

Subject-Adaptive Sparse Linear Models for Interpretable Personalized Health Prediction from Multimodal Lifelog Data

Dohyun Bu^{†*}, Jisoo Han^{†*}, Soohwa Kwon^{‡*}, Yulim So^{§*}, Jong-Seok Lee[†]

[†]Department of Industrial & Systems Engineering, KAIST, Daejeon, South Korea

[‡]Department of Systems Management Engineering, Sungkyunkwan University, Suwon, South Korea

[§]Department of Industrial Engineering, Sungkyunkwan University, Suwon, South Korea

Emails: {dohyun.bu, jongseok.lee}@kaist.ac.kr; {jshan1004, sooah0201, you0715}@skku.edu

Abstract—Improved prediction of personalized health outcomes—such as sleep quality and stress—from multimodal lifelog data could have meaningful clinical and practical implications. However, state-of-the-art models, primarily deep neural networks and gradient-boosted ensembles, sacrifice interpretability and fail to adequately address the significant inter-individual variability inherent in lifelog data. To overcome these challenges, we propose the Subject-Adaptive Sparse Linear (SASL) framework, an interpretable modeling approach explicitly designed for personalized health prediction. SASL integrates ordinary least squares regression with subject-specific interactions, systematically distinguishing global from individual-level effects. We employ an iterative backward feature elimination method based on nested F -tests to construct a sparse and statistically robust model. Additionally, recognizing that health outcomes often represent discretized versions of continuous processes, we develop a regression-then-thresholding approach specifically designed to maximize macro-averaged F1 scores for ordinal targets. For intrinsically challenging predictions, SASL selectively incorporates outputs from compact LightGBM models through confidence-based gating, enhancing accuracy without compromising interpretability. Evaluations conducted on the CH-2025 dataset—which comprises roughly 450 daily observations from ten subjects—demonstrate that the hybrid SASL-LightGBM framework achieves predictive performance comparable to that of sophisticated black-box methods, but with significantly fewer parameters and substantially greater transparency, thus providing clear and actionable insights for clinicians and practitioners.

I. INTRODUCTION

Sleep is fundamental to physical health, emotional stability, and overall quality of life [1]–[3]. Chronic sleep deprivation has been linked to severe health consequences, including hypertension, diabetes, obesity, depression, and heightened stress, which further exacerbates sleep issues, forming a detrimental cycle [4], [5]. Thus, accurately monitoring and predicting both daily sleep quality and fatigue could deliver substantial clinical and health benefits.

Recent advancements in wearable technology and smartphone-based sensing have enabled continuous, fine-

grained lifelogging, transforming everyday activities into high-resolution digital phenotypes [6]–[8]. These multimodal streams—including accelerometry, geolocation, app usage logs, and physiological signals—have shown significant promise in modeling and predicting emotional states, physical activities, sleep quality, and stress [9].

However, two key challenges remain: (1) **lack of interpretability** and (2) **inter-individual variability**. For lack of interpretability, state-of-the-art methods predominantly rely on deep neural networks and gradient-boosted ensembles [10]–[14]. Although these models can achieve strong predictive accuracy, their complexity often renders predictions *uninterpretable*, limiting practical utility in clinical settings where understanding *why* a prediction was made is crucial. For inter-individual variability, sensor–outcome relationships vary markedly across individuals; as a result, population-level models may miss subject-specific patterns, reducing accuracy and reliability in personalized contexts.

Furthermore, many health outcomes derived from lifelog data, such as sleep quality and fatigue levels, represent discretized forms of continuous states, naturally placing them within the scope of ordinal regression [15], [16]. Treating such outcomes as standard classification problems may ignore their underlying continuous structure, while conventional regression methods fail to align with discrete evaluation metrics such as macro-averaged F1 (macro-F1), thus motivating a hybrid modeling approach.

To address these challenges, we propose the **Subject-Adaptive Sparse Linear (SASL)** framework. SASL retains ordinary least squares regression (OLS) at its core, ensuring interpretability, while incorporating only minimal complexity to achieve competitive predictive performance. The proposed framework explicitly models inter-individual variability through subject-specific interactions, differentiating global effects from personalized signals. An iterative nested F -test procedure [17], designed to be statistically rigorous, ensures a parsimonious feature set by systematically removing statistically insignificant predictors.

Specifically, SASL employs a regression-then-thresholding

* These authors contributed equally to this work and are listed in alphabetical order.

strategy, first predicting continuous latent scores, and then deciding thresholds that maximize macro-F1 score, preserving the ordinal characteristics of the target variables. Moreover, recognizing that certain outcomes (e.g., sleep efficiency, S2) pose inherently greater modeling challenges, SASL selectively integrates predictions from a compact LightGBM [18] model, invoked via confidence-based gates, thereby improving overall accuracy without significantly compromising transparency.

Our key contributions are as follows:

- We present an interpretable, statistically rigorous, subject-adaptive linear modeling approach that leverages sparse and carefully selected predictors to distinguish global from subject-specific effects.
- We propose a regression-then-thresholding strategy for ordinal targets that is specifically aligned with macro-F1 evaluation.
- We develop a selective hybrid scheme that activates compact black-box components only for intrinsically challenging predictions, thereby balancing transparency with state-of-the-art accuracy.

The remainder of this paper is organized as follows. Section II reviews related work. Section III presents the proposed SASL framework. Section IV reports the empirical evaluation on the CH-2025 dataset [19] and the interpretability analyses. Finally, Section V concludes.

II. RELATED WORKS

A. Sensor-based Health Monitoring

Recent advances in mobile and wearable sensors have enabled continuous, unobtrusive monitoring of human behavior and physiology in natural settings. Multimodal lifelog data—comprising physical activity (e.g., accelerometer, step count), physiological signals (e.g., heart rate), environmental context (e.g., light, GPS), and mobile usage patterns—offer a rich foundation for modeling sleep, fatigue, and stress.

Prior studies have applied time-series models to capture temporal dependencies in lifelog data, addressing prediction tasks such as cognitive impairment, app churn, and stress detection [20]–[22]. Transformer-based architectures have demonstrated strong performance in modeling long-range dependencies and integrating heterogeneous modalities [23], [24]. In addition, gradient-boosted ensembles have been employed to model health indicators such as blood glucose and sleep quality from contextual features [25], [26]. CNN-based approaches have also been explored, converting sensor signals into structured or visual representations to enhance activity recognition and anomaly detection [27], [28]. However, many of these advances rely on black-box architectures, which can make interpretability challenging and may limit clinical trust, actionable insight, and subject-specific personalization.

B. Statistical Modeling

Statistical modeling has played a central role in health prediction tasks involving behavioral and physiological data, offering interpretable structures and well-defined inference procedures. Ordinary least squares (OLS) regression [29],

which models linear relationships by minimizing squared errors, has been widely used in tasks such as fatigue monitoring and sleep quality estimation, where model transparency and simplicity are essential [17]. To improve flexibility while retaining interpretability, recent studies have introduced extensions such as mixed-effects models and subject-specific interaction terms [15]. These methods allow statistical models to account for individual variation while maintaining a global structure.

In addition, various feature elimination strategies have been proposed to enhance statistical robustness in high-dimensional settings. Among them, backward elimination procedures based on nested F -tests have been used to construct sparse linear models with reliable explanatory power [30]. For prediction problems involving ordinal targets, regression-then-thresholding frameworks have also gained attention, as they align naturally with metrics such as the classification accuracy [31].

Taken together, these strands suggest that integrating expressive black-box models with sparse, subject-adaptive linear components and ordinal-aware decision rules may yield accurate, transparent predictions from lifelog data.

III. METHOD

Our proposed SASL pipeline consists of five primary stages: (III-A) data preprocessing and full model specification; (III-B) backward feature elimination; (III-C) time-aware evaluation and inference; (III-D) special treatment for S2; and (III-E) confidence-based post-processing. Collectively, they balance accuracy and transparency by coupling sparse, subject-adaptive OLS with leakage-aware validation, ordinal-aware thresholds, and a narrowly applied, high-confidence LightGBM assist.

A. Data Preprocessing and Full Model Specification

We first summarize each daily observation using interpretable statistical features. These include one-hot encoded weekdays (Mon, Tue, Wed, Thu, Fri, Sat, Sun), a numeric day-of-week index (`dow`: Monday=0 to Sunday=6), the ratio of screen-on time per day, calories expended per day, smartphone charging duration per day, and the proportions of each smartphone-detected activity type (e.g., driving, walking, running, stationary). All continuous variables are standardized for stable linear regression fitting.

To model inter-individual variability, subject identifiers (`id1` to `id10`) are encoded using one-hot vectors. We introduce interaction terms between subject IDs and each statistical feature, explicitly capturing subject-specific deviations. For instance, a single observation initially encoded as `{subject id, screen-on ratio}` is augmented to include `{id1, id2, screen-on ratio, id1 × screen-on ratio, id2 × screen-on ratio}`, allowing the regression model to estimate personalized coefficients for each individual. The resulting full model is equivalent to fitting separate linear models for each subject.

B. Backward Feature Elimination

To prune the inflated design matrix produced by the subject–interaction terms, we adopt an *iterative nested F-test* elimination scheme detailed in Algorithm 1. Here, $F_{\text{cdf}}(\cdot)$ is the cumulative distribution function of an F -distribution, $\text{cols}(\mathbf{X})$ returns the number of columns of matrix $\mathbf{X}$, and $\text{randChoice}(\mathcal{C}, \cdot)$ randomly selects one index from $\mathcal{C}$. The procedure begins with the full model and evaluates each candidate predictor by contrasting a reduced specification—obtained by omitting that single term—against the current model via the statistic

$$F = \frac{(\text{SSE}_{\text{reduced}} - \text{SSE}_{\text{full}})/1}{\text{SSE}_{\text{full}} / (N - p)}, \quad (1)$$

where SSE, N , and p denote the residual sum of squares, the number of observations, and the number of features in the full design matrix. The feature whose associated p -value exceeds the significance threshold α by the greatest margin is removed. The model is subsequently refitted, and the test is repeated. Ties—caused chiefly by multi-collinearity—are resolved using randomized priority, after which the resulting specification is vetted on validation data; the tie-breaking seed that maximizes ROC-AUC is retained. The loop stops once no term can be removed without violating the α criterion, yielding a statistically parsimonious predictor set $\mathbf{X}_{\text{reduced}}$ and an audit trail $\mathcal{H}$ of all deletion steps.

Algorithm 1: Backward Elimination via Nested F -Tests

Input: $\mathbf{X}_{\text{full}} \in \mathbb{R}^{N \times p}$, $\mathbf{y} \in \mathbb{R}^N$, α , s

Output: $\mathbf{X}_{\text{reduced}}$

```

1 Helper:  $\text{SSE}(\mathbf{X}, \mathbf{y})$  returns the SSE of the OLS fit on
    $\{\mathbf{x}_i, y_i\}_{i=1}^N$  where  $[\mathbf{x}_i; y_i] \in [\mathbf{X}; \mathbf{y}]$ .
2  $\mathbf{X} \leftarrow \mathbf{X}_{\text{full}}$ ,  $\mathcal{H} \leftarrow \emptyset$ ,  $t \leftarrow 0$ 
3 while  $\text{cols}(\mathbf{X}) > 1$  do
4    $\text{SSE}_{\text{full}} \leftarrow \text{SSE}(\mathbf{X}, \mathbf{y})$ 
5    $\text{df}_{\text{full}} \leftarrow N - \text{cols}(\mathbf{X})$ 
6   foreach predictor index  $j$  in  $\mathbf{X}$  do
7      $\mathbf{X}_{-j} \leftarrow \mathbf{X}$  without column  $j$ 
8      $\text{SSE}_{\text{reduced}}^{(j)} \leftarrow \text{SSE}(\mathbf{X}_{-j}, \mathbf{y})$ 
9      $F_j \leftarrow \frac{(\text{SSE}_{\text{reduced}}^{(j)} - \text{SSE}_{\text{full}})/1}{\text{SSE}_{\text{full}}/\text{df}_{\text{full}}}$ 
10     $\text{pval}_j \leftarrow 1 - F_{\text{cdf}}(F_j; 1, \text{df}_{\text{full}})$ 
11     $\mathcal{C} \leftarrow \{j : \text{pval}_j = \max_j \text{pval}_j\}$ 
12     $j^* \leftarrow \text{randChoice}(\mathcal{C}, \text{seed} = s)$ 
13    if  $\max_j \text{pval}_j > \alpha$  then
14       $t \leftarrow t + 1$ 
15       $\mathbf{X} \leftarrow \mathbf{X}_{-j^*}$ 
16    else
17      break
18 return  $\mathbf{X}_{\text{reduced}} \leftarrow \mathbf{X}$ 

```

Figure 1 visualizes how the backward F -test pares the full design matrix down to a sparse specification that keeps only statistically indispensable subject-specific terms. The simple global model, the unreduced full model, and the SASL fit are shown against the underlying ground-truth functions for the four synthetic subjects.




Fig. 1. Toy regression example after backward feature elimination. SASL (red solid) preserves a single subject-specific slope β_3 while discarding the remaining subject ID interactions, striking a compromise between the under-fitted global model $y = \mu + \beta x$ (brown dash-dot) and the over-parameterized full model (orange dashed dot). Here μ and β denote the *global* intercept and coefficient shared by all subjects, whereas μ_i and β_i are the additional intercept and slope for subject i . The term $x^{(i)}$ represents the interaction between the feature x and the one-hot indicator for subject i , meaning that it equals x for that subject and 0 for others.

C. Time-aware Evaluation and Inference

As sleep- and fatigue-related outcomes exhibit strong temporal continuity, random shuffling would introduce temporal leakage into training, biasing the evaluation and overstating performance. To avoid this, we employed a chronological two-fold cross-validation scheme. Specifically, each subject’s timeline was split at its median timestamp, yielding two segments: an early segment and a late segment. In fold-1, we trained on early segments and validated on late segments across all subjects; fold-2 reversed these roles. Thus, each fold preserved temporal ordering and subject balance, enabling a realistic assessment of model performance.

During backward feature elimination, multiple predictors occasionally exhibited equally non-significant effects, causing ties. To resolve such ties systematically, we introduced a priority-tuning step. We considered a set of random seeds $\mathcal{S}$, each defining a unique priority order for tie-breaking among predictors. For each seed $s \in \mathcal{S}$, we first applied backward feature elimination to the entire training set, obtaining a reduced feature subset $\mathbf{X}_{\text{reduced}}^{(s)}$. Next, for each chronological fold $f \in \{1, 2\}$, we fit an OLS model using the reduced subset, evaluated performance via ROC-AUC on the validation

split of the corresponding fold, and computed the seed-specific average performance:

$$\text{Score}(s) = \frac{1}{2} \left(\text{ROC_AUC}_{\text{fold-1}}^{(s)} + \text{ROC_AUC}_{\text{fold-2}}^{(s)} \right). \quad (2)$$

The final seed s^* was selected as

$$s^* = \arg \max_{s \in S} \text{Score}(s), \quad (3)$$

producing a deterministic and performance-optimized reduced feature set for subsequent modeling.

With the feature subset fixed by s^* , we then implemented a regression-then-thresholding strategy specifically aligned with the macro-F1 metric. For each candidate threshold τ (or threshold pair (τ_1, τ_2) for ternary targets), we separately computed the macro-F1 score on each chronological fold. Rather than directly selecting the threshold that maximizes the average macro-F1, we searched for a stable plateau region where both folds maintained consistently high macro-F1 scores. Formally, for a binary target, the discretization was defined as:

$$\hat{y} = \mathbb{1} [\hat{z} > \tau^*], \quad (4)$$

where $\hat{z}$ denotes the OLS prediction, and τ^* denotes the chosen robust threshold. For the ternary target (S1), the discretization was defined as:

$$\hat{y} = \mathbb{1} [\hat{z} > \tau_1^*] + \mathbb{1} [\hat{z} > \tau_2^*], \quad (5)$$

where thresholds τ_1^* , τ_2^* were likewise selected based on plateau stability.

Finally, a single OLS model was fit on the entire training set, using the finalized reduced feature set to estimate the regression coefficients $\hat{\beta}^{(s^*)}$. Test instances were scored as:

$$\hat{z} = \left(\hat{\beta}^{(s^*)} \right)^\top \mathbf{X}_{\text{test}}, \quad (6)$$

which were then discretized using the selected robust thresholds. These final predictions were then refined via a confidence-based gating step, incorporating predictions from a complementary LightGBM classifier, yielding the competition submission.

D. Special Treatment for S2 (Sleep Efficiency)

Unlike the other five targets, sleep efficiency (S2) captures nocturnal-only signals that are difficult to infer accurately from the coarse day-level summaries used in SASL. To address this inherent complexity, we developed a dedicated pipeline incorporating detailed temporal features and a robust modeling strategy specifically for S2. Minute-level sensor data were first re-indexed into consistent “analysis days” spanning from 16:00 to the following day’s 16:00 (KST) to align with individual sleep cycles. After appropriate domain-specific imputations, we computed 172 daily statistical aggregates, including interaction terms (e.g., *heart_rate* × *steps*), individual-specific baseline deviations, and event flags indicating nocturnal disturbances (*night_bright*). Additionally, subjects were grouped into 10 behavioral archetypes using k -means clustering on key routine features, which were then encoded as categorical features.

To manage feature dimensionality, we applied Recursive Feature Elimination (RFE) with a LightGBM classifier, reducing the input set to the top 30 informative features. The final model was trained via stratified 5-fold cross-validation repeated over multiple seeds (20 total folds), incorporating early stopping based on validation ROC-AUC scores to mitigate overfitting. Optimal classification thresholds maximizing macro-F1 scores were identified separately within each fold; the final threshold was set to the mean of the fold-specific values.

E. Confidence-based Post-processing

To complement the linear model’s interpretability with nuanced non-linear relationships, we implemented a confidence-based post-processing step that leverages predictions from a LightGBM classifier. Confidence scores were computed as the maximum predicted class probability output by the LightGBM model. When predictions from SASL and LightGBM models differed, we selectively adopted the LightGBM model’s prediction only if it demonstrated exceptionally high confidence (≥ 0.97 for all targets except S2, where ≥ 0.943 was used). This hybrid approach preserved overall stability and interpretability from the linear model while correcting select predictions where the LightGBM classifier demonstrated strong confidence, thereby modestly enhancing final predictive performance.

IV. EVALUATION

Experiments were conducted on the public **CH-2025** benchmark dataset [19], which consists of approximately 450 daily records, each summarized into a 2,160-dimensional feature vector across ten subjects. Following the time-aware evaluation approach described earlier (Section III-C), we constructed two chronological folds for each subject by splitting their timeline at the median timestamp. One fold trained on early segments and validated on later segments, while the other reversed this order. All model eliminations—including feature elimination, priority tuning, and threshold optimizations—were performed strictly within these folds, ensuring no leakage of temporal information into training. The final performance results are based on the competition’s held-out public test set, evaluated exactly once. Macro-F1 was computed on this public test split.

Table I summarizes the key results. The baseline SASL framework alone achieved a macro-F1 of 0.6258. Incorporating a dedicated LightGBM component for the most challenging target (S2, sleep efficiency) increased performance to 0.6311. A further improvement to 0.6387 was attained by selectively integrating high-confidence LightGBM model predictions for other targets as part of a confidence-based post-processing step. This final hybrid approach secured the fourth position among 1,032 participants on the competition’s public leaderboard.

The optimized decision thresholds identified through the plateau search (Section III-C) exhibited strong temporal robustness. Specifically, we adopted thresholds of $\tau_{Q1}=0.424$, $\tau_{Q2}=0.578$, $\tau_{Q3}=0.603$, $\tau_{S3}=0.650$ for binary outcomes, and $\tau_1=0.900$, $\tau_2=1.125$ for the ternary outcome S1.

ation, SASL effectively captured inter-individual variability while maintaining sparsity and transparency. Additionally, selectively integrating a compact LightGBM model through confidence-based gating further enhanced prediction accuracy without significantly compromising interpretability.

Nevertheless, avenues for future research remain. First, since the competition provided discretized labels, analyses using the original continuous or more finely graded ordinal labels could yield richer insights. Second, the current feature set predominantly consisted of daily summary statistics; future models incorporating detailed temporal and sequential dynamics could reveal deeper behavioral insights.

Despite these considerations, SASL exhibited notable strengths: achieving interpretability comparable to classical linear models, predictive accuracy on par with sophisticated black-box methods, and significantly fewer parameters. Achieving a top-1% placement on the public leaderboard further underscores the potential of interpretable, sparse modeling frameworks in personalized health prediction contexts.

ACKNOWLEDGMENT

This work was supported in part by the National Research Foundation of Korea (NRF) funded by the Korean Government (Ministry of Science and ICT (MSIT)) under Grant RS-2024-00361377, and in part by the Institute for Information and Communications Technology Planning & Evaluation (IITP) funded by the Korean Government (MSIT) under Grant RS-2025-02214591.

REFERENCES

- [1] M. Lee, H.-G. Kwak, H.-J. Kim, D.-O. Won, and S.-W. Lee, "Se-riessleepnet: an eeg time series model with partial data augmentation for automatic sleep stage scoring," *Frontiers in Physiology*, vol. 14, p. 1188678, 2023.
- [2] V. Clement-Carbonell, I. Portilla-Tamarit, M. Rubio-Aparicio, and J. J. Madrid-Valero, "Sleep quality, mental and physical health: a differential relationship," *International journal of environmental research and public health*, vol. 18, no. 2, p. 460, 2021.
- [3] N. Seighali, A. Abdollahi, A. Shafiee, M. J. Amini, M. M. Teymouri Athar, O. Safari, P. Faghfour, A. Eskandari, O. Rostaii, A. H. Salehi *et al.*, "The global prevalence of depression, anxiety, and sleep disorder among patients coping with post covid-19 syndrome (long covid): a systematic review and meta-analysis," *BMC psychiatry*, vol. 24, no. 1, p. 105, 2024.
- [4] V. L. Martire, D. Caruso, L. Palagini, G. Zoccoli, and S. Bastianini, "Stress & sleep: A relationship lasting a lifetime," *Neuroscience & Biobehavioral Reviews*, vol. 117, pp. 65–77, 2020.
- [5] R. Robillard, K. Dion, M.-H. Pennestri, E. Solomonova, E. Lee, M. Saad, A. Murkar, R. Godbout, J. D. Edwards, L. Quilty *et al.*, "Profiles of sleep changes during the covid-19 pandemic: Demographic, behavioural and psychological factors," *Journal of sleep research*, vol. 30, no. 1, p. e13231, 2021.
- [6] C. Gurrin, A. F. Smeaton, A. R. Doherty *et al.*, "Lifelogging: Personal big data," *Foundations and Trends® in information retrieval*, vol. 8, no. 1, pp. 1–125, 2014.
- [7] S. H. Jain, B. W. Powers, J. B. Hawkins, and J. S. Brownstein, "The digital phenotype," *Nature biotechnology*, vol. 33, no. 5, pp. 462–463, 2015.
- [8] R. Ribeiro, A. Trifan, and A. J. Neves, "Lifelog retrieval from daily digital data: narrative review," *JMIR mHealth and uHealth*, vol. 10, no. 5, p. e30517, 2022.
- [9] P. Chriskos, C. A. Frantzidis, C. M. Nday, P. T. Gkivogkli, P. D. Bamidis, and C. Kourtidou-Papadeli, "A review on current trends in automatic sleep staging through bio-signal recordings and future challenges," *Sleep medicine reviews*, vol. 55, p. 101377, 2021.
- [10] J. Lee, S. Yu, D. Kim, and K. Choi, "Multi-patching: Life-log classification with the reconstructed representation of multivariate time series," in *2024 15th International Conference on Information and Communication Technology Convergence (ICTC)*, 2024, pp. 798–803.
- [11] Y. Cho, N. Kwon, B. Yoon, and J. Kim, "Multi-modal sensor fusion for predicting sleep and stress patterns," in *2024 15th International Conference on Information and Communication Technology Convergence (ICTC)*, 2024, pp. 804–809.
- [12] J. Park, Y. Kim, J. Song, and H. Lee, "Contextual health state inference from lifelog data using llm," in *2024 15th International Conference on Information and Communication Technology Convergence (ICTC)*, 2024, pp. 816–821.
- [13] J. Kim, M. Ma, E. Choi, K. Cho, and C. Lee, "Tram: Enhancing user sleep prediction with transformer-based multivariate time series modeling and machine learning ensembles," in *2024 15th International Conference on Information and Communication Technology Convergence (ICTC)*, 2024, pp. 822–827.
- [14] Y. Na, S. Oh, S. Ko, and H. Lee, "Pixleepflow: A pixel-based lifelog framework for predicting sleep quality and stress level," in *2024 15th International Conference on Information and Communication Technology Convergence (ICTC)*, 2024, pp. 810–815.
- [15] T. Hothorn, T. Kneib, and P. Bühlmann, "Conditional transformation models," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 76, no. 1, pp. 3–27, 2014.
- [16] L. Kook, L. Herzog, T. Hothorn, O. Dürr, and B. Sick, "Deep and interpretable regression models for ordinal outcomes," *Pattern Recognition*, vol. 122, p. 108263, 2022.
- [17] A. Bommert, X. Sun, B. Bischl, J. Rahnenführer, and M. Lang, "Benchmark for filter methods for feature selection in high-dimensional classification data," *Computational Statistics & Data Analysis*, vol. 143, p. 106839, 2020.
- [18] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, "Lightgbm: A highly efficient gradient boosting decision tree," *Advances in neural information processing systems*, vol. 30, 2017.
- [19] ETRI, "ETRI Lifelog Dataset 2024," Online, 2025. [Online]. Available: <https://nanum.etri.re.kr/share/human/ETRILifelogDataset2024>
- [20] J. Kim and J. Lee, "Prediction of cognitive impairment using sleep lifelog data and lstm model," *Mathematics*, vol. 12, no. 20, p. 3208, 2024.
- [21] E. Elbasani and J.-D. Kim, "Life-log anomaly detection (llad) based on recurrent neural networks (rnn-lstm)," *Journal of Healthcare Engineering*, vol. 2021, p. 8829403, 2021.
- [22] H. Kwon, H. H. Kim, J. An, J.-H. Lee, and Y. R. Park, "Lifelog data-based prediction model of digital health care app churn using rnn-lstm," *Journal of Medical Internet Research*, vol. 23, no. 1, p. e22184, 2021.
- [23] D. Nguyen, T. Nguyen, L. Tran, H. Le, N. Nguyen, M. Do, and D. Nguyen, "Concept-based and embedding-based models in lifelog retrieval: An empirical comparison," *International Journal of Multimedia Information Retrieval*, vol. 14, no. 2, pp. 13–25, 2024.
- [24] S. Park, H. Kim, and J. Lee, "Tram: Enhancing user sleep prediction with transformer-based multivariate time series modeling and machine learning ensembles," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 5, 2025, pp. 4567–4574.
- [25] Y. Choi, M. Lee, and S. Kim, "Prediction of blood glucose using contextual lifelog data," in *Proceedings of the IEEE International Conference on Biomedical and Health Informatics*, 2023, pp. 123–128.
- [26] H. Jung, K. Park, and J. Choi, "A study on ml based sleep score model using lifelog data," *Healthcare Informatics Research*, vol. 30, no. 1, pp. 45–52, 2024.
- [27] K. Kim, J. Jang, H. Park, J. Jeong, and D. Shin, "Detecting abnormal behaviors in dementia patients using lifelog data: A machine learning approach," *Information*, vol. 14, no. 8, p. 433, 2023.
- [28] S. Lee, J. Kim, and H. Park, "Multi-modal lifelog data fusion for improved human activity recognition," in *Proceedings of the ACM International Conference on Multimedia Retrieval*, 2022, pp. 201–209.
- [29] C. R. Rao, "A treatise on ordinary least squares estimation of parameters of linear model," *Sankhya: The Indian Journal of Statistics, Series A*, vol. 35, no. 1, pp. 1–20, 1973.
- [30] C. Chen, S. T. Weiss, and Y.-Y. Liu, "Graph convolutional network-based feature selection for high-dimensional and low-sample size data," *Bioinformatics*, vol. 39, no. 4, p. btad135, 2023.
- [31] E. Frank and M. Hall, "A simple approach to ordinal classification," in *European conference on machine learning*. Springer, 2001, pp. 145–156.