

# Who We Are, Where We Are: Mental Health at the Intersection of Person, Situation, and Large Language Models

Nikita Soni<sup>1</sup>, August Håkan Nilsson<sup>2</sup>, Syeda Mahwish<sup>1</sup>,  
Vasudha Varadarajan<sup>1</sup>, H. Andrew Schwartz<sup>1</sup>, Ryan L. Boyd<sup>3</sup>

<sup>1</sup>Department of Computer Science, Stony Brook University

<sup>2</sup>Oslo Metropolitan University, Oslo Business School

<sup>3</sup>University of Texas at Dallas, Dept. of Psychology

{nisoni, smahwish, vvaradarajan, has}@cs.stonybrook.edu,  
august.nilsson1907@gmail.com, boyd@utdallas.edu

## Abstract

Mental health is not a fixed trait but a dynamic process shaped by the interplay between individual dispositions and situational contexts. Building on interactionist and constructionist psychological theories, we develop interpretable models to predict well-being and identify adaptive and maladaptive self-states in longitudinal social media data. Our approach integrates person-level psychological traits (e.g., resilience, cognitive distortions, implicit motives) with language-inferred situational features derived from the Situational 8 DIAMONDS framework. We compare these theory-grounded features to embeddings from a psychometrically-informed language model that captures temporal and individual-specific patterns. Results show that our principled, theory-driven features provide competitive performance while offering greater interpretability. Qualitative analyses further highlight the psychological coherence of features most predictive of well-being. These findings underscore the value of integrating computational modeling with psychological theory to assess dynamic mental states in contextually sensitive and human-understandable ways.

## 1 Introduction

Understanding mental health through language has long been a foundational goal in clinical psychology and computational social science. Human expression — especially as manifested through digital communication — provides a unique window into internal states, social interactions, and psychological well-being. The CLPsych 2025 shared task builds on prior work in computational linguistics and clinical psychology, extending the analysis of mental health from static assessments to a dynamic, temporally anchored perspective. Seq2Psych, our interdisciplinary team, approaches this challenge by integrating psychological theory with computational methods, ensuring that our models are both

empirically grounded and practically applicable. Our work emphasizes not only predictive accuracy but also a principled, theory-grounded approach to interpretability, facilitating a nuanced explanation of how self-states fluctuate over time.

**Primary Contributions** made in this work include: (1) proposal of a theory-driven baseline that combines language-inferred person-level traits (e.g., well-being, cognitive distortions, resilience) with situational context features derived from the Situational 8 DIAMONDS framework (Rauthmann et al., 2014); (2) use of a human-centered language model (Soni et al., 2022, 2024c) trained on temporal user histories, to generate person-contextualized embeddings aligned with psychometric theory; (3) evaluation of these representations — individually and in hybrid configurations — for predicting well-being and identifying adaptive/maladaptive self-states in longitudinal text data; (4) an analysis of the most predictive psychological features to highlight interpretable connections between language, context, and mental health outcomes.

## 2 Background

Traditional models of psychological assessment often rely on static categories — diagnostic labels that imply stable traits or enduring conditions. However, integrative psychological theories emphasize that mental states are inherently dynamic, shaped by a complex interplay between individual dispositions and situational contexts (Buss, 1987; Ekehammar, 1974; Boyd and Markowitz, 2024). The constructionist view of emotions, for example, posits that emotional experiences emerge from interactions between an individual’s traits, cognitive processes, and environmental affordances (Barrett, 2017). Likewise, interactionist approaches in personality psychology highlight that adaptiveness or maladaptiveness of a given behavior is highly contingent upon situational fit (Mischel and Shoda,

1995; Fleeson, 2004).

The notion of situational fit is central to understanding mental health dynamics. Psychological well-being is not merely an individual trait but a function of how well a person’s responses align with their context. A behavior that is adaptive in one situation may be maladaptive in another. For instance, hypervigilant behaviors may be adaptive in some contexts (e.g., military personnel in combat situations), but highly *maladaptive* in another (e.g., a classroom or casual social gathering; Vyas et al. 2023). This perspective aligns with the broader movement in psychology that views well-being as a dynamic process rather than a fixed state (see, e.g., Hollenstein, 2015)

## 2.1 A Principled Baseline: Integration of Person-Level Traits and Situational Context

To model mental health dynamics in a principled manner, our approach combines person-level traits with psychological dimensions of the situation. Specifically, we leverage:

### Psychological Characteristics of the Situation

Using a large language model, we annotated each post for the psychological characteristics of its context, based on the Situational 8 DIAMONDS (S8D) framework (Rauthmann et al., 2014). This framework captures eight psychosocial aspects of a given situation — **D**uty, **I**ntellect, **A**dversity, **M**ating, **p**ositivity, **N**egativity, **D**eception, and **S**ociality — that shape the meaning and context of person-environment transactions.

**Person-Level Traits (PLT)** We employed existing models and methods to assess and estimate individual differences in implicit motives (Nilsson et al., 2024), depression and anxiety (Gu et al.), harmony in life and satisfaction with life (Kjell et al., 2022a) valence (Eijsbroek, 2023), cognitive distortions (Varadarajan et al., 2025), and resilience (Mahwish et al.). These traits serve as stable psychological anchors that interact dynamically with situational contexts in well-established fashions (Mejía and Hooker, 2015; Ungar, 2013; Joiner Jr. and Timmons, 2009).

By combining these features, we constructed a baseline model that aligns with psychological theory, providing an interpretable reference point against which more data-driven approaches can be evaluated. Our method represents a true interdisciplinary effort in computational social sci-

ence, bridging insights from personality psychology, emotion theory, and NLP to advance the study of mental health dynamics in digital contexts.

## 2.2 Human Language Modeling: HaRT for Person-Contextual Embeddings

Our principled baseline approach offers a clear explanatory mechanism for predicting well-being and distinguishing between adaptive and maladaptive self-states. However, we anticipate that more advanced language models will enhance predictive accuracy and provide a richer representation of language and individuals. HaRT, trained on the Human Language Modeling (HuLM) task — which predicts the next word based on prior words, incorporating a latent user representation derived from their temporal historical language — enables a person-contextualized understanding of language (Soni et al., 2022). Grounded in psychometric theory on the stability of psychological traits (Watson, 2004) HuLM processes an author’s language collectively, recognizing that linguistic patterns are best understood within the context of the individual themselves, over time (Soni et al., 2024b; Ganesan et al., 2024). This approach is particularly well-suited for our tasks, given the dataset’s longitudinal structure, where language is nested within individuals, and has proven to be effective in mental health assessments (Ganesan et al., 2022; Varadarajan et al., 2024b), psychological assessments (Soni et al., 2025), and user attributes assessments (Soni et al., 2024a).

## 3 Data & Tasks

**Dataset.** The CLPsych 2025 shared task (Tseriotou et al., 2025) provided annotated evidence for adaptive and maladaptive self-states (Slonim, 2024) as spans of texts from posts written by individuals historically in addition to a score representing the overall well-being in a post. The data consists of 30 users (timelines) with a total of 343 posts of which 199 posts were annotated.

**Shared Tasks.** The shared tasks focus on the longitudinal modeling of changes in individual’s mood and states (Shing et al., 2018; Zirikly et al., 2019; Tsakalidis et al., 2022). We participate in 2 sub-tasks targeted at post-level judgments: a) predicting the overall well-being, and b) identifying evidence for adaptive and maladaptive self-states.

## 4 Methods

We extracted two categories of features: a theory-informed baseline — comprising Situational 8 DIAMONDS (S8D) and Person-Level Traits (PLT; see Section 2.1) — and person-contextualized embeddings. While PLT features were computed at both sentence and post levels, S8D were limited to post-level annotations due to their reliance on broader context.

**Situational 8 DIAMONDS (S8D).** We used Deepseek-R1 (DeepSeek-AI et al., 2025) with few-shot prompting to infer scores for each of the eight situational dimensions at the post level. Each dimension was prompted separately using two manually annotated exemplars tailored to its psychological construct (see § A.2). Scores ranged from 1 (not present) to 9 (highly present), reflecting the inferred prominence of each situational characteristic.

**Person-Level Traits (PLT).** We extracted 19 features across four subdomains:

*Implicit Motives.* Following Nilsson et al. (b), we applied fine-tuned RoBERTa-Large models to estimate three subconscious motives — achievement, affiliation, and power — at the sentence level. These predictions were averaged and adjusted for word count to yield post-level scores (details in Appendix A.3).

*Mental Health.* Using the Language-Based Assessment Model Library (Nilsson et al., a), we inferred six psychological dimensions: valence (Eijsbroek, 2023), harmony in life, satisfaction with life (Kjell et al., 2022b), anxiety, and two depression indices (Gu et al., 2024). Features were extracted at the sentence level and averaged to generate post-level estimates.

*Resilience.* We implemented the Resilience through Language Modeling (ReLM) framework (Mahwish et al.) to compute scores for nine resilience-related facets (e.g., optimism, coping toolkit) at both sentence and post levels. See Appendix A.4 for details.

*Cognitive Distortions.* Drawing on prior work (Varadarajan et al., 2025), we used pre-trained models to estimate levels of cognitive distortion, a known correlate of maladaptive emotional states (Mann et al., 2002; Bathina et al., 2021), at both sentence and post levels.

**Person-Contextual Embeddings: HaRT.** We fine-tuned HaRT (Human-aware Recurrent Transformer) (Soni et al., 2022) to predict continuous

	5-fold Ridge CV	
	$r \uparrow$	$MSE \downarrow$
S8D	0.528	2.556
Dist	0.365	3.059
ReLM	0.533	2.538
PLT	<b>0.629</b>	<b>2.149</b>
S8D + ReLM + Dist	0.623	2.178
S8D + PLT	0.622	2.174

Table 1: Pearson correlation ( $r$ ) and Mean Squared Error (MSE) when training a ridge regression model using different “principled” baseline features to predict continuous Well-being scores using nested 5-fold cross-validation.

well-being scores at the post level and binary adaptive/maladaptive labels at the sentence level. To do so, we split the CLPsych training data into internal training and validation sets. HaRT processes users’ historical posts in sequence, enabling the generation of temporally informed, person-specific embeddings at both the sentence and post levels.

We evaluated these embeddings across three tasks using 5-fold nested cross-validation: (a) continuous well-being prediction via ridge regression, (b) adaptive label prediction using logistic regression, and (c) maladaptive label prediction using logistic regression. We chose 5-fold CV to mitigate overfitting given the small sample size. For all classifiers, we used a penalty range of [10, 0, -1, -0.10, 0.10]. For span identification, we predicted label probabilities and applied thresholds of 0.45 (adaptive) and 0.4 (maladaptive) to extract evidence-level annotations.

## 5 Results & Discussion

**Well-being Scores.** Situational characteristics (S8D) inferred from posts were predictive of annotated well-being scores (see Table 1). When combined with PLT features, our “principled” baseline — grounded in interactionist theory — yielded improved performance. The psychometric theory-inspired HaRT model outperformed baselines on the internal validation set, although we observed signs of overfitting in 5-fold CV (see Table 2). Nonetheless, official results showed similar trends (Table A2), with the theory-driven S8D +PLT baseline outperforming the theory-agnostic HaRT<sub>WB-FT + Ridge</sub> variant.

**Adaptive and Maladaptive States.** HaRT models performed well in the binary classification of adaptive and maladaptive self-states (Table A4),

	Internal Val Set		5-fold Ridge CV	
	$r \uparrow$	$MSE \downarrow$	$r \uparrow$	$MSE \downarrow$
HaRT <sub>WB-FT</sub>	0.684	1.828	0.876	0.828
HaRT <sub>WB-FT</sub> + S8D + ReLM + Dist	-	-	0.883	0.787
HaRT <sub>WB-FT</sub> + S8D + PLT	-	-	0.884	0.783

Table 2: Pearson correlation ( $r$ ) and Mean Squared Error (MSE) when fine-tuning HaRT using internal train and validation splits and further training a ridge regression model using resulting embeddings and principled baseline features. Note: Pearson  $r$  values for the 5-fold ridge CV numbers are likely inflated due to partial data contamination across the fine-tuning and cross-validation datasets. The internal validation set was used while finetuning the HaRT model, after which the weights were separately used as inputs for well-being task. The internal validation numbers have been omitted for the post-finetuned model.

	5-fold Log.reg. CV	
	$F1_{macro}$	AUC
Dist <sub>adaptive</sub>	<b>0.54</b>	0.75
HaRT <sub>WB-FT</sub>	0.50	0.74
HaRT <sub>WB-FT</sub> + ReLM + Dist <sub>adaptive</sub>	0.53	<b>0.76</b>
PLT <sub>adaptive</sub>	0.48	0.66
HaRT <sub>WB-FT</sub> + PLT <sub>adaptive</sub>	0.52	<b>0.76</b>
Dist <sub>maladaptive</sub>	0.56	0.73
HaRT <sub>WB-FT</sub>	0.56	0.73
HaRT <sub>WB-FT</sub> + ReLM + Dist <sub>maladaptive</sub>	0.57	0.76
PLT <sub>maladaptive</sub>	0.49	0.70
HaRT <sub>WB-FT</sub> + PLT <sub>maladaptive</sub>	<b>0.58</b>	<b>0.77</b>

Table 3: Macro F1 and AUC results when training a logistic regression model to predict binary adaptive and maladaptive labels separately over sentences split from posts.

with additional gains observed when combined with PLT features (Table 3). While Dist and PLT features alone showed reasonable performance, they produced minimal variation in predicted probabilities across examples (Figure 1). In contrast, HaRT-based models exhibited greater sensitivity to language variation and were more effective at identifying adaptive and maladaptive evidence spans. Additional supporting results and probability distributions can be found in Appendix § A.6.

## 5.1 Discussion.

Our interactionist, theory-based “principled” baseline approach effectively predicts annotated well-being scores. However, it struggled to capture ev-

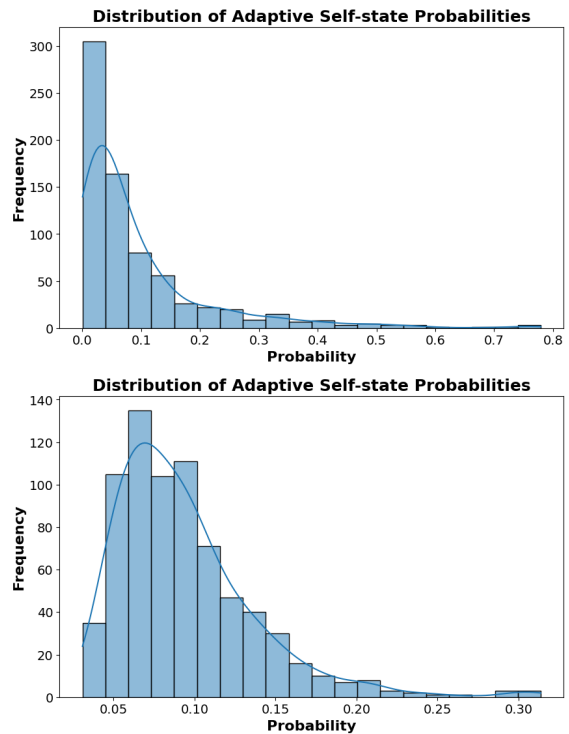


Figure 1: Distribution of probabilities to predict adaptive state for a given sentence. On the top is using HaRT<sub>WB-FT</sub> + PLT features, and the bottom is using PLT features in Logistic Regression models.


idence of adaptive and maladaptive states within posts, highlighting the challenge of disentangling self-states from their situational context — an issue that even human observers can find difficult to assess accurately (Uleman et al., 1996; Nisbett et al., 1973; Ross, 1977).

To further explore the predictive power of our principled baseline features, we conduct a qualitative analysis of well-being correlations. As shown in Figure 2, the top three features positively associated with well-being scores are: ‘satisfaction with life’ (from mental health in PLT), ‘positivity in the situation’ (from S8D), and ‘harmony in life’ (from mental health in PLT). Conversely, the top three features negatively correlated with well-being scores include: ‘higher power belief’ (resilience from PLT), ‘depression scale’ (mental health from PLT), and overall ‘resilience score’ (from PLT).

Additionally, Figure 2 shows that the ridge regression model, leveraging our principled baseline features, assigns positive importance to ‘sociality’, ‘positivity’, and ‘intellect’ (all from S8D), while attributing negative importance to ‘daily lived experience’ (from resilience in PLT), ‘need for affiliation’



**Pearson correlation between Well-being and S8D + PLT features**



**Ridge regression Beta coefficient for Well-being using S8D and PLT**




Figure 2: Qualitative analysis of features in our principled baseline consisting of psychological characteristics of the situation and person-level traits. Left: Pearson correlation coefficients; Right: Ridge regression beta coefficients for predicting well-being with the S8D and PLT features.

(from implicit motives in PLT)<sup>1</sup>, and ‘belief in a higher power’ (from resilience in PLT).

These findings align with prior research indicating that well-being is closely tied to life satisfaction and positive social interactions (Diener et al., 1999; Cacioppo and Cacioppo, 2014). Similarly, negative associations with depression, present-focus, and certain aspects of religiosity are consistent with existing psychological literature on mental health dynamics (Beck, 1967; Himmelstein et al., 2018; McCullough and Larson, 1999; Braam and Koenig, 2019). A more detailed discussion of these results

<sup>1</sup>High expression of affiliation-related language may indicate a frustrated, rather than fulfilled, need for social connection. Classic motivation theories (McClelland, 1987) suggest that individuals who frequently verbalize their affiliation motive may be experiencing social deprivation or unmet interpersonal needs. This aligns with research on compensatory behaviors, where socially disconnected individuals often amplify affiliative overtures to seek connection (Baumeister and Leary, 1995; Richman and Leary, 2009).

can be found in the Appendix § A.5.

## 6 Conclusion

Mental health is not a static trait but a dynamic outcome shaped by ongoing interactions between person and context. In this work, we operationalized this psychological insight by combining person-level traits and situational characteristics — core tenets of interactionist and constructionist theory — to model well-being and adaptive self-states in language.

Our theory-driven baseline, built from the Situational 8 DIAMONDS and language-inferred psychological traits, demonstrated strong performance while offering interpretable, psychologically grounded predictions. Features like positivity, satisfaction with life, and harmony in life emerged as key indicators of well-being, while markers of cognitive distortion and unmet affiliation needs were linked to maladaptive patterns. HaRT’s person-contextualized embeddings added value in modeling temporal variation, particularly for adaptive and maladaptive evidence detection.

These findings highlight the value of bridging computational models with psychological theory — not only to improve prediction, but to ensure outputs are meaningful and human-understandable. Future work should explore how different contexts modulate trait adaptiveness, and how language-based systems might support more flexible, resilient self-states over time.

By integrating theory and computation, we move toward systems that understand individuals not as fixed entities, but as contextually situated and dynamically evolving.

## Limitations

While this work offers a psychologically grounded and interpretable approach to modeling mental health from language, several limitations must be acknowledged — both technical and conceptual.

First, our analyses are constrained by the scale and structure of the CLPsych 2025 dataset. With only 30 users and fewer than 200 annotated posts, the generalizability of our findings is limited. Although we used robust cross-validation and avoided tuning on the test set, future work should evaluate these models on larger, more diverse, and demographically representative datasets.

Second, the ground truth labels themselves are inherently interpretive, reflecting human judgments

of well-being and self-states based on textual evidence. This raises a broader epistemological question: to what extent can self-states be reliably inferred from language alone? Our work assumes that linguistic expressions are sufficient proxies for psychological states — a premise that, while useful for modeling, must be critically examined in clinical and applied settings.

Third, our feature extraction relies on pretrained models and heuristics that may carry latent biases or be insensitive to cultural or contextual nuance. For example, expressions of distress or resilience may vary across communities, and models trained on general corpora may fail to capture such variation meaningfully.

Finally, although we draw from psychological theory, our models remain correlational. They can identify linguistic markers of mental health but in no way are able to definitively speak to underlying mechanisms, causal relationships, or interventions. Future work should incorporate longitudinal clinical assessments to validate language-based features against real-world outcomes.

These limitations do not undermine the value of this work, but rather highlight the need for computational psychology to remain grounded in balancing predictive power with interpretive care, and data-driven modeling with theoretical accountability.

## Ethical Considerations

Modeling mental health from language data presents profound ethical challenges. While the tools developed in this work aim to advance understanding of mental states in contextually sensitive and interpretable ways, their misuse — or even well-intentioned use without adequate safeguards — poses real risks to privacy, autonomy, and well-being.

First and foremost is the issue of consent. Although the data used in this study were shared with participant permission as part of a structured research challenge, this controlled environment does not reflect broader real-world settings in which language-based models might be applied. Any future deployment must ensure that individuals are aware of — and have control over — how their language data are interpreted, stored, and acted upon.

Second, language-based inferences about mental health are probabilistic and inherently contain some degree of uncertainty. Over-reliance on

model outputs — particularly in clinical, legal, or surveillance contexts — could lead to misdiagnoses, stigmatization, or unwarranted interventions. The interpretability of our features helps mitigate this risk, but human oversight and psychological expertise remain essential in any applied use.

Third, there are critical concerns around representation and bias. Our models are trained on English-language data from social media forums, which may reflect particular cultural, demographic, and socioeconomic perspectives. As a result, model outputs may not generalize across populations and could even reinforce existing inequities if deployed indiscriminately. Expanding the diversity of training data and engaging with cultural psychology are necessary steps forward.

Finally, as researchers in computational social science, we must remain vigilant about the institutional and commercial pressures that can shape how mental health technologies are built and used. The potential to infer mental states from language at scale invites both promise and peril. Ethical research in this space demands more than compliance — it requires an ongoing commitment to transparency, self-reflection, and the prioritization of human dignity.

## Acknowledgments

We are grateful to the organizers of the CLPsych 2025 Shared Task for their efforts in curating the dataset, designing an impactful task, and fostering interdisciplinary collaboration.

This research is supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via the HIATUS Program contract #2022-22072200005, a grant from the NIH-NIAAA (R01 AA028032), and a grant from the CDC/NIOSH (U01 OH012476), Developing and Evaluating Artificial Intelligence-based Longitudinal Assessments of PTSD in 9/11 Responders. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, any other government organization, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

## References

- Lisa Feldman Barrett. 2017. *How emotions are made: The secret life of the brain*. How emotions are made: The secret life of the brain. Houghton Mifflin Harcourt, Boston, MA.
- Krishna C Bathina, Marijn Ten Thij, Lorenzo Lorenz-Luaces, Lauren A Rutter, and Johan Bollen. 2021. Individuals with depression express more distorted thinking on social media. *Nature human behaviour*, 5(4):458–466.
- Roy F. Baumeister and Mark R. Leary. 1995. **The need to belong: Desire for interpersonal attachments as a fundamental human motivation**. *Psychological Bulletin*, 117(3):497–529. Place: US Publisher: American Psychological Association.
- Aaron T. Beck. 1967. *Depression: Clinical, experimental, and theoretical aspects*, no additional printings listed edition edition. see notes for publisher info.
- Ryan L. Boyd and David M. Markowitz. 2024. **Verbal behavior and the future of social science**. *American Psychologist*, pages 1–23. Place: US Publisher: American Psychological Association.
- Arjan W. Braam and Harold G. Koenig. 2019. **Religion, spirituality and depression in prospective studies: A systematic review**. *Journal of Affective Disorders*, 257:428–438.
- David M. Buss. 1987. **Selection, evocation, and manipulation**. *Journal of Personality and Social Psychology*, 53(6):1214–1221.
- John T. Cacioppo and Stephanie Cacioppo. 2014. **Social relationships and health: The toxic effects of perceived social isolation**. *Social and Personality Psychology Compass*, 8(2):58–72. [\\_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/spc3.12087](https://onlinelibrary.wiley.com/doi/pdf/10.1111/spc3.12087).
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojuan Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanxia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yudian Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. **DeepSeek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning**. *arXiv preprint*. ArXiv:2501.12948 [cs].
- Ed Diener, Eunkook M. Suh, Richard E. Lucas, and Heidi L. Smith. 1999. **Subjective well-being: Three decades of progress**. *Psychological Bulletin*, 125(2):276–302. Place: US Publisher: American Psychological Association.
- Veerle C Eijlsbroek. 2023. **A comparison of the assessments from a person and the person’s AI-model using computational language assessments**.
- Bo Ekehammar. 1974. **Interactionism in personality from a historical perspective**. *Psychological Bulletin*, 81(12):1026–1048. Place: US Publisher: American Psychological Association.
- William Fleeson. 2004. **Moving personality beyond the person-situation debate: The challenge and the opportunity of within-person variability**. *Current Directions in Psychological Science*, 13(2):83–87. Publisher: SAGE Publications Inc.
- Adithya V Ganesan, Siddharth Mangalik, Vasudha Varadarajan, Nikita Soni, Swanie Juhng, João Sedoc, H Andrew Schwartz, Salvatore Giorgi, and Ryan L Boyd. 2024. From text to context: Contextualizing language with humans, groups, and communities for socially aware nlp. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 5: Tutorial Abstracts)*, pages 26–33.
- Adithya V Ganesan, Vasudha Varadarajan, Juhi Mittal, Shashanka Subrahmanya, Matthew Matero, Nikita

- Soni, Sharath Chandra Guntuku, Johannes Eichstaedt, and H Andrew Schwartz. 2022. Wwbp-sqt-lite: Multi-level models and difference embeddings for moments of change identification in mental health forums. In *Proceedings of the Eighth Workshop on Computational Linguistics and Clinical Psychology*, pages 251–258.
- Zhuojun Gu, Katarina Kjell, H. Andrew Schwartz, and Oscar Kjell. Natural language response formats for assessing depression and worry with large language models: A sequential evaluation with model pre-registration. *Assessment*.
- Zhuojun Gu, Katarina Kjell, H Andrew Schwartz, Oscar Kjell, et al. 2024. Natural language response formats for assessing depression and worry with large language models: A sequential evaluation with model pre-registration.
- Philip Himmelstein, Scott Barb, Mark A. Finlayson, and Kymberly D. Young. 2018. Linguistic analysis of the autobiographical memories of individuals with major depressive disorder. *PLOS ONE*, 13(11):e0207814.
- Tom Hollenstein. 2015. This time, it's real: Affective flexibility, time scales, feedback loops, and the regulation of emotion. *Emotion Review*, 7(4):308–315. Publisher: SAGE Publications.
- Thomas E. Joiner Jr. and Katherine A. Timmons. 2009. Depression in its interpersonal context. In *Handbook of depression, 2nd ed*, pages 322–339. The Guilford Press, New York, NY, US.
- Oscar N. E. Kjell, Sverker Sikström, Katarina Kjell, and H. Andrew Schwartz. 2022a. Natural language analyzed with AI-based transformers predict traditional subjective well-being measures approaching the theoretical upper limits in accuracy. *Scientific Reports*, 12(1):3918.
- Oscar NE Kjell, Sverker Sikström, Katarina Kjell, and H Andrew Schwartz. 2022b. Natural language analyzed with ai-based transformers predict traditional subjective well-being measures approaching the theoretical upper limits in accuracy. *Scientific reports*, 12(1):3918.
- Kurt Kroenke, Robert L. Spitzer, and Janet B. W. Williams. 2011. *Patient Health Questionnaire-9*. Institution: American Psychological Association.
- Syeda Mahwish, Ryan L. Boyd, Vasudha Varadarajan, Roman Kotov, Benjamin J. Luft, H. Andrew Schwartz, and Sean A. P. Clouston. Measuring resilience using language modeling (ReLM): A computational approach to observing resilience.
- Ruth E Mann, Anthony R Beech, T Ward, DR Laws, and SM Hudson. 2002. Cognitive distortions, schemas, and implicit theories. *Sexual Deviance: Issues, Theories and Treatment*, pages 135–153.
- David C. McClelland. 1987. *Human motivation*. CUP Archive.
- Michael E. McCullough and David B. Larson. 1999. Religion and depression: A review of the literature. *Twin Research and Human Genetics*, 2(2):126–136.
- Shannon T. Mejía and Karen Hooker. 2015. Emotional well-being and interactions with older adults' close social partners: Daily variation in social context matters. *Psychology and Aging*, 30(3):517–528. Place: US Publisher: American Psychological Association.
- Walter Mischel and Yuichi Shoda. 1995. A cognitive-affective system theory of personality: Reconceptualizing situations, dispositions, dynamics, and invariance in personality structure. *Psychological Review*, 102(2):246–268. Place: US Publisher: American Psychological Association.
- August Nilsson, Zhuojun Gu, Veerle C Eijsbroek, Katarina Kjell, Salvatore Giorgi, Roman Kotov, H. Andrew Schwartz, and Oscar N. E. Kjell. a. The Language-Based Assessment Model (L-BAM) library: Open model sharing for independent validation and broader application.
- August Nilsson, J Malte Runge, Oscar Kjell, Adithya V Ganesan, Carl Viggo Nilsson, et al. 2024. Automatic implicit motives codings are as accurate as humans' and 99% faster.
- August Nilsson, J. Malte Runge, Oscar Kjell, Nikita Soni, Adithya V. Ganesan, and Carl Viggo Nilsson. b. Automatic implicit motives codings are as accurate as humans' and 99% faster. *Journal of Personality and Social Psychology: Personality Processes and Individual Differences*.
- Richard E. Nisbett, Craig Caputo, Patricia Legant, and Jeanne Marecek. 1973. Behavior as seen by the actor and as seen by the observer. *Journal of Personality and Social Psychology*, 27(2):154–164. Place: US Publisher: American Psychological Association.
- Lenore S. Radloff. 1977. The CES-D Scale: A self-report depression scale for research in the general population. *Applied Psychological Measurement*, 1(3):385–401. Place: US Publisher: Sage Publications.
- John F. Rauthmann, David Gallardo-Pujol, Esther M. Guillaume, Elysia Todd, Christopher S. Nave, Ryne A. Sherman, Matthias Ziegler, Ashley Bell Jones, and David C. Funder. 2014. The Situational Eight DIAMONDS: A taxonomy of major dimensions of situation characteristics. *Journal of Personality and Social Psychology*, 107(4):677–718.
- Laura Smart Richman and Mark R. Leary. 2009. Reactions to discrimination, stigmatization, ostracism, and other forms of interpersonal rejection. *Psychological Review*, 116(2):365–383.
- Lee Ross. 1977. The intuitive psychologist and his shortcomings: Distortions in the attribution Process. In Leonard Berkowitz, editor, *Advances in Experimental Social Psychology*, volume 10, pages 173–220. Academic Press.



- Han-Chin Shing, Suraj Nair, Ayah Zirikly, Meir Friedenberg, Hal Daumé III, and Philip Resnik. 2018. [Expert, crowdsourced, and machine assessment of suicide risk via online postings](#). In *Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic*, pages 25–36, New Orleans, LA. Association for Computational Linguistics.
- Dana Atzil Slonim. 2024. Self-other dynamics (sod): A transtheoretical coding manual.
- Nikita Soni, Niranjana Balasubramanian, H Schwartz, and Dirk Hovy. 2024a. Comparing pre-trained human language models: Is it better with human context as groups, individual traits, or both? In *Proceedings of the 14th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis*, pages 316–328.
- Nikita Soni, Pranav Chitale, Khushboo Singh, Niranjana Balasubramanian, and H. Andrew Schwartz. 2025. [Evaluation of llms-based hidden states as author representations for psychological human-centered nlp tasks](#). Preprint, arXiv:2503.00124.
- Nikita Soni, Lucie Flek, Ashish Sharma, Diyi Yang, Sara Hooker, and H. Andrew Schwartz. 2024b. Proceedings of the 1st human-centered large language modeling workshop. In *Proceedings of the 1st Human-Centered Large Language Modeling Workshop*.
- Nikita Soni, Matthew Matero, Niranjana Balasubramanian, and H. Andrew Schwartz. 2022. [Human language modeling](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 622–636, Dublin, Ireland. Association for Computational Linguistics.
- Nikita Soni, H. Andrew Schwartz, João Sedoc, and Niranjana Balasubramanian. 2024c. Large human language models: A need and the challenges. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8623–8638.
- Adam Tsakalidis, Jenny Chim, Iman Munire Bilal, Ayah Zirikly, Dana Atzil-Slonim, Federico Nanni, Philip Resnik, Manas Gaur, Kaushik Roy, Becky Inkster, Jeff Leintz, and Maria Liakata. 2022. [Overview of the CLPsych 2022 shared task: Capturing moments of change in longitudinal user posts](#). In *Proceedings of the Eighth Workshop on Computational Linguistics and Clinical Psychology*, pages 184–198, Seattle, USA. Association for Computational Linguistics.
- Talia Tseriotou, Jenny Chim, Ayal Klein, Aya Shamir, Guy Dvir, Iqra Ali, Cian Kennedy, Guneet Singh Kohli, Anthony Hills, Ayah Zirikly, Dana Atzil-Slonim, and Maria Liakata. 2025. Overview of the clpsych 2025 shared task: Capturing mental health dynamics from social media timelines. In *Proceedings of the 10th Workshop on Computational Linguistics and Clinical Psychology*. Association for Computational Linguistics.
- James S. Uleman, Leonard S. Newman, and Gordon B. Moskowitz. 1996. [People as flexible interpreters: Evidence and issues from spontaneous trait inference](#). In Mark P. Zanna, editor, *Advances in Experimental Social Psychology*, volume 28, pages 211–279. Academic Press.
- Michael Ungar. 2013. [Resilience, trauma, context, and culture](#). *Trauma, Violence, & Abuse*, 14(3):255–266. Publisher: SAGE Publications.
- Vasudha Varadarajan, Allison Lahnala, Adithya V. Ganesan, Gourab Dey, Siddharth Mangalik, Ana-Maria Bucur, Nikita Soni, Rajath Rao, Kevin Lanning, Isabella Vallejo, Lucie Flek, H. Andrew Schwartz, Charles Welch, and Ryan L. Boyd. 2024a. [Archetypes and entropy: Theory-driven extraction of evidence for suicide risk](#). In *Proceedings of the 9th Workshop on Computational Linguistics and Clinical Psychology (CLPsych 2024)*, pages 278–291, St. Julians, Malta. Association for Computational Linguistics.
- Vasudha Varadarajan, Allison Lahnala, Adithya V. Ganesan, Gourab Dey, Siddharth Mangalik, Ana-Maria Bucur, Nikita Soni, Rajath Rao, Kevin Lanning, Isabella Vallejo, et al. 2024b. Archetypes and entropy: Theory-driven extraction of evidence for suicide risk. In *Proceedings of the 9th Workshop on Computational Linguistics and Clinical Psychology (CLPsych 2024)*, pages 278–291.
- Vasudha Varadarajan, Allison C. Lahnala, Sujeeth Vankudari, Syeda Mahwish, Akshay Raghavan, Scott Feltman, Camilo Ruggero, Roman Kotov, and H. Andrew Schwartz. 2025. Linking language-based distortion detection to mental health outcomes. In *Proceedings of The 10th Workshop on Computational Linguistics and Clinical Psychology*, Albuquerque, USA. Association of Computational Linguistics.
- Karishma Vyas, Dominic Murphy, and Neil Greenberg. 2023. [Cognitive biases in military personnel with and without PTSD: a systematic review](#). *Journal of Mental Health*, 32(1):248–259. Publisher: Routledge. eprint: <https://doi.org/10.1080/09638237.2020.1766000>.
- David Watson. 2004. [Stability versus change, dependability versus error: Issues in the assessment of personality over time](#). *Journal of Research in Personality*, 38(4):319–350.
- Ayah Zirikly, Philip Resnik, Özlem Uzuner, and Kristy Hollingshead. 2019. [CLPsych 2019 shared task: Predicting the degree of suicide risk in Reddit posts](#). In *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*, pages 24–33, Minneapolis, Minnesota. Association for Computational Linguistics.

## A Appendix

### A.1 Official Submissions and results

Table A1 shows the official shared task results for Task A.1 (evidence extraction), and Table A2 shows the official results for Task A.2 (wellbeing prediction).

### A.2 Situational 8 DIAMONDS Prompts

We used the items associated with each psychosocial situation from Rauthmann et al. (2014) to define each of the S8D to Deepseek-R1. We individually prompted for each of the S8D and provided two personalized exemplars and annotations for each of the exemplars (2-shot). Some few-shot examples were manually created and some were picked from the CLPsych data. The prompts template would read as follows:

*Instruction: You are an expert in situational perception and psychological analysis. Your task is to evaluate a given block of text for the (insert situation) dimension from the Situational 8 DIAMONDS taxonomy. Individuals who score higher in the (insert situation) dimension relate to the following situations:*

*{{S8D items corresponding to relevant dimension were inserted here}}*

*Your task is to provide the following in a structured JSON format: Rating: Assign a numerical rating for the mating dimension on a scale of 1 to 9 (where 1 = Not at all present and 9 = Highly present).*

*Reasoning: Provide a justification for the rating based on the text. Span Extraction: Identify specific phrases in the text that support your rating.*

*Below are two examples with respective input texts and corresponding outputs to illustrate the task:*

*{{Example texts with corresponding annotations were inserted here}}*

*Now, evaluate the following input text:*

*{{Text requiring annotation was inserted here}}*

We then curated two examples catered towards each situation and hand-annotated the example to

provide as a guideline for output. We provided two messages from the CLPsych 2025 dataset for the situations Duty and Intellect, although for the remaining six S8D, we curated our own examples. An example annotation for the situation Adversity (including manually creates few-shot examples) goes as follows:

"Example 1: "I can't believe how unfair my manager is being. I worked overtime all last week, skipped my breaks, and still got blamed for a project delay that wasn't even my fault. He called me out in front of the entire team, making it sound like I was slacking off. I tried to defend myself, but he just dismissed me and moved on. It's exhausting constantly feeling like I have to prove myself, only to be treated like I'm incompetent."

**\*\*Output:\*\*** "adversity": 8, "reasoning": "The individual is experiencing direct blame from their manager for a project delay that was not their fault. They describe being publicly criticized in front of colleagues and dismissed when attempting to defend themselves. The tone reflects frustration and exhaustion from repeated unfair treatment, which strongly aligns with the Adversity dimension.", "supporting spans": "I worked overtime all last week, skipped my breaks, and still got blamed for a project delay that wasn't even my fault.", "He called me out in front of the entire team, making it sound like I was slacking off.", "I tried to defend myself, but he just dismissed me and moved on.", "It's exhausting constantly feeling like I have to prove myself, only to be treated like I'm incompetent."

Example 2: "Every time I try to express my opinion, my older brother just shuts me down. He talks over me, mocks what I say, and makes me feel like I'm too stupid to contribute. It's like my thoughts don't matter in my own family. Even when I call him out on it, he just laughs and says I'm being too sensitive. I don't know how to get him to take me seriously."

**\*\*Output:\*\*** "adversity": 7, "reasoning":

	A.1 Adaptive/Maladaptive					
	recall			weighted recall		
	overall	adaptive	maladaptive	overall	adaptive	maladaptive
HaRT <sub>WB-FT</sub> + ReLM + Dist (LinSVC)	0.108	0.099	0.116	0.103	0.097	0.109
HaRT <sub>WB-FT</sub> + Log.reg.	0.105	0.077	0.132	0.102	0.075	0.129
HaRT <sub>adaptive-FT</sub>	<b>0.276</b>	<b>0.245</b>	<b>0.308</b>	-	-	-
HaRT <sub>maladaptive-FT</sub>	-	-	-	<b>0.236</b>	<b>0.238</b>	<b>0.235</b>

Table A1: Task A.1 Adaptive and Maladaptive Evidence task. We found that finetuning HaRT for independent adaptive and maladaptive evidence classification can yield significant boosts over traditional principled baselines.

	A.2 Well-being				F1
	timeline MSE ↓				
	overall	minimal impairment	impaired	serious impairment	macro
HaRT <sub>WB-FT</sub>	3.73	2.76	1.95	5.92	0.15
HaRT <sub>WB-FT</sub> + Ridge	3.42	2.95	<b>1.33</b>	5.1	0.17
HaRT <sub>WB-FT</sub> + S8D + ReLM + Dist (Ridge)	3.27	2.63	1.38	4.98	0.19
HaRT <sub>WB-FT</sub> + S8D + PLT (Ridge)	3.22	1.40	2.60	4.86	0.19
S8D + ReLM + Dist (Ridge)	3.02	<b>1.21</b>	1.79	5.25	0.17
S8D + PLT (Ridge)	<b>2.78</b>	1.84	2.14	<b>3.89</b>	0.19

Table A2: Task A.2 Wellbeing task. We found that unlike Task A.1, finetuning HaRT to the wellbeing prediction task need not consistently offer boosts, instead, principled and theoretical methods can offer significant advantages, with a small number of interpretable dimensions without compromising on the accuracy. The first 3 rows in this table were our official submissions while others are presented for additional analysis. We note that using HaRT fine-tuned for the respective adaptive and maladaptive binary classifications may provide benefits over using HaRT<sub>WB-FT</sub> in all combinations, however, due to time constraints we do not have empirical results for the same.

"The individual describes repeated experiences of being dismissed, mocked, and dominated by their older brother. The situation involves verbal criticism, a power imbalance, and an inability to be taken seriously, all of which strongly align with the Adversity dimension.", "supporting spans": "Every time I try to express my opinion, my older brother just shuts me down.", "He talks over me, mocks what I say, and makes me feel like I'm too stupid to contribute.", "Even when I call him out on it, he just laughs and says I'm being too sensitive."

### A.3 Implicit Motives and Mental Health in PLT

To construct our person-level trait (PLT) features, we extracted both implicit motivational needs and core mental health dimensions from participants' language using pre-existing, validated models.

**Implicit Motives.** Following classic motivational theory (McClelland, 1987), we define three core implicit motives reflected in language: (1) the need for achievement, indicated by references to striving for excellence; (2) the need for affiliation, reflected in efforts to initiate or maintain friendly relationships; and (3) the need for power, expressed as influence or control over others or institutions. We used RoBERTa-based models from prior work (Nilsson et al., 2024), trained on expert-coded Picture Story Exercises, to infer these motives at the sentence level. Sentence-level predictions were then aggregated to post level using word count-adjusted averaging procedures.

**Mental Health Dimensions.** We further extracted six features representing key aspects of mental health using models from the Language-Based Assessment Model Library (Nilsson et al., a). These include:

- **Valence:** Trained on annotated Facebook posts rated for emotional positivity or negativity. The model's out-of-sample correlation with human ratings was  $r = .81$ .
- **Harmony in Life & Satisfaction with Life:** Trained on open-text responses rated using validated scales (Kjell et al., 2022a). The models achieved out-of-sample correlations of  $r = .73$  and  $r = .71$ , respectively.
- **Depression:** Two separate models were used — one trained to the Patient Health Questionnaire-9 (PHQ-9; Kroenke et al., 2011), and the other to the Center for Epidemiologic Studies Depression Scale (CES-D; Radloff, 1977). These models yielded correlations of  $r = .66$  and  $r = .73$ , respectively.
- **Anxiety:** Trained on worry-based language mapped to the Generalized Anxiety Disorder 7-item scale (GAD-7), with a correlation of  $r = .63$ .

All models were previously pre-registered for their respective source projects, evaluated using nested cross-validation, and applied out-of-sample to the present dataset. These features form part of our psychologically interpretable PLT baseline.

### A.4 Resilience in PLT

Traditional views of resilience often reduce it to the absence of psychopathology or the ability to recover from stress. However, contemporary psychological frameworks emphasize a broader understanding: resilience as a multidimensional capacity for adaptive functioning in the face of adversity. To operationalize this richer perspective, we used the ReLM (Resilience using Language Modeling) framework (Mahwish et al.), which integrates an archetype-based approach to assess resilience from language.

ReLM captures nine core facets of resilience: optimism, flexibility mindset, sense of social support (SoS), continued activities of daily living (CADL), cognitive reappraisal, emotional maturity, uncertainty tolerance, belief in a higher power, and coping toolkit. Each facet is represented by four prototype statements — brief exemplar sentences derived from a synthesis of resilience literature and analysis of archival interviews with individuals who have demonstrated stability in the face of trauma. For example, a prototype for flexibility mindset reads: "I always try new things because I'm open to exploring."

To assess individual alignment with each facet, ReLM embeds both prototype statements and participant text using Sentence RoBERTa and computes their semantic similarity (see: Varadarajan et al., 2024a). The resulting scores quantify how strongly a participant's language reflects each dimension of resilience. A higher score indicates greater expression of that facet.



Finally, a composite resilience score is computed by applying exploratory factor analysis to the nine facet scores. Across multiple datasets, a single-factor solution consistently explained 49–56% of the variance (Mahwish et al.), supporting the use of a unified resilience metric. This composite score provides a theoretically grounded and interpretable estimate of an individual’s language-based resilience profile.

### A.5 Continued Discussion

The ‘daily lived experience’ facet in this dataset exhibits a strong negative association with well-being. This relationship likely stems from the facet’s focus on individuals persisting through routine tasks despite ongoing stressors. Participant statements such as “I’ve adjusted to my new environment—it was hard at first, but I’m improving” and “I’ve been leaving the house more” (rewritten for anonymity) illustrate gradual adaptation and effort. However, because well-being in this dataset is framed in terms of symptom absence and unimpaired functioning, the ‘daily lived experience’ facet presents a paradox. While it reflects resilience — people continuing daily tasks despite struggles — it also signals underlying difficulty. The very act of pushing forward in the face of these challenges may indicate diminished well-being, as it suggests persistent symptomatology masked by forced functionality.

The Belief in a Higher Power facet reflects trust in external forces during times of struggle, as illustrated by participant statements like “I accept things as they are... I trust that things will get better over time” and “I know that life will work itself out” (anonymized). These responses suggest a surrender of personal control to fate or higher being—a mechanism that may offer emotional relief. However, the model’s negative weighting of this facet stems from a tension between its definition of well-being (rooted in agency, engagement, and lack of symptom) and a resilience strategy that relies on external control. While faith can provide comfort, the framework may be interpreting passive reliance on higher powers as maladaptive in contexts where well-being is tied to active mastery of one’s circumstances.

### A.6 Supplementary results and figures

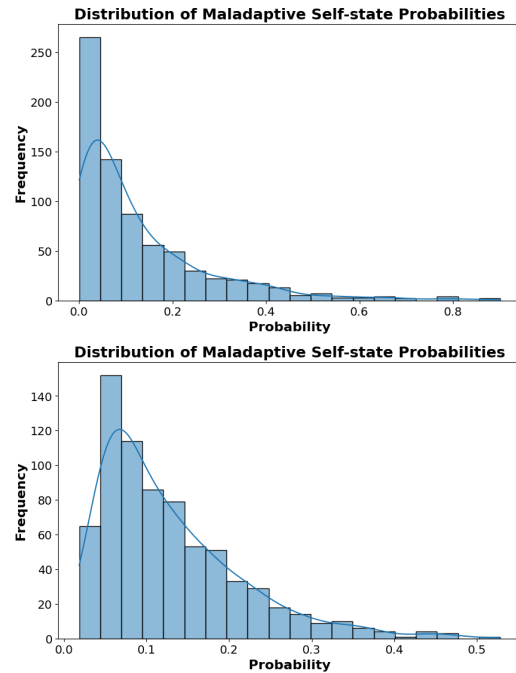


Figure A1: Distribution of probabilities to predict maladaptive state for a given sentence. On the top is using HaRT<sub>WB-FT</sub> + PLT features, and the bottom is using PLT features in Logistic Regression models.




Figure A2: Distribution of probabilities to predicting adaptive state for a given sentence. On the top is using HaRT<sub>WB-FT</sub> + ReLM + Dist<sub>adaptive</sub> features, and the bottom is using ReLM + Dist<sub>adaptive</sub> features in Logistic Regression models.

	<b>Internal Val Set</b>			
	$F1_{macro}$	$F1_{wtd}$	$AUC$	$Acc$
Dist <i>adaptive</i>	0.53	0.85	0.75	0.89
HaRT <i>adaptive</i> -FT	0.67	0.89	0.66	0.89
Dist <i>maladaptive</i>	0.56	0.82	0.76	0.84
HaRT <i>maladaptive</i> -FT	0.60	0.82	0.60	0.82

Table A3: Task A.1 Additional results on internal validation set when predicting binary adaptive and maladaptive labels separately over sentences split from posts.

	<b>5-fold Log.reg. CV</b>			
	$F1_{macro}$	$F1_{wtd}$	$AUC$	$Acc$
Dist <i>adaptive</i>	0.54	0.87	0.75	0.90
HaRT <sub>WB-FT</sub>	0.50	0.87	0.74	0.90
HaRT <sub>WB-FT</sub> + ReLM + Dist <i>adaptive</i>	0.53	0.87	0.76	0.91
PLT <i>adaptive</i>	0.48	0.86	0.66	0.91
HaRT <sub>WB-FT</sub> + PLT <sub>adaptive</sub>	0.52	0.87	0.76	0.90
Dist <i>maladaptive</i>	0.56	0.81	0.73	0.85
HaRT <sub>WB-FT</sub>	0.56	0.81	0.73	0.85
HaRT <sub>WB-FT</sub> + ReLM + Dist <i>maladaptive</i>	0.57	0.82	0.76	0.86
PLT <i>maladaptive</i>	0.49	0.80	0.70	0.86
HaRT <sub>WB-FT</sub> + PLT <sub>maladaptive</sub>	0.58	0.83	0.77	0.86

Table A4: Additional results when training a logistic regression model to predict binary adaptive and maladaptive labels separately over sentences split from posts.




Figure A3: Distribution of probabilities to predicting maladaptive state for a given sentence. On the top is using HaRT<sub>WB-FT</sub> + ReLM + Dist *maladaptive* features, and the bottom is using ReLM + Dist *maladaptive* features in Logistic Regression models.