011) is used for the machine learning models, with the mixed effects model also using the *merf* Python library (*MERF*). The permutation tests were implemented using *MLxtend* (Raschka, 2018).

D. Supplementary Results: Comparison to Standard Machine Learning Baselines

To further evidence the benefits of the mixed effects random forest (*MERF*) model over standard baselines, Figure 2 displays an extended results set that includes the MAE of standard machine learning models, in addition to the mixed effects random forest and simple average baselines discussed in Section 4.3. The scenarios are consistent with those described in Section 4.2. Given the performance of all of these machine learning models are similar to (or worse) than the group level averages, they are excluded from the main body in the interest of brevity. However, we discuss their details here for completeness.

Hyperparameter tuning is performed on all of the machine learning models (with the exception of the mixed effects random forest) before the testing set MAE metric is generated. *Linear Regression* refers to a standard regression model with linear parameters, and different hyperparameters are assessed for the strength and nature of regularisation (i.e.,

L1 vs L2). Its notably poor performance may result from either the data setting – where there are many more features than observations (i.e., the $p \gg n$ context) – or its inability to model nonlinear relations between features, or both. *Decision Tree* refers to a standard decision tree regression, and hyperparameters of its max depth, as well as the number of samples required to form splits and leaves are tuned. *KNN* refers to a standard k-nearest neighbours regression, and the number of neighbours optimised in the hyperparameter tuning. Finally, *SVR* refers to a support vector regression model with a radial basis function kernel (RBF), and the kernel coefficient values (*gamma*) and L2 regularisation strength (*C*) are tuned. The models were implemented using Scikit-learn (Pedregosa et al., 2011).