LEARNING BRAIN-STATE DYNAMICS FOR OBJECTIVE ASSESSMENT OF STRESS RESILIENCE VIA EEG

Anonymous authors
Paper under double-blind review

ABSTRACT

Stress resilience is a relatively stable psychological trait that reflects an individual's capacity to maintain or rapidly recover adaptive functioning under adversity and sustained pressure. Although EEG signals offer a promising basis for objective resilience assessment beyond self-report measures, estimating stress resilience from EEG remains challenging because resilience is a long-term trait, whereas EEG predominantly reflects real-time, transient neural activity that is entangled with and obscures the slow trait signal. To address this, we propose KoMaNet, the Koopman Manifold Network. It suppresses nuisance at three levels to highlight the underlying trait. A self-supervised state network encodes each clip by imitating an expert descriptor, yielding a stable representation of the brain state. A Koopman dynamic network models the stable dynamics of brain states with a conditioned Koopman operator over the recording context. A manifold subject network classifies subjects by their dynamics on a manifold. We collect the RESIST dataset, comprising overnight-sleep and afternoon-nap paradigms, for stress resilience assessment. KoMaNet achieves AUCs of 0.740 and 0.700 on RESIST Night and RESIST Day, respectively. Model-attribution analyses identify discriminative patterns in occipital delta and frontal theta/alpha/beta bands. On RESIST Night, high-resilience subjects show smaller deviations from stationary dynamics than low-resilience subjects after sleep onset.

1 INTRODUCTION

Stress resilience refers to the capacity to maintain, or rapidly recover, adaptive psychological and behavioral functioning under adversity, uncertainty, and sustained pressure Kalisch et al. (2017). Individuals with higher resilience are generally better able to regulate affect, maintain cognitive flexibility, and act rationally during emergencies or prolonged stress exposure. This is particularly important in executive leadership, air traffic control, firefighting, emergency medicine, and other safety-critical professions, where impaired stress coping may lead to costly or even catastrophic consequences. In human resource management, developing reliable ways to assess stress resilience is therefore both a scientific challenge and a practical need.

Despite its importance, stress resilience is still primarily assessed through self-report questionnaires and psychometric scales, such as the 10-item Connor-Davidson Resilience Scale, known as CD-RISC-10, and related survey instruments Campbell-Sills & Stein (2007). These tools are widely used because they are simple and low-cost. However, such assessments suffer from intrinsic limitations that become especially problematic in real-world selection and monitoring scenarios. First, self-report measures depend heavily on introspection and honest disclosure, and are therefore vulnerable to response bias and deliberate faking Viswesvaran & Ones (1999). In competitive contexts such as recruitment, promotion, or certification for high-stakes occupations, individuals may intentionally exaggerate their stress tolerance and adaptive capacity in order to appear more suitable for a role Morgeson et al. (2007). Second, questionnaire scores often fail to capture an individual's condition in a timely and objective manner. People working in high-pressure professions are frequently evaluated repeatedly over long periods. Familiarity effects, habitual responding, and stable self-presentation strategies may make responses reflect a learned answering pattern or long-term self-image rather than the person's current resilience-related condition Morgeson et al. (2007). Cru-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

cially, these failures stem not from the resilience construct itself but from the conditions of measurement, where respondents are motivated to manage their answers. This motivates a complementary paradigm. Labels collected once in a non-evaluative research setting supervise an objective proxy in situations where self-report breaks down.

EEG signals provide a promising direction for such an objective assessment paradigm because they can be recorded continuously and capture aspects of neural activity that are difficult to obtain from self-report alone. Existing evidence supports this possibility. Prior EEG studies have linked resilience and related traits to frequency-band patterns, hemispheric asymmetry, and temporal flexibility in EEG patterns Curtis & Cicchetti (2007); Zhao et al. (2017); Paban et al. (2019); Kuper et al. (2019).

However, turning such correlates into a robust predictive model is challenging. A practical difficulty in studying this problem is the lack of publicly available EEG datasets with direct stress resilience annotations. To address this gap, we construct a self-collected dataset, RESIST, which comprises two sleep-EEG paradigms of the same resilience assessment, an overnight sleep subset called RESIST Night and an afternoon nap subset called RESIST Day. In RESIST, the questionnaires are administered in a controlled research setting, which serves as our benchmark for stress resilience assessment. In this setting, participants have no incentive to exaggerate their stress tolerance and adaptive capacity, and they are not repeatedly assessed with the same instrument.

Another challenge lies in the modeling itself. Task-specific EEG deep models have recently made strong progress on motor imagery Lawhern et al. (2018); Jiang et al. (2024b); Ingolfsson et al. (2020), emotion recognition Song et al. (2022), and sleep staging Supratak et al. (2017); Supratak & Guo (2020); Phan et al. (2019), and EEG foundation models Yang et al. (2026); Wang et al. (2024a); Jiang et al. (2024a); Wang et al. (2025); Zhou et al. (2025) further achieve cross-task general representations; yet stress resilience, a long-term trait, remains understudied. These models excel at recognizing discriminative patterns locked to a specific stimulus or state. Such transient activity fluctuates with events, whereas stress resilience is a long-term property that manifests in how brain-state dynamics deviate from and return to their stable evolution over long stretches of time. This aligns with the definition of resilience as the capacity to maintain, or rapidly recover, adaptive functioning under adversity. The difficulty is that this deviation-recovery regularity is buried among numerous transient events Delorme et al. (2007); Jung et al. (2000), and existing models can hardly isolate it. Consequently, a computational framework that explicitly references the stable dynamics of brain states, measures an individual's deviation and recovery, and aggregates them into a stable subject-level trait statistic remains underexplored.

To address this challenge, we propose KoMaNet, the Koopman Manifold Network, which comprises three modules. First, the Self-supervised State Network performs clip-level modeling. It learns a clip encoder that, from the current clip, predicts the expert state descriptor of its neighboring clips, yielding a stable representation of the locally sustained brain state. Second, the Koopman Dynamic Network models the dynamics across neighboring clips, inspired by the progress of deep learning with Koopman theory Zhang et al. (2025); Wu et al. (2025); Aswani & Jabari (2025); Zhang (2025); Xu et al. (2025). It captures how the state evolves in time, where the stable, persistent dynamics of each state are absorbed by a state-conditioned Koopman operator and the residual models an individual's deviation from this stable dynamics. Finally, the Manifold Subject Network performs subject-level modeling. It gathers a subject's dynamic representations over the whole time span, embeds both the stable dynamics and the deviation, and compares subjects on a manifold. Together, the three modules turn transient deviations into a stable statistic that discriminates the long-term trait.

We make three contributions. First, we introduce RESIST, which comprises 160 subjects and about 630 hours of sleep EEG. To the best of our knowledge, it is the first large-scale EEG dataset for stress resilience prediction and provides a benchmark for objective resilience decoding from EEG signals. Second, we propose KoMaNet, the first network for stress resilience prediction. It combines a self-supervised state network, a Koopman dynamic network, and a manifold subject network and outperforms state-of-the-art EEG models. Third, we validate KoMaNet on RESIST as the primary benchmark for stress resilience assessment. We reveal the spatial and temporal patterns of stress resilience and show that resilience can be decoded from both low-density overnight recordings and short-duration afternoon naps.

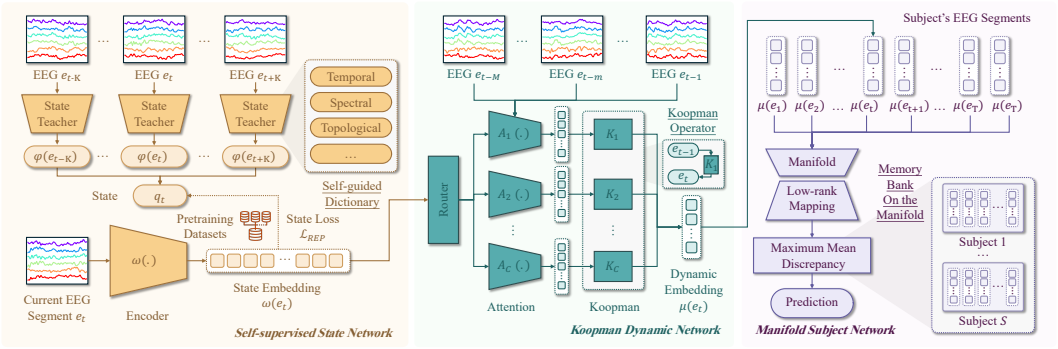


Figure 1: Overview of KoMaNet. The Self-supervised State Network encodes each clip into a stable representation of the locally sustained brain state. The Koopman Dynamic Network captures how the state evolves over time and models both its stable dynamics and an individual's deviation from them. The Manifold Subject Network represents each subject as a distribution of dynamic embeddings and compares subjects on a distribution manifold.

2 Related Work

Koopman theory studies nonlinear dynamical systems through a linear operator acting on observables. Deep Koopman models learn a higher-dimensional latent representation in which nonlinear state evolution can be approximated by linear dynamics. For example, deep autoencoders can learn compact representations of Koopman eigenfunctions and support the linearization of nonlinear dynamics (Lusch et al., 2018). SKOLR connects structured Koopman operators to linear recurrent updates for time-series forecasting (Zhang et al., 2025). K^2VAE combines a Koopman model with Kalman filtering to reduce error accumulation in long-horizon predictions (Wu et al., 2025).

Recently, Koopman theory has been used to understand latent representations across domains. Koopman Autoencoders model how neural network representations evolve through layers as dynamical systems and reveal progressive topological simplification (Aswani & Jabari, 2025). For multi-domain settings, shared-embedding Koopman frameworks use Procrustes analysis to align the spectra of domain-specific operators and transfer knowledge across heterogeneous time series (Zhang, 2025). ResKoopNet directly minimizes spectral residuals to obtain more complete Koopman spectral information and address the spectral inclusion problem of earlier dictionary-based approaches (Xu et al., 2025). Deep Koopman models are well suited to studying non-stationary dynamics and long-horizon temporal dependencies, yet their potential for stress resilience prediction remains unexplored.

3 Methodology

We propose KoMaNet, a framework for subject-level stress resilience prediction from EEG. As illustrated in Fig. 1, KoMaNet comprises three modules, the Self-supervised State Network, the Koopman Dynamic Network, and the Manifold Subject Network.

3.1 Self-supervised State Network

The Self-supervised State Network learns a clip encoder that maps each raw EEG clip to a compact, stable state embedding.

Let ω be the encoder, which maps a clip $\mathbf{e} \in \mathbb{R}^{S\times L}$ with S channels and length L to a d-dimensional state embedding,

$$\omega : \mathbb{R}^{S\times L} \to \mathbb{R}^d, \qquad \mathbf{e} \mapsto \omega(\mathbf{e}). \tag{1}$$

We train the encoder on an unlabeled pretraining set. Let $\mathcal{U} = \{\mathbf{U}_i\}_{i=1}^{N_u}$ be the unlabeled pretraining set, where each recording $\mathbf{U}_i = (\mathbf{e}_{i,1}, \ldots, \mathbf{e}_{i,T_i})$ is an ordered sequence of clips.

To model the locally sustained brain state and supervise this embedding without labels, we imitate a fixed teacher map

$$\varphi : \mathbb{R}^{S\times L} \to \mathbb{R}^m. \tag{2}$$

The teacher map summarizes the physiological state of a clip through a dictionary of features, spanning per-channel spectral power and inter-channel correlations.

Because a resilience trait is slow and stable, we let the embedding imitate the average of the descriptor over the following K clips,

$$\mathbf{q}_{i,t} = \frac{1}{K}\sum_{j=0}^{K-1} \varphi(\mathbf{e}_{i,t+j}), \tag{3}$$

and a shared decoder $D_\omega : \mathbb{R}^d \to \mathbb{R}^m$ reconstructs this target from the embedding.

We then train the encoder with a weighted state-reconstruction loss

$$\mathcal{L}_{\text{rep}} = \mathbb{E}_{(i,t)}\big\| D_\omega(\omega(\mathbf{e}_{i,t})) - \mathbf{q}_{i,t}\big\|_\alpha^2, \tag{4}$$

$$\|\mathbf{v}\|_\alpha^2 = \sum_{r=1}^{m} \alpha_r v_r^2, \tag{5}$$

where α_r weights descriptor dimension r to down-weight unstable channels. Averaging over the window suppresses isolated transients that do not persist in the average, giving clip-level robustness.

After pretraining, we freeze ω. Every task clip $\mathbf{e}_{n,t}$ is then represented by the state embedding $\omega(\mathbf{e}_{n,t})$ and passed to the Koopman Dynamic Network.

3.2 Koopman Dynamic Network

The Koopman Dynamic Network models the temporal evolution of the state embeddings. Let $\mathcal{D} = \{(\mathbf{X}_n, y_n)\}_{n=1}^N$ be the task dataset, where $\mathbf{X}_n = (\mathbf{e}_{n,1}, \ldots, \mathbf{e}_{n,T_n})$ is the clip sequence of subject n and y_n is the resilience label of subject n. The encoder turns subject n into an embedding sequence $(\omega(\mathbf{e}_{n,1}), \ldots, \omega(\mathbf{e}_{n,T_n}))$.

Different brain states follow different dynamics. We therefore partition each recording into several conditions and learn one attention module and one Koopman operator per condition. A fixed stage-partition function π labels each clip with a discrete control state,

$$c_{n,t} = \pi(\mathbf{e}_{n,t}) \in \{1, \ldots, C\}, \tag{6}$$

indexing the coarse brain state. Conditioning on $c_{n,t}$ lets the model compare state deviation and recovery within the same brain state rather than confounding them across brain states.

To supply temporal context, an attention module A_c summarizes the preceding window of M embeddings into a context state

$$\mathbf{h}_{n,t-1} = A_{c_{n,t}}\big(\omega(\mathbf{e}_{n,t-M}), \ldots, \omega(\mathbf{e}_{n,t-1})\big). \tag{7}$$

We then route each clip to a state-specific operator. Keeping a bank of C Koopman operators $\{\mathbf{K}_c\}_{c=1}^C$ with $\mathbf{K}_c \in \mathbb{R}^{d\times d}$, a hard router selects the operator matching the current control state,

$$\mathbf{K}_{n,t} = \mathbf{K}_{c_{n,t}}. \tag{8}$$

Each operator captures the typical one-step linear evolution of the embedding under the corresponding brain state and is spectrally stabilized,

$$\mathbf{K}_c = \rho\,\frac{\mathbf{W}_c}{\sigma_{\max}(\mathbf{W}_c)}, \qquad \rho(\mathbf{K}_c) \le \rho \le 1, \tag{9}$$

where $\rho(\cdot)$ is the spectral radius and $\sigma_{\max}$ is the spectral norm, so that no mode is amplified and near-unit modes can persist.

The selected operator then advances the context state to predict the next embedding, $\hat{\omega}(\mathbf{e}_{n,t}) = \mathbf{K}_{c_{n,t}}\mathbf{h}_{n,t-1}$. The transition residual

$$\mathbf{r}_{n,t} = \omega(\mathbf{e}_{n,t}) - \hat{\omega}(\mathbf{e}_{n,t}) \tag{10}$$

measures how far the actual dynamics deviate from the linear evolution typical of state $c_{n,t}$.

Finally, we assemble a per-transition dynamic embedding $\mu(\mathbf{e}_{n,t}) = [\,\omega(\mathbf{e}_{n,t});\ \mathbf{r}_{n,t}\,]$, which carries both the state embedding and the state-conditioned dynamics.

At training time, we minimize a one-step prediction loss

$$\mathcal{L}_{\text{dyn}} = \mathbb{E}_{(n,t)}\, \ell_\delta\big(\hat{\omega}(\mathbf{e}_{n,t}),\, \omega(\mathbf{e}_{n,t})\big), \tag{11}$$

where ℓ_δ is the Huber loss with threshold δ,

$$\ell_\delta(\mathbf{u}, \mathbf{v}) = \begin{cases} \frac{1}{2}\|\mathbf{u}-\mathbf{v}\|^2, & \|\mathbf{u}-\mathbf{v}\| \le \delta, \\ \delta\left(\|\mathbf{u}-\mathbf{v}\| - \frac{\delta}{2}\right), & \|\mathbf{u}-\mathbf{v}\| > \delta, \end{cases} \tag{12}$$

3.3 Manifold Subject Network

The Manifold Subject Network represents each subject as the empirical distribution of the subject's dynamic embeddings and classifies subjects on a distribution manifold.

We first represent each subject as a distribution, the empirical distribution of its dynamic embeddings $\mu(\mathbf{e}_{n,t})$ over a set $\mathcal{S}_n$ of transitions,

$$\widehat{P}_n = \frac{1}{|\mathcal{S}_n|} \sum_{t \in \mathcal{S}_n} \delta_{\mu(\mathbf{e}_{n,t})}. \tag{13}$$

We map each dynamic embedding $\mathbf{v}$ through a learnable low-rank projection

$$\psi(\mathbf{v}) = \mathbf{W}\mathbf{v}, \qquad \mathbf{W} \in \mathbb{R}^{r \times D}. \tag{14}$$

We then measure the discrepancy between two subject distributions. With an RBF kernel $\kappa(\mathbf{a}, \mathbf{b}) = \exp(-\gamma\|\psi(\mathbf{a}) - \psi(\mathbf{b})\|_2^2)$,

$$\begin{aligned} \mathrm{MMD}^2(\widehat{P}_n, \widehat{P}_r) = \mathbb{E}_{\mathbf{a},\mathbf{a}' \sim \widehat{P}_n} \kappa(\mathbf{a}, \mathbf{a}') + \mathbb{E}_{\mathbf{b},\mathbf{b}' \sim \widehat{P}_r} \kappa(\mathbf{b}, \mathbf{b}') \\ - 2\, \mathbb{E}_{\mathbf{a} \sim \widehat{P}_n, \mathbf{b} \sim \widehat{P}_r} \kappa(\mathbf{a}, \mathbf{b}). \end{aligned} \tag{15}$$

Finally, storing the training subjects $\mathcal{T}$ as landmarks, we place any subject n by its similarity to each landmark,

$$\mathbf{a}_n = \left[\, \exp\{-\eta\, \mathrm{MMD}^2(\widehat{P}_n, \widehat{P}_r)\} \right]_{r \in \mathcal{T}} \in \mathbb{R}^{|\mathcal{T}|}. \tag{16}$$

For the high-versus-low resilience classification, a linear map g yields the subject-level prediction $\hat{p}_n = \sigma(g(\mathbf{a}_n))$. We fit the network with a class-balanced binary cross-entropy loss

$$\mathcal{L}_{\text{cls}} = -\sum_{n \in \mathcal{T}} w_{y_n} \big[y_n \log \hat{p}_n + (1 - y_n) \log(1 - \hat{p}_n) \big], \tag{17}$$

where the weights w_{y_n} are inversely proportional to class frequency.

4 Experiments

4.1 Benchmark Dataset

RESIST comprises 160 subjects in two sleep-EEG cohorts. Both RESIST Night and RESIST Day use the 10-item Connor-Davidson Resilience Scale, known as CD-RISC-10, with scores ranging from 0 to 40 Campbell-Sills & Stein (2007). Figure 2 shows the score distributions.

The RESIST Night subset comprises 100 subjects, including 46 males and 54 females with a mean age of 29.5 years ranging from 21 to 45. Overnight sleep EEG signals were recorded with a Nox-A1 ambulatory polysomnography system from Nox Medical Inc. in Reykjavik, Iceland. We used six electrodes, F3, F4, C3, C4, O1, and O2, placed according to the international 10-20 system. For each subject, we analyze a six-hour segment. The recordings were notch filtered at 50 Hz to suppress power-line interference, band-pass filtered at 0.5-30 Hz, resampled to 100 Hz, and z-score

normalized per clip Wang et al. (2024b). As supervision signals, we used questionnaire-derived resilience scores based on the CD-RISC-10. We binarize the continuous CD-RISC-10 score at a threshold of 20.

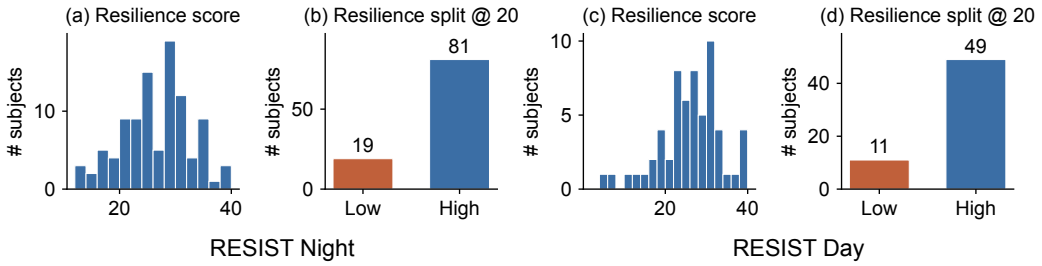


Figure 2: The figure shows the subject distributions for RESIST Night, which includes 100 subjects, and RESIST Day, which includes 60 afternoon nap subjects.

Table 1: We compare KoMaNet with task-specific, foundation, and Koopman-based EEG models on RESIST Night and RESIST Day. For each metric, the best result is shown in **bold**, and the second-best result is underlined.

	RESIST Night				RESIST Day			
Method	BAcc↑	AUC↑	NLL↓	Cov@50↑	BAcc↑	AUC↑	NLL↓	Cov@50↑
TCNet	$0.597_{\pm 0.052}$	$0.585_{\pm 0.075}$	3.72	0.32	$0.556_{\pm 0.086}$	$0.545_{\pm 0.070}$	3.77	0.33
LMDA	$0.623_{\pm 0.030}$	$0.617_{\pm 0.027}$	3.60	0.33	$0.578_{\pm 0.041}$	$0.556_{\pm 0.046}$	3.68	0.30
TSCeption	$\underline{0.664_{\pm 0.018}}$	$0.648_{\pm 0.031}$	<u>3.42</u>	<u>0.38</u>	$\underline{0.592_{\pm 0.014}}$	$0.575_{\pm 0.074}$	<u>3.50</u>	<u>0.34</u>
CSPNet	$0.641_{\pm 0.034}$	$0.628_{\pm 0.034}$	3.52	0.35	$0.585_{\pm 0.025}$	$\underline{0.596_{\pm 0.061}}$	3.56	0.33
TSLANet	$0.584_{\pm 0.013}$	$0.620_{\pm 0.016}$	3.64	0.30	$0.566_{\pm 0.017}$	$0.548_{\pm 0.083}$	3.62	0.29
Conformer	$0.623_{\pm 0.017}$	$0.615_{\pm 0.016}$	3.62	0.33	$0.588_{\pm 0.045}$	$0.590_{\pm 0.109}$	3.54	0.32
EEGNet	$0.658_{\pm 0.024}$	$0.655_{\pm 0.041}$	3.48	0.36	$0.583_{\pm 0.060}$	$0.582_{\pm 0.096}$	3.58	0.31
BIOT	$0.586_{\pm 0.030}$	$0.611_{\pm 0.050}$	3.75	0.29	$0.512_{\pm 0.011}$	$0.515_{\pm 0.020}$	3.86	0.29
EEGPT	$0.594_{\pm 0.028}$	$0.640_{\pm 0.055}$	3.55	0.34	$0.510_{\pm 0.011}$	$0.518_{\pm 0.020}$	3.72	0.30
LaBraM	$0.598_{\pm 0.040}$	$0.615_{\pm 0.060}$	3.70	0.28	$0.518_{\pm 0.026}$	$0.520_{\pm 0.030}$	3.88	0.28
CBraMod	$0.605_{\pm 0.045}$	$0.636_{\pm 0.055}$	3.58	0.31	$0.515_{\pm 0.026}$	$0.517_{\pm 0.030}$	3.83	0.29
CSBrain	$0.590_{\pm 0.050}$	$0.605_{\pm 0.058}$	3.72	0.30	$0.510_{\pm 0.020}$	$0.512_{\pm 0.030}$	3.90	0.30
Ko-PatchTST	$0.586_{\pm 0.076}$	$0.612_{\pm 0.063}$	3.95	0.24	$0.531_{\pm 0.045}$	$0.528_{\pm 0.083}$	3.95	0.24
Ko-Informer	$0.579_{\pm 0.110}$	$0.601_{\pm 0.088}$	4.05	0.22	$0.524_{\pm 0.045}$	$0.518_{\pm 0.073}$	4.10	0.22
Ko-Autoformer	$0.588_{\pm 0.071}$	$0.607_{\pm 0.116}$	3.90	0.25	$0.535_{\pm 0.022}$	$0.532_{\pm 0.088}$	3.92	0.25
DeepKoopman	$0.610_{\pm 0.136}$	$0.634_{\pm 0.112}$	3.66	0.30	$0.556_{\pm 0.039}$	$0.548_{\pm 0.067}$	3.62	0.28
Koopa	$0.618_{\pm 0.146}$	$\underline{0.670_{\pm 0.088}}$	3.60	0.32	$0.548_{\pm 0.081}$	$0.560_{\pm 0.092}$	3.75	0.27
ResKoopNet	$0.601_{\pm 0.135}$	$0.638_{\pm 0.085}$	3.70	0.28	$0.540_{\pm 0.012}$	$0.545_{\pm 0.058}$	3.80	0.26
SKOLR	$0.583_{\pm 0.066}$	$0.605_{\pm 0.069}$	3.98	0.23	$0.528_{\pm 0.038}$	$0.522_{\pm 0.068}$	3.98	0.24
KoMaNet	$\mathbf{0.724}_{\pm 0.018}$	$\mathbf{0.740}_{\pm 0.019}$	**3.25**	**0.58**	$\mathbf{0.635}_{\pm 0.020}$	$\mathbf{0.700}_{\pm 0.021}$	**3.36**	**0.55**

The RESIST Day subset complements the overnight recordings with a controlled afternoon-nap setting. Nap EEG was acquired with a 129-channel EGI HydroCel Geodesic Sensor Net from EGI/Philips at 250 Hz. To exploit the high-density montage, we retain all 128 scalp electrodes of the EGI net, band-pass filter at 0.5-30 Hz, resample to 100 Hz, and z-score normalize per clip. Each nap lasts 30 minutes. Of the recorded participants, 60 have both EEG and a valid CD-RISC-10 score. We binarize the CD-RISC-10 score at a threshold of 20.

4.2 Comparison Experiment

We adopt subject-wise five-fold cross-validation and average the results over five random seeds. We report balanced accuracy and AUC as classification metrics, and NLL and Cov@50 as regression

metrics. NLL is the Gaussian negative log-likelihood of the resilience score under each model's predictive distribution, where lower is better. Cov@50 is the fraction of subjects whose true score falls within the model's central 50% predictive interval, where closer to the nominal 0.5 is better.

We compare KoMaNet with task-specific deep EEG models, EEG foundation models, and Koopman-based time-series models on RESIST Night and RESIST Day, as shown in Table 1. The task-specific models are TCNet Ingolfsson et al. (2020), LMDA Miao et al. (2023), TSCeption Ding et al. (2022), CSPNet Jiang et al. (2024b), TSLANet Eldele et al. (2024), Conformer Song et al. (2022), and EEGNet Lawhern et al. (2018). The foundation models BIOT Yang et al. (2023), EEGPT Wang et al. (2024a), LaBraM Jiang et al. (2024a), CBraMod Wang et al. (2025), and CSBrain Zhou et al. (2025) are evaluated as frozen probes.

The Koopman-based models are Koopman-PatchTST, Koopman-Informer, and Koopman-Autoformer Forootani et al. (2026), DeepKoopman Lusch et al. (2018), Koopa Liu et al. (2023), ResKoopNet Xu et al. (2025), and SKOLR Zhang et al. (2025). We train each baseline on its designed EEG length, select hyperparameters by grid search, and aggregate chunk predictions into subject-level decisions by voting.

Among the baselines, TSCeption has the highest balanced accuracy on both subsets. Koopa has the highest AUC on RESIST Night, while CSPNet has the highest AUC on RESIST Day.

4.3 Ablation Study

We ablate each of the three KoMaNet modules. For the Self-supervised State Network, we drop the self-supervised objective in Eq. 4 and train the encoder by supervised chunk classification. For the Koopman Dynamic Network, we remove the state-conditioned Koopman operators and residual described in Eqs. 11 and 9 and feed the raw state embeddings to the subject manifold. For the Manifold Subject Network, we replace the distribution manifold with a mean-pool of each subject's chunk embeddings and a linear classifier.

As shown in Fig. 3, removing any single module lowers both balanced accuracy and AUC yet stays above chance, showing that the three modules are complementary. The Koopman Dynamic Network yields the largest gain on RESIST Night, underscoring the stable dynamics and deviation, whereas the Self-supervised State Network yields the largest gain on RESIST Day, where pretraining matters most under weaker dynamics and less data.

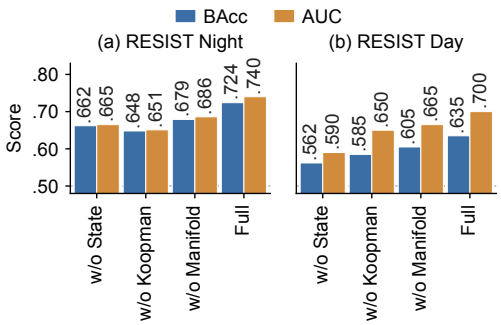


Figure 3: Ablation of the three KoMaNet modules on RESIST Night and RESIST Day. Removing any single module lowers balanced accuracy and AUC yet stays above chance.

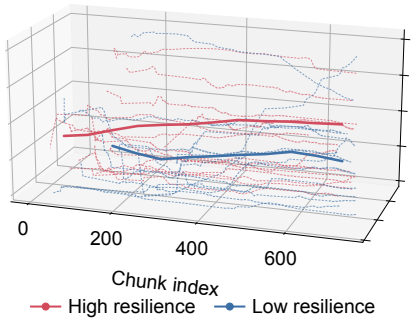


Figure 4: 3D visualization of the Koopman dynamic embeddings on RESIST Night, showing that as the night progresses the representations settle to a steady state.

4.4 Case Study

To further validate that the Koopman Dynamic Network captures the long-term dynamics relevant to psychological traits, we visualize the dynamic embeddings produced by the learned control-state Koopman operators. Red and blue denote the high- and low-resilience groups.

As shown in Fig. 4, the two group-mean trajectories stay separated along the vertical PCA axis throughout the night, and the dense cores of the two point clouds remain apart from sleep onset to the end of the recording, revealing a difference in their steady states. The high-resilience subjects trace a smoother accumulated trajectory, reflecting a smaller state-conditioned residual. The control-state Koopman operators thus organize the representations into a stable long-term structure on which the low- and high-resilience groups become distinguishable.

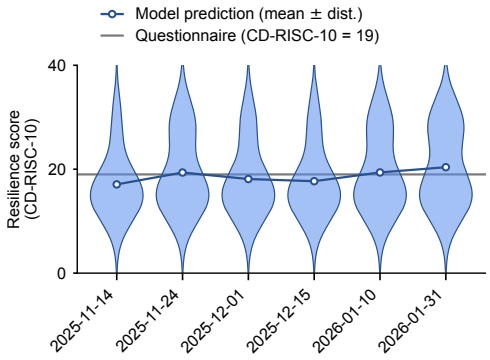


Figure 5: Cross-session stability of the predicted resilience for one subject with EEG recorded on six dates over about 2.5 months.

Because resilience is a slowly varying trait, the value decoded for one person should stay close over time and remain near that person's questionnaire score. To test this, we use one subject with a self-reported CD-RISC-10 of 19 obtained after the final acquisition session, whose afternoon-nap EEG was recorded on six different dates spanning roughly 2.5 months, and we predict resilience across sessions.

As shown in Figure 5, the predicted resilience is highly stable across the six dates. The per-session posterior mean stays within a narrow band and never departs from the questionnaire value by more than two points, with the predictive distribution centered on the ground-truth score throughout. This indicates that the model's resilience readout is temporally consistent for a given individual.

We further quantify this difference through the state-conditioned residual $\mathbf{r}_{n,t}$ defined in Eq. 10, which measures how far each transition departs from the stationary dynamics predicted for its brain state and on which the subject-level decision is based. As shown in Fig. 6, before sleep onset the two groups show little separation in their deviation. At sleep onset, the groups diverge. Once subjects consolidate into NREM sleep, the low-resilience deviation rises and remains consistently and substantially above the high-resilience deviation through the rest of the night. The reliably resilience-discriminative signal is thus carried by stable NREM sleep dynamics rather than by transient wake activity, consistent with the design of KoMaNet.

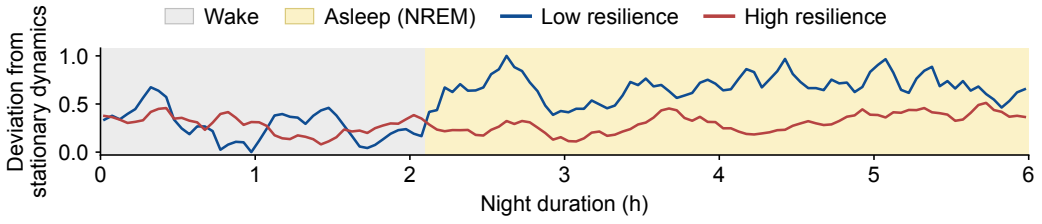
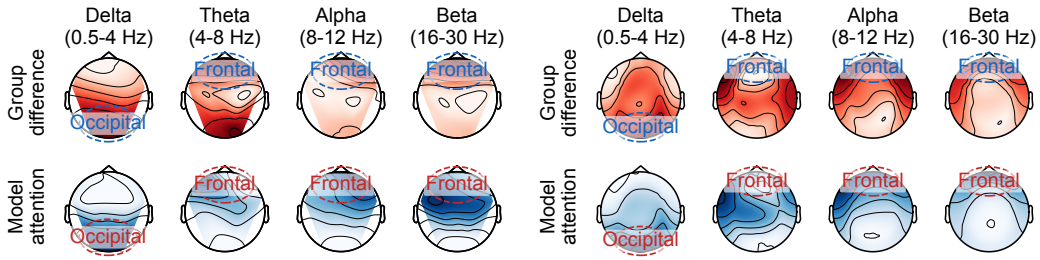


Figure 6: Deviation from stationary dynamics over the night on the RESIST Night subset, showing that low-resilience subjects deviate more than high-resilience subjects after sleep onset.

4.5 Interpretability Experiments

We probe which electrode-band patterns distinguish the high- and low-resilience groups. For each electrode-band pair, we quantify the group difference using Cohen's d of relative band power. We

also compute a discriminative attention score following Chung et al. (2024). For visualization, we normalize both measures to the range from 0 to 1 within each frequency band. Darker colors represent values closer to 1, and lighter colors represent values closer to 0. To obtain the attention score, we perturb each electrode-band pair in the raw EEG, propagate the perturbation through the trained KoMaNet, and average the resulting change in the subject-level low-resilience prediction within each group. We perform the analysis on RESIST Night using six-channel PSG and on RESIST Day using a 128-channel EGI net.



(a) RESIST Night (b) RESIST Day

Figure 7: EEG topography is shown across the $\delta/\theta/\alpha/\beta$ bands for RESIST Night in the left panel and RESIST Day in the right panel. In each panel, the top row shows the difference between the low- and high-resilience groups, and the bottom row shows the model's discriminative attention. Darker red indicates a larger normalized signed d within each frequency band, while darker blue indicates greater normalized attention. The color scale does not show the sign of d.

As shown in the left panel of Fig. 7, on RESIST Night the model's attention concentrates on frontal theta/alpha/beta and occipital delta. The top row also shows group differences in these bands and regions. Frontal activity is a well-documented correlate of cortical arousal and top-down emotion regulation Quaedflieg et al. (2020); KeunhoYoo et al. (2024), while posterior delta has been linked to homeostatic and motivational-affective processing Knyazev (2012); Wang et al. (2024c); Ben Simon et al. (2026). The same frontal and occipital pattern recurs in the right panel of Fig. 7 for RESIST Day, though it is more spatially diffuse, indicating a consistent if weaker cross-paradigm signal.

5 Limitations

Self-report resilience reflects an individual's self-perception, which is subject to cognitive bias and does not necessarily correspond to their true underlying resilience. Reconciling this gap between perceived and actual resilience lies beyond the scope of this work. In future work, we plan to validate the learned assessment against behavioral stress-reactivity measures as an external criterion.

6 Conclusions

We construct RESIST, a self-collected EEG dataset for stress resilience assessment with overnight-sleep and afternoon-nap recordings, and propose KoMaNet to predict resilience from brain-state dynamics. KoMaNet achieves AUCs of 0.740 on RESIST Night and 0.700 on RESIST Day, exceeding the strongest listed baseline on each subset. Model-attribution analyses highlight discriminative occipital delta and frontal theta/alpha/beta patterns. On RESIST Night, high-resilience subjects show smaller deviations from stationary dynamics after sleep onset than low-resilience subjects. Together, these results support the use of stable brain-state dynamics for EEG-based resilience assessment across both recording paradigms.

Reproducibility Statement

Section 3 describes KoMaNet, and Section 4 specifies the RESIST cohorts, evaluation protocol, baselines, and metrics. The appendix provides implementation details, the combined-cohort com-

parison, distribution analysis, and the external benchmark. The project code repository contains the model implementation and experiment scripts.

ETHICS STATEMENT

The RESIST data collection was approved by the ethics committee of the authors' institution. All participants provided written informed consent before EEG recording and questionnaire assessment. The study was conducted in accordance with the Declaration of Helsinki.

AI USE STATEMENT

Generative AI assisted with polishing and clarifying prose. The authors reviewed the AI-assisted work and take responsibility for the final content.

REFERENCES

Nishant Suresh Aswani and Saif Jabari. Koopman autoencoders learn neural representation dynamics. In *ICML*, 2025.

Anahit Babayan, Miray Erbey, Deniz Kumral, Janis D Reinelt, Andrea MF Reiter, Josefin Röbbig, H Lina Schaare, Marie Uhlig, Alfred Anwander, Pierre-Louis Bazin, et al. A mind-brain-body dataset of MRI, EEG, cognition, emotion, and peripheral physiology in young and old adults. *Scientific Data*, 6(1):1–21, 2019.

Eti Ben Simon, Vyoma D Shah, Olivia Murillo, Zavecz Zsofia, and Matthew P Walker. Impaired slow-wave sleep accounts for brain aging-related increases in anxiety. *Communications Psychology*, 2026.

Hanshu Cai, Zhenqin Yuan, Yiwen Gao, Shuting Sun, Na Li, Fuze Tian, Han Xiao, Jianxiu Li, Zhengwu Yang, Xiaowei Li, et al. A multi-modal open dataset for mental-disorder analysis. *Scientific Data*, 9(1):178, 2022.

Laura Campbell-Sills and Murray B Stein. Psychometric analysis and refinement of the connor-davidson resilience scale (CD-RISC): Validation of a 10-item measure of resilience. *Journal of Traumatic Stress*, 20(6):1019–1028, 2007.

Chris Xie Chen, Shirley Xin Li, Chung Shun Ho, Joey Wing Yan Chan, Leo King Wai Chan, Tatia Mei-Chun Lee, Yun Kwok Wing, and Ngan Yin Chan. Associations of psychological resilience with macro-and microstructures in nrem and rem sleep in adolescents. *International Journal of Clinical and Health Psychology*, 25(2):100570, 2025.

Hyunseung Chung, Sumin Jo, Yeonsu Kwon, and Edward Choi. Time is not enough: Time-frequency based explanation for time-series black-box models. In *International Conference on Information and Knowledge Management*, pp. 394–403, 2024.

W Tanner Creel and Richard E Hartman. Eeg entropy as a biomarker of emotion regulation and resilience. *Neuroscience Reports*, 2025.

W John Curtis and Dante Cicchetti. Emotion and resilience: A multilevel investigation of hemispheric electroencephalogram asymmetry and emotion regulation in maltreated and nonmaltreated children. *Development and Psychopathology*, 19(3):811–840, 2007.

Arnaud Delorme, Terrence Sejnowski, and Scott Makeig. Enhanced detection of artifacts in EEG data using higher-order statistics and independent component analysis. *Neuroimage*, 34(4):1443–1449, 2007.

Yi Ding, Neethu Robinson, Su Zhang, Qiuhao Zeng, and Cuntai Guan. TSception: Capturing temporal dynamics and spatial asymmetry from EEG for emotion recognition. *IEEE Transactions on Affective Computing*, 14(3):2238–2250, 2022.

Emadeldeen Eldele, Mohamed Ragab, Zhenghua Chen, Min Wu, and Xiaoli Li. TSLANet: Rethinking transformers for time series representation learning. In *International Conference on Learning Representations*, 2024.

Ali Forootani, Mohammad Khosravi, and Masoud Barati. Deepkoopformer: a koopman enhanced transformer based architecture for time series forecasting. *Nature Artificial Intelligence*, 2(1):35, 2026.

Sahen Gupta and Jayasankara Reddy. Eeg signatures of resilience across individuals with high and low anxiety. *NeuroRegulation*, 12(1):12–28, 2025.

Rumaisa Abu Hasan, Tong Boon Tang, Muhamad Saiful Bahri Yusoff, and Syed Saad Azhar Ali. Electroencephalography-based recognition of low mental resilience using multi-condition decision-level fusion approach. *Journal of Translational Engineering in Health and Medicine*, 2025.

Thorir Mar Ingolfsson, Michael Hersche, Xiaying Wang, Nobuaki Kobayashi, Lukas Cavigelli, and Luca Benini. EEG-TCNet: An accurate temporal convolutional network for embedded motor-imagery brain–machine interfaces. In *International Conference on Systems, Man, and Cybernetics*, pp. 2958–2965, 2020.

Weibang Jiang, Liming Zhao, and Bao-liang Lu. Large brain model for learning generic representations with tremendous eeg data in bci. In *International Conference on Learning Representations*, 2024a.

Xue Jiang, Lubin Meng, Xinru Chen, Yifan Xu, and Dongrui Wu. CSP-Net: Common spatial pattern empowered neural networks for EEG-based motor imagery classification. *Knowledge-Based Systems*, 305:112668, 2024b.

Tzyy-Ping Jung, Scott Makeig, Colin Humphries, Te-Won Lee, Martin J Mckeown, Vicente Iragui, and Terrence J Sejnowski. Removing electroencephalographic artifacts by blind source separation. *Psychophysiology*, 37(2):163–178, 2000.

Ioannis Kakkos, Elias Tzavellas, Eleni Feleskoura, Stamatis Mourtakos, Eleftherios Kontopodis, Ioannis Vezakis, Theodosis Kalamatianos, Emmanouil Synadinakis, George K Matsopoulos, Ioannis Kalatzis, et al. Eeg-based assessment of cognitive resilience via interpretable machine learning models. *AI*, 6(6):112, 2025.

Raffael Kalisch, Dewleen G Baker, Ulrike Basten, Marco P Boks, George A Bonanno, Eddie Brummelman, Andrea Chmitorz, Guillén Fernàndez, Christian J Fiebach, Isaac Galatzer-Levy, et al. The resilience framework as a strategy to combat stress-related disorders. *Nature Human Behaviour*, 1:784–790, 2017.

Kenny KeunhoYoo, Bowen Xiu, George Nader, Ariel Graff, Philip Gerretsen, Reza Zomorrodi, and Vincenzo De Luca. Power spectral analysis of resting-state eeg to monitor psychological resilience to stress. *Psychiatry Research Communications*, 4(3):100175, 2024.

Sirvan Khalighi, Teresa Sousa, José Moutinho Santos, and Urbano Nunes. ISRUC-Sleep: A comprehensive public dataset for sleep researchers. *Computer methods and programs in biomedicine*, 124:180–192, 2016.

Gennady G Knyazev. EEG delta oscillations as a correlate of basic homeostatic and motivational processes. *Neuroscience & Biobehavioral Reviews*, 36(1):677–695, 2012.

Niclas Kuper, Wiebke Käckenmester, and Jan Wacker. Resting frontal EEG asymmetry and personality traits: A meta–analysis. *European Journal of Personality*, 33(2):154–175, 2019.

Vernon J Lawhern, Amelia J Solon, Nicholas R Waytowich, Stephen M Gordon, Chou P Hung, and Brent J Lance. EEGNet: a compact convolutional neural network for EEG-based brain–computer interfaces. *Journal of Neural Engineering*, 15(5):056013, 2018.

Yong Liu, Chenyu Li, Jianmin Wang, and Mingsheng Long. Koopa: Learning non-stationary time series dynamics with koopman predictors. *NeurIPS*, 36:12271–12290, 2023.

Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9(1):4950, 2018.

Zhengqing Miao, Meirong Zhao, Xin Zhang, and Dong Ming. Lmda-net: A lightweight multi-dimensional attention network for general eeg-based brain-computer interfaces and interpretability. *NeuroImage*, 276:120209, 2023.

Frederick P Morgeson, Michael A Campion, Robert L Dipboye, John R Hollenbeck, Kevin Murphy, and Neal Schmitt. Reconsidering the use of personality tests in personnel selection contexts. *Personnel psychology*, 60(3):683–729, 2007.

Veronique Paban, Julien Modolo, Ahmad Mheich, and Mahmoud Hassan. Psychological resilience correlates with EEG source-space brain network flexibility. *Network Neuroscience*, 3(2):539–550, 2019.

Huy Phan, Fernando Andreotti, Navin Cooray, Oliver Y Chén, and Maarten De Vos. SeqSleepNet: end-to-end hierarchical recurrent neural network for sequence-to-sequence automatic sleep staging. *Transactions on Neural Systems and Rehabilitation Engineering*, 27(3):400–410, 2019.

Peter Putman, Bart Verkuil, Elsa Arias-Garcia, Ioanna Pantazi, and Charlotte van Schie. Eeg theta/beta ratio as a potential biomarker for attentional control and resilience against deleterious effects of stress on attention. *Cognitive, Affective, & Behavioral Neuroscience*, 14(2):782–791, 2014.

CWEM Quaedflieg, TR Schneider, J Daume, AK Engel, and L Schwabe. Stress impairs intentional memory control through altered theta oscillations in lateral parietal cortex. *Journal of Neuroscience*, 40(40):7739–7748, 2020.

Yonghao Song, Qingqing Zheng, Bingchuan Liu, and Xiaorong Gao. EEG conformer: Convolutional transformer for EEG decoding and visualization. *Transactions on Neural Systems and Rehabilitation Engineering*, 31:710–719, 2022.

Charles D Spielberger, Fernando Gonzalez-Reigosa, Angel Martinez-Urrutia, Luiz FS Natalicio, and Diana S Natalicio. The state-trait anxiety inventory. *Interamerican Journal of Psychology*, 5 (3 & 4), 1971.

Akara Supratak and Yike Guo. TinySleepNet: An efficient deep learning model for sleep stage scoring based on raw single-channel EEG. In *EMBC*, pp. 641–644, 2020.

Akara Supratak, Hao Dong, Chao Wu, and Yike Guo. DeepSleepNet: A model for automatic sleep stage scoring based on raw single-channel EEG. *Transactions on Neural Systems and Rehabilitation Engineering*, 25(11):1998–2008, 2017.

Chockalingam Viswesvaran and Deniz S Ones. Meta-analyses of fakability estimates: Implications for personality measurement. *Educational and Psychological Measurement*, 59(2):197–210, 1999.

Guangyu Wang, Wenchao Liu, Yuhong He, Cong Xu, Lin Ma, and Haifeng Li. EEGPT: Pretrained transformer for universal and reliable representation of EEG signals. *Advances in Neural Information Processing Systems*, 2024a.

Jiquan Wang, Sha Zhao, Haiteng Jiang, Shijian Li, Tao Li, and Gang Pan. Generalizable sleep staging via multi-level domain alignment. In *AAAI Conference on Artificial Intelligence*, volume 38, pp. 265–273, 2024b.

Jiquan Wang, Sha Zhao, Zhiling Luo, Yangxuan Zhou, Haiteng Jiang, Shijian Li, Tao Li, and Gang Pan. CBraMod: A criss-cross brain foundation model for EEG decoding. In *International Conference on Learning Representations*, 2025.

Xiaoyu Wang, Li Lin, Lei Zhan, Xianghong Sun, Zheng Huang, and Liang Zhang. Resting state EEG delta-beta amplitude-amplitude coupling: a neural predictor of cortisol response under stress. *Cognitive Neurodynamics*, 18(6):3995–4007, 2024c.

Xingjian Wu, Xiangfei Qiu, Hongfan Gao, Jilin Hu, Bin Yang, and Chenjuan Guo. K^2VAE: A koopman-kalman enhanced variational autoencoder for probabilistic time series forecasting. In *ICML*, 2025.

Yuanchao Xu, Kaidi Shao, Nikos K Logothetis, and Zhongwei Shen. Reskoopnet: Learning koopman representations for complex dynamics with spectral residuals. In *ICML*, 2025.

Chaoqi Yang, M. Brandon Westover, and Jimeng Sun. BIOT: Biosignal transformer for cross-data learning in the wild. *Advances in Neural Information Processing Systems*, 36:78240–78260, 2023.

Liuyin Yang, Qiang Sun, Ang Li, and Marc M Van Hulle. Are EEG foundation models worth it? comparative evaluation with traditional decoders in diverse BCI tasks. In *ICLR*, 2026.

Bailing Zhang. Koopman framework with self-supervised spectral alignment for multi-domain time-series modeling and prediction. *Neurocomputing*, pp. 131109, 2025.

Yitian Zhang, Liheng Ma, Antonios Valkanas, Boris N Oreshkin, and Mark Coates. SKOLR: Structured koopman operator linear RNN for time-series forecasting. In *ICML*, 2025.

Guozhen Zhao, Yan Ge, Biying Shen, Xingjie Wei, and Hao Wang. Emotion analysis for personality inference from EEG signals. *Transactions on Affective Computing*, 9(3):362–371, 2017.

Yuchen Zhou, Jiamin Wu, Zichen Ren, Zhouheng Lou, Baoyu Wang, Geyu Chen, Wenhui Zhao, et al. CSBrain: A cross-scale spatiotemporal brain foundation model for EEG decoding. *Advances in Neural Information Processing Systems*, 2025.

APPENDIX

A IMPLEMENTATION DETAILS

The clip encoder ω is a lightweight task-specific convolutional network pretrained on an external corpus. For RESIST Night, we use TinySleepNet Supratak & Guo (2020) pretrained on the public ISRUC corpus Khalighi et al. (2016). For RESIST Day, we use CSPNet Jiang et al. (2024b) pretrained on the public MODMA corpus Cai et al. (2022). For the resting-state recordings, we use an EEGNet backbone pretrained on an external resting-state corpus. Each encoder maps a clip to a d=16-dimensional state embedding and is then frozen. The teacher map φ concatenates per-electrode neurophysiological descriptors, such as spectral band power and functional connectivity, into a fixed-length vector. The control state is the sleep stage for the overnight recordings, an evenly split phase bin for the resting recordings, and a data-driven k-means cluster when reliable staging is unavailable. We train KoMaNet for 20 epochs with a batch size of 128 using the Adam optimizer with a learning rate of 0.001 and no weight decay. The Huber threshold is δ=1. The teacher window is K=3 for the sleep recordings and K=1 for the resting recordings.

All experiments are run on a server running Ubuntu 22.04.5 LTS. It has two Intel Xeon Platinum 8168 CPUs with 24 cores and 48 threads each, for a total of 96 threads, and 376 GB of RAM. Its NVIDIA Tesla V100-PCIE GPUs each have 32 GB of memory. Because KoMaNet is lightweight, each run uses a single V100 GPU. The models are implemented in Python 3 with PyTorch 2.5.1 and CUDA 12.1.

We evaluate both readouts of KoMaNet. For the binary classification of low and high resilience, we report balanced accuracy, denoted BAcc↑, and the area under the ROC curve, denoted AUC↑. Both metrics are appropriate under the strong class imbalance. For probabilistic regression of the continuous resilience score, we report the Gaussian negative log-likelihood, denoted NLL↓. It scores the predicted distribution $\mathcal{N}(m_n, \sigma_n^2+\sigma_L^2)$ against the noisy label. We also report the central interval coverage, denoted Cov@50. It measures the fraction of subjects whose labels fall inside their predicted 50% intervals. Its nominal value is 0.5.

B Comparison Experiments

The combined RESIST setting merges RESIST Night and RESIST Day using their six shared electrodes, F3, F4, C3, C4, O1, and O2, and a 30-minute sleep duration for each subject. As for RESIST Night, we use TinySleepNet pretrained on the ISRUC corpus as the encoder. We use the same subject-wise five-fold cross-validation over five random seeds, evaluation metrics, and baseline families described in Section 4.

Table 2: We compare methods in the combined RESIST setting, which uses the six shared electrodes and a 30-minute sleep duration. For each metric, the best result is shown in **bold**, and the second-best result is underlined.

	RESIST			
Method	BAcc↑	AUC↑	NLL↓	Cov@50↑
TCNet	$0.531_{\pm 0.018}$	$0.535_{\pm 0.083}$	3.72	0.32
LMDA	$0.565_{\pm 0.034}$	$0.552_{\pm 0.061}$	3.62	0.30
TSCeption	$0.578_{\pm 0.010}$	$0.568_{\pm 0.018}$	3.55	0.34
CSPNet	$\underline{0.581_{\pm 0.010}}$	$\underline{0.569_{\pm 0.007}}$	$\underline{3.52}$	$\underline{0.35}$
TSLANet	$0.552_{\pm 0.016}$	$0.532_{\pm 0.059}$	3.64	0.29
Conformer	$0.563_{\pm 0.032}$	$0.538_{\pm 0.020}$	3.66	0.31
EEGNet	$0.578_{\pm 0.026}$	$0.545_{\pm 0.045}$	3.58	0.32
BIOT	$0.557_{\pm 0.011}$	$0.562_{\pm 0.005}$	3.65	0.30
EEGPT	$0.520_{\pm 0.011}$	$0.530_{\pm 0.020}$	3.62	0.30
LaBraM	$0.568_{\pm 0.016}$	$0.548_{\pm 0.008}$	3.60	0.28
CBraMod	$0.540_{\pm 0.011}$	$0.535_{\pm 0.020}$	3.70	0.29
CSBrain	$0.534_{\pm 0.011}$	$0.522_{\pm 0.020}$	3.72	0.30
Ko-PatchTST	$0.531_{\pm 0.089}$	$0.522_{\pm 0.152}$	3.90	0.24
Ko-Informer	$0.524_{\pm 0.068}$	$0.515_{\pm 0.089}$	4.05	0.22
Ko-Autoformer	$0.528_{\pm 0.044}$	$0.518_{\pm 0.111}$	3.88	0.24
DeepKoopman	$0.558_{\pm 0.101}$	$0.540_{\pm 0.074}$	3.64	0.28
Koopa	$0.548_{\pm 0.079}$	$0.552_{\pm 0.199}$	3.70	0.27
ResKoopNet	$0.540_{\pm 0.108}$	$0.545_{\pm 0.145}$	3.75	0.26
SKOLR	$0.535_{\pm 0.101}$	$0.522_{\pm 0.155}$	3.85	0.25
KoMaNet	$\mathbf{0.613}_{\pm 0.009}$	$\mathbf{0.577}_{\pm 0.012}$	**3.40**	**0.54**

Table 2 reports the combined-cohort results. KoMaNet achieves a balanced accuracy of $0.613_{\pm 0.009}$ and an AUC of $0.577_{\pm 0.012}$, above the strongest listed baseline on each metric. Its NLL is 3.40, and its central 50% interval coverage is 0.54. This setting evaluates six-electrode, 30-minute recordings across the two cohorts.

The combined-cohort ablation gives BAcc/AUC of $0.566/0.548$ without the Self-supervised State Network, $0.558/0.539$ without the Koopman Dynamic Network, and $0.581/0.556$ without the Manifold Subject Network, compared with $0.613/0.577$ for the full model. Each module contributes to the combined-cohort classification results.

C Distribution Analysis

We further conduct a distribution analysis to show whether the model's predicted resilience distribution overlaps the label's uncertainty distribution. We model the observed label as $y \sim \mathcal{N}(m_{\text{true}}, \sigma_L^2)$ with the standard error of measurement $\sigma_L = \text{SD}\sqrt{1 - r_{xx}}$, where SD is the cohort score spread and r_{xx} is the test-retest reliability of the CD-RISC-10 questionnaire, for which $r_{xx} \approx 0.75$. As shown in Figure 8 for RESIST Night, for the low- and high-resilience groups separately, we superimpose all subjects' predicted distributions into one aggregate density and their label distributions into another, and compare the two. For each group g, we average the subject-level Gaussian densities to obtain the predicted density $p_g(s)$ and questionnaire-label density $q_g(s)$. We define their overlap

as $O_g = \int_0^{40} \min\{p_g(s), q_g(s)\}\, \mathrm{d}s$, the integral of the smaller density over the CD-RISC-10 score range. Both predicted densities closely track their label densities, with the low-resilience mass concentrating below the cutoff and the high-resilience mass above it, while the low-resilience minority is predicted with higher uncertainty.

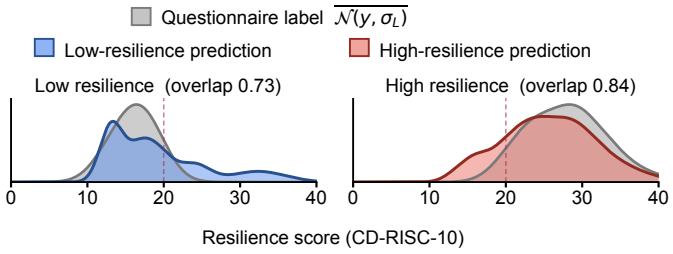


Figure 8: The model's predicted resilience distribution and the questionnaire label distribution on RESIST Night, showing that the predictions match the self-reported resilience scores.

D External Benchmark Experiment

To assess whether KoMaNet generalizes beyond our self-collected cohort, we additionally evaluate it on the public LEMON dataset using a related stress-relevant affective trait. The LEMON dataset is a public dataset comprising 227 healthy participants, including a young group of 108 males and 45 females with a mean age of 25.1 years and an elderly group of 37 males and 37 females with a mean age of 67.6 years Babayan et al. (2019). In this work, LEMON is not used as a second resilience benchmark, but rather as an external public dataset for evaluating the generalization of the proposed framework on a related stress-relevant affective trait. For EEG, 62-channel resting-state signals were recorded using active ActiCAP electrodes placed according to the 10-10 system. The recordings were band-pass filtered at 0.3-50 Hz, resampled to 100 Hz, and normalized using z-score standardization Wang et al. (2024b). For LEMON, we use scores from the State-Trait Anxiety Inventory, known as STAI Spielberger et al. (1971), as the target for a related affective trait prediction task. We follow the same training and evaluation protocol as in the benchmark experiment and compare against the task-specific and foundation EEG baselines. These baselines run in their native per-chunk regime with majority voting at the subject level, while KoMaNet aggregates the chunk sequence through the learned Koopman dynamics.

As shown in Table 3, KoMaNet also generalizes to the external LEMON benchmark, which exhibits greater inter-subject variability, achieving a balanced accuracy of 0.624 and an AUC of 0.617. Resting-state classification of a stable affective trait is intrinsically hard, and every baseline remains below KoMaNet. KoMaNet improves balanced accuracy by roughly 6 points over the strongest baseline. This demonstrates that the learned predictive and dynamical inductive biases, rather than the per-chunk encoder alone, remain effective on a larger and more heterogeneous public cohort. On LEMON, all three modules also contribute, though by a smaller margin than on RESIST Night. Removing the Koopman component lowers balanced accuracy from 0.624 to 0.574, removing the manifold to 0.588, and removing the pretraining to 0.599. The narrower gap reflects the weaker temporal dynamics of resting-state EEG compared with the rich long-range structure of overnight sleep, where the Koopman component is most beneficial.

Table 3: We compare KoMaNet with task-specific and foundation EEG models on the external LEMON benchmark. For each metric, the best result is shown in **bold**, and the second-best result is underlined.

Method	BAcc.↑	AUC↑	NLL↓	Cov@50↑
TCNet	$0.535_{\pm 0.030}$	$0.540_{\pm 0.050}$	3.96	0.35
LMDA	$0.548_{\pm 0.033}$	$0.552_{\pm 0.050}$	4.16	0.25
TSCeption	$0.540_{\pm 0.045}$	$0.545_{\pm 0.055}$	4.19	0.28
CSPNet	$\underline{0.560_{\pm 0.050}}$	$\underline{0.573_{\pm 0.055}}$	4.15	0.24
TSLANet	$0.538_{\pm 0.023}$	$0.542_{\pm 0.050}$	$\underline{3.73}$	0.27
Conformer	$0.552_{\pm 0.043}$	$0.565_{\pm 0.055}$	4.09	0.33
EEGNet	$0.545_{\pm 0.045}$	$0.550_{\pm 0.055}$	3.94	0.33
BIOT	$0.552_{\pm 0.032}$	$0.558_{\pm 0.050}$	4.05	0.29
EEGPT	$0.548_{\pm 0.028}$	$0.560_{\pm 0.030}$	3.83	$\underline{0.37}$
LaBraM	$0.550_{\pm 0.045}$	$0.555_{\pm 0.055}$	4.17	0.26
CBraMod	$0.542_{\pm 0.045}$	$0.548_{\pm 0.050}$	4.02	0.26
CSBrain	$0.538_{\pm 0.050}$	$0.544_{\pm 0.055}$	4.12	0.36
KoMaNet	$\mathbf{0.624}_{\pm 0.010}$	$\mathbf{0.617}_{\pm 0.014}$	**3.55**	**0.51**

E EEG BIOMARKERS OF STRESS RESILIENCE

Existing EEG studies have associated resilience and related stress-adaptive traits with several families of EEG markers. Spectral power and hemispheric asymmetry have been linked to resilience-related affect regulation and attentional control Curtis & Cicchetti (2007); Gupta & Reddy (2025); Putman et al. (2014). Sleep EEG studies further suggest that microstructural sleep patterns are related to resilience capacity Chen et al. (2025). Other work has examined resting-state power, source-space network organization, EEG entropy, and effective connectivity, showing that resilience is reflected not only in local oscillatory activity but also in broader patterns of EEG coordination and flexibility KeunhoYoo et al. (2024); Paban et al. (2019); Creel & Hartman (2025); Hasan et al. (2025); Kakkos et al. (2025). These analyses largely characterize correlational biomarkers, yet lack predictive models.