

CHiRPE: A Step Towards Real-World Clinical NLP with Clinician-Oriented Model Explanations

Stephanie Fong^{1,2}, Zimu Wang^{2,3}, Guilherme C. Oliveira², Xiangyu Zhao², Yiwen Jiang², Jiahe Liu², Beau-Luke Colton¹, Scott Woods⁴, Martha E. Shenton⁵, Barnaby Nelson¹, Zongyuan Ge², Dominic Dwyer^{1,2,*}

¹Orygen and The University of Melbourne ²AIM for Health Lab, Monash University

³University of Liverpool ⁴Yale School of Medicine, Yale University

⁵Brigham and Women’s Hospital, Harvard Medical School

stephanie.fong@unimelb.edu.au, dominic.dwyer@orygen.org.au

Abstract

The medical adoption of NLP tools requires interpretability by end users, yet traditional explainable AI (XAI) methods are misaligned with clinical reasoning and lack clinician input. We introduce CHiRPE (Clinical High-Risk Prediction with Explainability), an NLP pipeline that takes transcribed semi-structured clinical interviews to: (i) predict psychosis risk; and (ii) generate novel SHAP explanation formats co-developed with clinicians. Trained on 944 semi-structured interview transcripts across 24 international clinics of the AMP-SCZ study, the CHiRPE pipeline integrates symptom-domain mapping, LLM summarisation, and BERT classification. CHiRPE achieved over 90% accuracy across three BERT variants and outperformed baseline models. Explanation formats were evaluated by 28 clinical experts who indicated a strong preference for our novel concept-guided explanations, especially hybrid graph-and-text summary formats. CHiRPE demonstrates that clinically-guided model development produces both accurate and interpretable results. Our next step is focused on real-world testing across our 24 international sites.

1 Introduction

Mental illness is a major contributor to the global health burden (McGorry et al., 2025), with psychotic disorders posing particular concern due to their relatively high mortality risk (Walker et al., 2015). Psychosis is often preceded by a prodromal Clinical High-Risk for Psychosis (CHR-P) phase (Wang et al., 2022), during which early detection can enable interventions that may delay or prevent progression, improving long-term outcomes.

The PSYCHS (Positive SYmptoms and Diagnostic Criteria for the CAARMS-Harmonized SIPS) is

a clinician-administered, semi-structured interview (Woods et al., 2024) for detecting CHR-P. However, its routine use is constrained by substantial time and administrative effort (Oliver et al., 2022). To address these limitations, recent advances in NLP offer a viable path forward. Automated analysis of interview transcripts can support scalable and consistent screening of disorder markers, thereby alleviating the reliance on labour-intensive manual assessment (Molina et al., 2025; Na et al., 2025). However, their “black-box” nature significantly impedes adoption in clinical contexts where transparency and interpretability are essential (Topol, 2019; Cinà et al., 2022).

Post-hoc explainable AI (XAI) techniques, such as SHapley Additive exPlanations (SHAP; Lundberg and Lee, 2017), have been proposed to bridge this gap. In NLP, SHAP assigns each input feature (e.g. word or token) a value representing its average influence on the model’s prediction (Shapley et al., 1953). While these attributions capture the model’s logic, they fundamentally misalign with real-world clinician reasoning (Lawrence et al., 2024). In parallel, large language models (LLMs) remain opaque (Lawrence et al., 2024) and sensitive to subtle prompts or context changes (Peng et al., 2023; Zhou et al., 2023; Chatterjee et al., 2024). Despite the importance of interpretability for clinical trust (Di Martino and Delmastro, 2023; Talebi et al., 2024), clinician involvement in XAI design remains limited (Ghassemi et al., 2021).

Motivated by this gap, we introduce CHiRPE (Clinical High-Risk Prediction with Explainability), a human-centred NLP framework that processes transcribed PSYCHS interviews to classify individuals as CHR-P and generates five SHAP explanation formats, three of which are co-developed with clinicians (see Figure 1). Trained on 943 tran-

*Corresponding author.

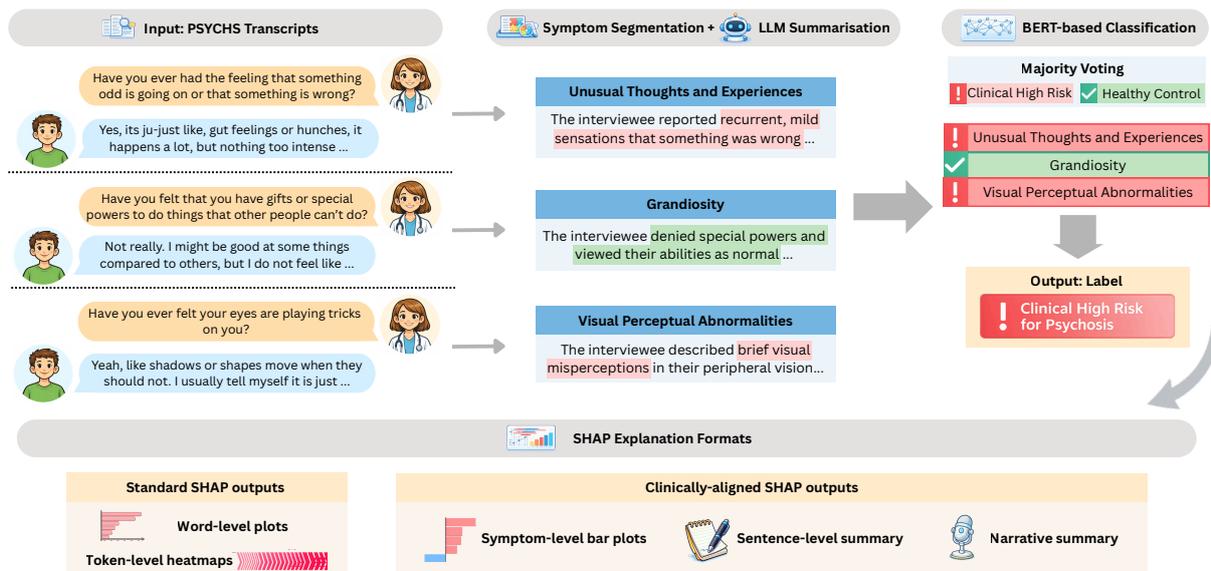


Figure 1: CHiRPE pipeline. Raw PSYCHS transcripts are segmented by symptom domain and summarised, then passed to a BERT-based classifier. The system outputs a CHR-P or Healthy label alongside SHAP explanations.

scripts across 24 international sites, the pipeline integrates symptom domain mapping, LLM summarisation, BERT-based classification, and SHAP-based explanations. Our key contributions are as follows: (i) an integrated pipeline CHR prediction that aligns model inputs with clinically meaningful constructs while preserving high classification accuracy; (ii) co-designed SHAP explanation formats grounded in how clinicians reason about psychosis risk; (iii) empirical clinician expert evaluation of these co-designed SHAP formats, showing that they outperform standard explanations in interpretability, clinical reasoning, and alignment.

2 Related Work

A systematic review of nine studies using shallow classifiers with handcrafted features or static embeddings reported CHR-P diagnostic accuracies between 56-95% (Molina et al., 2025), but most relied on small, single-site samples and internal validation, raising concerns about overfitting and generalisability. Recent AMP-SCZ work using Llama 3 “normality” features and Naïve Bayes achieved moderate accuracy ($AUC \approx 0.68$) on a smaller subset of the same PSYCHS interview data (Bilgrami et al., 2025), yet lacked interpretability.

Existing results demonstrate that the shift to deep learning enables richer linguistic modelling but introduces opacity, necessitating interpretability methods. While SHAP formats have shown promise in broader medical applications such as cerebral infarction (Nohara et al., 2022) and tho-

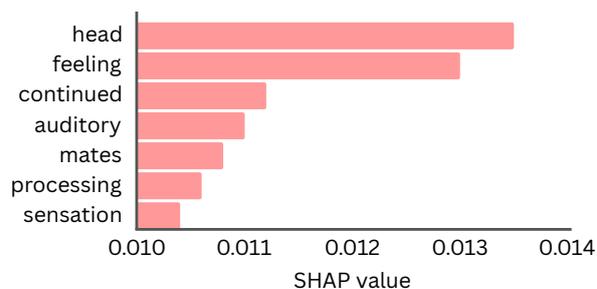


Figure 2: SHAP word-level bar plot showing the top contributing words for a CHR-P classification decision.

racic surgery (Hur et al., 2025), mental health NLP studies remain limited (Baki et al., 2022; Nie and Wu, 2025). Standard SHAP word-level bar plots (Figure 2) and token-level heatmaps (Figure 4) either present decontextualised words in isolation or span entire transcripts, which are both misaligned with the types of concept-driven summaries used in clinical reasoning to make decisions. To our knowledge, no mental health NLP studies have examined adapted SHAP formats that align with clinical reasoning, and it remains a critical gap that is preventing the clinical use of NLP tools.

3 CHiRPE

The CHiRPE pipeline is presented in Figure 1. It consists of the following: a) concept mapping to generate transcript segments; b) summarisation of concept segments; c) CHR detection classifier training; d) application of clinically-informed SHAP approaches. In this section, we introduce each step

in detail.

3.1 Symptom Domain Segmentation

The PSYCHS interview follows a fixed order of standardised questions organised into 15 attenuated positive symptom domains (see Appendix A). Interviewer utterances were mapped to domains using fuzzy string matching¹ with a threshold of 80% (justified in Appendix B).

3.2 Summarisation of Interview Segments

Each interview segment was rephrased into third-person using a two-step process with Mistral-7B-Instruct-v0.3² to better match BERT pretraining data. The process involved an initial clinician style rewrite followed by refinement for completeness and coherence (see Appendix C for prompts and Appendix F for prompt sensitivity analysis).

3.3 Classification Models

We fine-tuned BERT (Devlin et al., 2019), ClinicalBERT (Alsentzer et al., 2019), and MentalBERT (Ji et al., 2021), domain-specific models for clinical and mental health language, to classify each summary symptom segment as CHR-P or control. Segments exceeding 512 tokens were split into fixed-length chunks, and transcript-level labels were determined by majority voting across segments.

3.4 Explanation Generation with SHAP

As baselines, the standard SHAP visualisations of word-level bar plots for top impact words and colour-coded heatmaps were generated, as shown in Figures 2 and 4. In consultation with clinicians, we introduced new presentation formats, including sentence-level summaries, narrative explanations, and symptom-level plots.

Sentence-level Summary. For individuals classified as CHR, the sentence with the highest average net SHAP contribution toward a CHR prediction was extracted as a *Sentence-level Summary*. For example,

P4 Ideas of Guilt

They admit to constantly dwelling on past problems, which significantly impact their daily life.

¹<https://pypi.org/project/fuzzywuzzy/>

²<https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.3>

Narrative Summary. A concise *Narrative Summary* was generated using Qwen3-4B (Yang et al., 2025) for its strong contextual generation ability. Each summary included ± 1 sentence of context and a representative interviewee quote. The prompt is shown in Appendix E. An example summary is as follows:

P4 Ideas of Guilt

The interviewee reports persistent and intrusive thoughts centered on moral concerns and others' perspectives, particularly during emotional distress. They describe frequent rumination on past issues, which interferes with daily functioning, and recurrent episodes of intense, unexplained guilt occurring at least weekly.

“Even if it’s not actively thinking about the problem, like, ‘Oh, I wish I’d done something different,’ it affects little things in life that make you think about it.”

Symptom-Level Plots. SHAP values were aggregated into mean net scores per symptom domain and visualised as horizontal bar charts (Figure 3), showing each domain’s influence on CHR-P or control prediction.

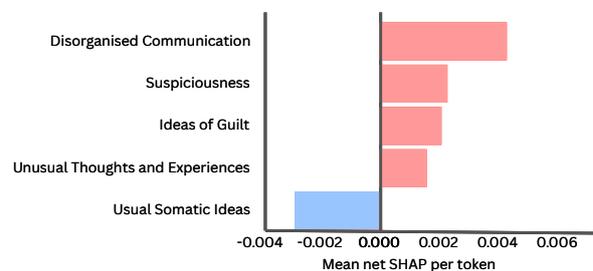


Figure 3: Symptom-level plots (red: CHR-P, blue: healthy control).

4 Data Analysis Methods

4.1 Performance Metrics

Accuracy, precision, recall, F1, and AUC were evaluated at the transcript level on the test set.

4.2 Clinical Expert Feedback

Clinical experts completed a mixed-method questionnaire on SHAP explanation preferences, including: a) word-level bar plots; b) colour-coded heatmaps; c) symptom-level bar plots; d) sentence-level summary; and e) narrative summaries (full questionnaire in Appendix H). Descriptive and

Model	Baseline				Summary Only				Segmentation Only				Proposed (Summ + Seg)			
	Acc	F1	Prec	Rec	Acc	F1	Prec	Rec	Acc	F1	Prec	Rec	Acc	F1	Prec	Rec
BERT	83.91	89.78	83.11	97.62	81.03	88.09	80.79	96.83	91.81	95.10	93.15	97.14	91.23	94.50	96.99	92.14
ClinicalBERT	81.03	88.00	81.21	96.03	82.76	88.97	82.88	96.03	87.51	92.56	91.97	93.21	90.64	94.20	95.59	92.86
MentalBERT	83.91	89.85	82.67	98.41	84.48	90.04	84.14	96.83	91.23	94.77	92.52	97.14	91.23	94.85	91.39	98.57

Table 1: Performance metrics of BERT-based models across ablation settings. Best results for each metric within each model are highlighted in bold.

inferential statistics were used to compare inter-pretability ratings across formats.

5 Data and Experimental Setup

5.1 Dataset

Participants were drawn from the Accelerating Medicines Partnership Schizophrenia³ (AMP-SCZ) study, focusing on CHR-P identification and transition to psychosis (Wannan et al., 2024). The dataset comprised 943 English PSYCHS transcripts from 581 unique participants aged 12-30 (M = 20.9, SD = 4.1), 63.3% of whom were female. Labels were assigned by trained research assistants, with 83.6% CHR-P and 16.4% Healthy Controls across 24 international sites of the AMP-SCZ study.

The dataset was partitioned into 64% for training, 16% for development (via nested 5-fold cross-validation on an 80% subset), and 20% for held-out testing. Splits were stratified by CHR status and grouped by participant ID to prevent data leakage. A fixed random seed ensured consistent splits across model variants.

5.2 Pre-processing

Audio recordings were human transcribed⁴, with all timestamps and personally identifiable information removed. A few-shot XLM-RoBERTa classifier (Conneau et al., 2020) distinguished interviewer and interviewee turns.

5.3 Models and Hyperparameters

BERT-based models were trained on a single A100 GPU within a high-performance computing cluster. All classification models were fine-tuned using weighted cross-entropy loss to address class imbalance. Hyperparameters, including learning rate, batch size, weight decay, and number of epochs, were optimised via inner-loop grid search (details in Appendix D).

³<https://www.ampscz.org/>

⁴<https://www.transcribeme.com/>

6 Results

6.1 Model Classification Performance

All three transformer-based models performed strongly on the held-out test set (Table 1; optimal hyperparameters in Appendix D), outperforming baselines without segmentation and summarisation. AUCs were 0.95 for BERT and 0.97 for both ClinicalBERT and MentalBERT, exceeding baseline AUCs of 0.94, 0.90, and 0.95, respectively.

6.2 Ablation Analysis

6.2.1 Segmentation and Summarisation

Ablation results show that symptom domain segmentation is the main driver of performance gains, while summarisation provides smaller but consistent improvements (see Table 1). In particular, segmentation reduces noise from long, heterogeneous transcripts and focuses the classifier on clinically meaningful information. By contrast, summarisation maintains comparable accuracy despite substantially reducing input length, improving modelling efficiency.

In the CHiRPE pipeline, segmentation and summarisation together achieve the best overall trade-off between F1, precision, and efficiency. The two components appear complementary: segmentation imposes clinically meaningful structure, while summarisation standardises and condenses content for SHAP attribution and downstream narrative generation, yielding both improved predictive performance and a more scalable explanation workflow.

6.2.2 Precision–Recall Trade off

As shown in Table 2, while the proposed setting (summarisation + segmentation) improves F1 and precision, it yields a modest drop (-4%) in recall and is accompanied by a substantial gain (+41%) in specificity. From a modelling perspective, this pattern suggests reduced bias from class imbalance (CHR = 83.6%, HC = 16.4%) and less overfitting to CHR cases. Clinically, given that CHR reflects risk rather than diagnosis, and that false positives

carry nontrivial clinical costs (De Pablo et al., 2021; Di Lisi et al., 2022), this shift indicates a more conservative and better calibrated decision boundary.

	Recall (CHR)	Specificity (HC)
Baseline	0.9603	0.3958
Proposed (Summ + Seg)	0.9286	0.8065

Table 2: Recall–specificity trade off for ClinicalBERT under ablation settings.

6.3 Clinical Expert Evaluation of Explanation Formats

Twenty-eight clinical experts, including psychiatrists, clinical psychologists, and trained research assistants, completed a 10-minute questionnaire evaluating the interpretability of SHAP outputs on a 5-point Likert scale (Table 3). Responses were anonymous, and participants gave informed consent for data usage (Appendix H).

	Explanation Format	Interpretability
Baseline	Word-level plots	2.18
	Token-level heatmaps	2.50
Proposed	Symptom bar plots	3.71
	Sentence-level summary	4.25
	Single narrative	4.08
	Graph + single narrative	4.32
	Multiple narratives	4.21

Table 3: Mean interpretability ratings for baseline and proposed explanation formats (n = 28).

6.3.1 Quantitative Clinical Expert Feedback

All novel SHAP formats outperformed the traditional SHAP word-level bar plots and heatmaps. The combination of symptom-level plots with narrative summaries received the highest interpretability ratings. A repeated-measures ANOVA confirmed significant differences across formats, $F(6, 28) = 26.485$, $p < .001$. Full ANOVA results are presented in Appendix I.

Additional ratings on Clinical Reasoning and Clinical Intuition Alignment showed the same pattern of results, with proposed formats consistently outperforming baseline visualisations. Full quantitative results are reported in Appendix I.

6.3.2 Qualitative Clinical Expert Feedback

Qualitative coding of 20 clinical expert comments identified eight themes, most commonly favouring

hybrid graph-text formats and concise presentations. Representative quotes are provided in Appendix I.

Theme	n	%
Graph + text summaries preferred	6	30
Concise formats (dot points, fragments)	4	20
Structured symptom tables needed	3	15
Clear definition of experiences	2	10
Clarify model inputs and outputs	2	10
Overview across multiple symptoms	1	5
Add risk severity scale for decisions	1	5
Data access and governance issues	1	5

Table 4: Themes from qualitative feedback.

7 Conclusion and Future Work

In conclusion, CHiRPE outperformed baseline models in classification and surpassed existing SHAP approaches in interpretability. The preference for hybrid graph–text summaries and concise text formats highlighted CHiRPE’s practical contributions to future explainable clinical AI designs. CHiRPE addresses a key gap in mental health NLP in light of evolving AI regulations related to interpretability (e.g., the EU AI Act).

While recent AMP-SCZ work using Llama 3 features achieved moderate accuracy (AUC \approx 0.68) compared to CHiRPE’s higher performance (AUC $>$ 0.95) (Bilgrami et al., 2025), differing methods and limited reporting in evaluation metrics preclude direct comparison. Prior work has also shown performance gains from domain-specific BERT models in mental health tasks (Ji et al., 2021; Turchin et al., 2023), but CHiRPE achieved comparable accuracy without task-specific pretraining. If replicated, this observation potentially suggests a need to shift focus from continual fine-tuning to aligning models with clinical reasoning, particularly in areas where diagnoses rely on subjective narratives rather than objective biomarkers (Oliver et al., 2022).

Future work will focus on refining CHiRPE’s text-based summaries and symptom-level visualisations and piloting an interface prototype (Figure 5 in Appendix J) across our network of 24 international sites. We will also incorporate patient perspectives to enhance concept mappings and extend the pipeline to additional mental health conditions, such as depression, mania, and anorexia.

Limitations

This study was limited to English transcripts collected from the 24 international sites. Future work

could extend concept mapping and model evaluation to multilingual data for broader applicability. Additional recruitment of participants in future studies could also strengthen the robustness and generalisability of findings. It is also worth exploring alternative models that are not BERT-based for this task.

Ethical Considerations

The transcript data used in this study were obtained from the Accelerating Medicines Partnership Schizophrenia (AMP-SCZ) consortium under a data use agreement. All data were de-identified and collected with informed consent for research purposes. Recruitment and consent procedures were IRB-approved (*Accelerating Medicines Partnership, 2023*). Participants will be reimbursed for their time and expenses associated with completing the research assessments. The amount of reimbursement varies by data collection site and the extent of study completion.

We used the pre-trained BERT model released by Google under the Apache License 2.0⁵ (*Devlin et al., 2019*), ClinicalBERT under the MIT license⁶ (*Alsentzer et al., 2019*), and MentalBERT under the Creative Commons Attribution Non Commercial 4.0 license⁷ (*Ji et al., 2021*). All models were fine-tuned on AMP-SCZ data solely for research purposes, in accordance with their respective licenses.

The development of explainable AI tools for sensitive mental health settings raises important ethical and privacy challenges, particularly around informed consent. In clinical contexts where ambient data capture and AI-driven interpretation are involved, ethical safeguards cannot be assumed—they must be explicitly designed, tested, and adapted to the needs of each case. As part of our design process, we are co-developing CHiRPE's consent procedures with both patients and clinicians to ensure transparency, agency, and appropriateness. This form of human-centred ethical design is expected to further enhance model performance, explanatory capacity, and ultimately the trust required for CHiRPE to contribute to the long-term sustainability of AI technologies in mental health care.

⁵<https://www.apache.org/licenses/LICENSE-2.0>

⁶<https://tlo.mit.edu/understand-ip/exploring-mit-open-source-license-comprehensive-guide>

⁷<https://creativecommons.org/licenses/by-nc/4.0/deed.en>

Data and Code Availability

The AMP-SCZ dataset used in this study is available to researchers through the National Data Archive of the National Institute of Health in the USA: <https://www.ampscz.org/scientists/data/>. All code used for data preprocessing, model training, and explanation generation is available at <https://github.com/stephaniesyfong/CHiRPE>.

Use of AI

ChatGPT and GitHub Copilot were used to assist with code debugging and language editing. All outputs were manually reviewed and verified by the authors.

Acknowledgements

The Accelerating Medicines Partnership Schizophrenia (AMP-SCZ) is a public-private partnership managed by the Foundation for the National Institutes of Health. The AMP-SCZ research program is supported by contributions from the AMP-SCZ public and private partners, including NIMH (U24MH124629, U01MH124631, and U01MH124639) and Wellcome (220664/Z/20/Z and 220664/A/20/Z). The views expressed in this article are personal views of the authors and may not be understood or quoted as being made on behalf of or reflecting the position of the Department of Health and Human Services, including the National Institutes of Health and the US Food and Drug Administration, the United States Government, or the European Medicines Agency or one of its committees or working parties. The research was supported by the University of Melbourne's Research Computing Services and the Petascale Campus Initiative. The authors report no conflict of interest. Dwyer was supported by an NHMRC Emerging Leadership II Investigator Grant (#2034943).

References

Accelerating Medicines Partnership. 2023. Accelerating medicines partnership@ schizophrenia observational study protocol. https://cdn.clinicaltrials.gov/large-docs/03/NCT05905003/Prot_000.pdf. Accessed May 18, 2025.

Emily Alsentzer, John R Murphy, Willie Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew

- McDermott. 2019. Publicly available clinical bert embeddings. *arXiv preprint arXiv:1904.03323*.
- Pinar Baki, Heysem Kaya, Elvan Çiftçi, Hüseyin Güleç, and Albert Ali Salah. 2022. A multimodal approach for mania level prediction in bipolar disorder. *IEEE Transactions on Affective Computing*, 13(4):2119–2131.
- Zarina Bilgrami, Cheryl Corcoran, Guillermo Cecchi, and Phillip Wolff. 2025. Leveraging ai and language analysis to predict psychosis risk in clinical high-risk individuals. *Biological Psychiatry*, 97(9):S68–S69.
- Anwoy Chatterjee, HSVNS Kowndinya Renduchintala, Sumit Bhatia, and Tanmoy Chakraborty. 2024. Posix: A prompt sensitivity index for large language models. *arXiv preprint arXiv:2410.02185*.
- Giovanni Cinà, Tabea Röber, Rob Goedhart, and Ilker Birbil. 2022. Why we do need explainable ai for healthcare. *arXiv preprint arXiv:2206.15363*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 8440–8451.
- Gonzalo Salazar De Pablo, Joaquim Radua, Joana Pereira, Iliaria Bonoldi, Vincenzo Arienti, Filippo Bessana, Livia Soardo, Anna Cabras, Lydia Fortea, Ana Catalan, and 1 others. 2021. Probability of transition to psychosis in individuals at clinical high risk: an updated meta-analysis. *JAMA psychiatry*, 78(9):970–978.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.
- Alessandro Di Lisi, Simona Pupo, Marco Menchetti, and Lorenzo Pelizza. 2022. Antipsychotic treatment in people at clinical high risk for psychosis: a narrative review of suggestions for clinical practice. *Journal of Clinical Psychopharmacology*, pages 10–1097.
- Flavio Di Martino and Franca Delmastro. 2023. Explainable ai for clinical and remote health applications: a survey on tabular and time series data. *Artificial Intelligence Review*, 56(6):5261–5315.
- Marzyeh Ghassemi, Luke Oakden-Rayner, and Andrew L Beam. 2021. The false hope of current approaches to explainable artificial intelligence in health care. *The Lancet Digital Health*, 3(11):e745–e750.
- Sujeong Hur, Yura Lee, Joongheum Park, Yeong Jeong Jeon, Jong Ho Cho, Duck Cho, Dobin Lim, Wonil Hwang, Won Chul Cha, and Junsang Yoo. 2025. Comparison of shap and clinician friendly explanations reveals effects on clinical decision behaviour. *npj Digital Medicine*, 8(1):578.
- Shaoxiong Ji, Tianlin Zhang, Luna Ansari, Jie Fu, Prayag Tiwari, and Erik Cambria. 2021. Mentalbert: Publicly available pretrained language models for mental healthcare. *arXiv preprint arXiv:2110.15621*.
- Hannah R Lawrence, Renee A Schneider, Susan B Rubin, Maja J Matarić, Daniel J McDuff, and Megan Jones Bell. 2024. The opportunities and risks of large language models in mental health. *JMIR Mental Health*, 11(1):e59479.
- Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30.
- Patrick McGorry, Hasini Gunasiri, Cristina Mei, Simon Rice, and Caroline X Gao. 2025. The youth mental health crisis: analysis and solutions. *Frontiers in Psychiatry*, 15:1517533.
- José Tomás García Molina, Pablo A Gaspar, and Alicia Figueroa-Barra. 2025. Automatic speech recognition in psychiatric interviews: a rocket to diagnostic support in psychosis. *Revista Colombiana de Psiquiatría (English ed.)*, 54(4):624–631.
- Hongbin Na, Yining Hua, Zimu Wang, Tao Shen, Beibei Yu, Lilin Wang, Wei Wang, John Torous, and Ling Chen. 2025. A survey of large language models in psychotherapy: Current landscape and future directions. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 7362–7376.
- Hui Nie and Xiaoyan Wu. 2025. A dual-channel prediction-interpretation framework with pre-trained language models and shap explainability. *Journal of Computer and Communications*, 13(3):116–137.
- Yasunobu Nohara, Koutarou Matsumoto, Hidehisa Soejima, and Naoki Nakashima. 2022. Explanation of machine learning models using shapley additive explanation and application for real data in hospital. *Computer Methods and Programs in Biomedicine*, 214:106584.
- Dominic Oliver, Maite Arribas, Joaquim Radua, Gonzalo Salazar de Pablo, Andrea De Micheli, Giulia Spada, Martina Maria Mensi, Magdalena Kotlicka-Antczak, Renato Borgatti, Marco Solmi, and 1 others. 2022. Prognostic accuracy and clinical utility of psychometric instruments for individuals at clinical high-risk of psychosis: a systematic review and meta-analysis. *Molecular Psychiatry*, 27(9):3670–3678.
- Hao Peng, Xiaozhi Wang, Jianhui Chen, Weikai Li, Yun-jia Qi, Zimu Wang, Zhili Wu, Kaisheng Zeng, Bin Xu, Lei Hou, and 1 others. 2023. When does in-context learning fall short and why? a study on specification-heavy tasks. *arXiv preprint arXiv:2311.08993*.

- Lloyd S Shapley and 1 others. 1953. A value for n-person games.
- Salmon Talebi, Elizabeth Tong, Anna Li, Ghiam Yamin, Greg Zaharchuk, and Mohammad RK Mofrad. 2024. Exploring the performance and explainability of fine-tuned bert models for neuroradiology protocol assignment. *BMC Medical Informatics and Decision Making*, 24(1):40.
- Eric J Topol. 2019. High-performance medicine: the convergence of human and artificial intelligence. *Nature medicine*, 25(1):44–56.
- Alexander Turchin, Stanislav Masharsky, and Marinka Zitnik. 2023. Comparison of bert implementations for natural language processing of narrative medical documents. *Informatics in Medicine Unlocked*, 36:101139.
- Elizabeth Reisinger Walker, Robin E McGee, and Benjamin G Druss. 2015. Mortality in mental disorders and global disease burden implications: a systematic review and meta-analysis. *JAMA psychiatry*, 72(4):334–341.
- Peng Wang, Chuan-Dong Yan, Xiao-Jie Dong, Lei Geng, Chao Xu, Yun Nie, and Sheng Zhang. 2022. Identification and predictive analysis for participants at ultra-high risk of psychosis: A comparison of three psychometric diagnostic interviews. *World journal of clinical cases*, 10(8):2420.
- Cassandra MJ Wannan, Barnaby Nelson, Jean Addington, Kelly Allott, Alan Anticevic, Celso Arango, Justin T Baker, Carrie E Bearden, Tashrif Billah, Sylvain Bouix, and 1 others. 2024. Accelerating medicines partnership® schizophrenia (amp® scz): rationale and study design of the largest global prospective cohort study of clinical high risk for psychosis. *Schizophrenia bulletin*, 50(3):496–512.
- Scott W Woods, Sophie Parker, Melissa J Kerr, Barbara C Walsh, S Andrea Wijtenburg, Nicholas Prunier, Angela R Nunez, Kate Buccilli, Catalina Mourgues-Codern, Kali Brummitt, and 1 others. 2024. Development of the psychs: Positive symptoms and diagnostic criteria for the caarms harmonized with the sips. *Early Intervention in Psychiatry*, 18(4):255–272.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Wenxuan Zhou, Sheng Zhang, Hoifung Poon, and Muhao Chen. 2023. Context-faithful prompting for large language models. *arXiv preprint arXiv:2303.11315*.

A PSYCHS template questions and symptom-domain matching

Sample list of P1–P3 standard template questions

P1 Unusual Thoughts and Experiences

- Have you ever had the feeling that something odd is going on or that something is wrong?
- Have you ever been confused at times whether something you have experienced is real or imaginary?
- Have you ever daydreamed a lot or found yourself preoccupied with stories, fantasies, or ideas?
- Has your experience of time ever seemed to have changed? Has it become unnaturally faster or unnaturally slower?
- Have you ever seemed to live through events exactly as you have experienced them before?
- Have familiar people or surroundings ever seemed strange?
- Have you felt that you or others or the world have changed in some way?
- Have you ever felt that you might not actually exist? Or that the world might not exist?
- Have you ever felt you can predict the future?
- Have you felt that things that were happening around you had a special meaning just for you?
- Have you ever felt the radio or TV or other electronic devices are communicating directly with you?
- Do you know what it means to be superstitious? Have you been superstitious?
- Have you ever felt that some person or force outside yourself has been controlling or interfering with your thoughts, feelings, actions or urges?
- Have you ever felt that ideas or thoughts that are not your own have been put into your head? Or that your own thoughts have been taken out of your head?
- Have your thoughts ever been broadcast so that other people know what you are thinking? Or ever said out loud so that other people can hear them?
- Have you ever thought that people might be able to read your mind? Or that you could read other people's minds?

P2 Suspiciousness

- Have you ever felt like people have been talking about you, laughing at you or thinking about you in a negative way?
- Have you ever found yourself feeling mistrustful or suspicious of other people?
- Have you ever felt that you have to pay close attention to what's going on around you in order to feel safe?
- Have you ever felt like you are being singled out or watched?
- Has anybody been giving you a hard time or trying to hurt you? Do you have a sense of who that might be? Do you feel they have hostile or negative intentions?

P3 Unusual Somatic Ideas

- Have you ever worried that something might be wrong with your body, your health, or a part of your body? Have you thought that it seems different to others in some way?
- Have you worried about your body shape?
- Have you ever worried that something odd is going on with your body that you can't explain?

Domain	Template Question	Matched Utterance (segment)
P1	Have you experienced odd or unusual beliefs that other people find strange?	Interviewer: I'm wondering if you've ever held beliefs that others might see as a bit unusual. Interviewee: [silence] Interviewer: For instance, do you ever feel like everyday events—say, a news headline—carry a secret message meant just for you? Interviewee: Sometimes I feel like casual conversations on TV are speaking directly to me.
P2	Have you ever heard or seen things that other people couldn't perceive?	Interviewer: Do you ever notice sounds or sights that others around you don't seem to pick up on? Interviewee: [silence] Interviewer: Like hearing soft voices in an empty room, or glimpsing shadows that no one else sees? Interviewee: Yes, sometimes I catch voices calling my name in an empty room.

Table 5: Symptom–question mapping example for PSYCHS domains P1 and P2 with interviewer segments.

B Segmentation threshold with Fuzzy Matching

A random subset of fifty transcripts was manually segmented into the 15 PSYCHS symptom domains by two human expert raters. The resulting gold-standard segments were compared with automatically generated ones using fuzzy matching at similarity thresholds of 70%, 80%, and 90%. The 80% threshold achieved best performance (see Table below) and was therefore used in our subsequent analyses.

Threshold	Precision	Recall	F1
70%	0.628	0.512	0.341
80%	0.869	0.881	0.817
90%	0.869	0.876	0.817

Table 6: Averaged Macro Performance at different thresholds.

C Prompt for Rephrasing and Summarisation of Transcripts

First Pass (Initial Draft Prompt)

You are an expert clinical interviewer. Summarise the following interview segment in a single third person paragraph, covering what was asked and the detailed response.

Interview segment: *<segment>*

Draft summary:

Second Pass (Refinement Draft Prompt)

Here is a transcript segment and an initial draft summary. Improve the summary by adding any important information from the segment that was missed. Keep third person narration, one coherent paragraph, and no bullet points.

Interview segment: *<segment>*

Draft summary: *<draft>*

Improved summary:

D Hyperparameter tables

We performed a grid search over the following ranges: learning rate 1e-5, 2e-5, batch size 8, 16, number of epochs 2, 3, and weight decay 0.0, 0.01 for all models.

Model	Learning Rate	Batch Size	Epochs	Weight Decay
BERT	2e-5	8	3	0.01
ClinicalBERT	2e-5	8	3	0.01
MentalBERT	2e-5	8	3	0.01

Table 7: Best hyperparameters found via grid search.

E Prompt for Generating Narrative Summaries

Description based on Excerpt Summary:

You are an expert clinical interviewer. Rewrite the excerpt into ONE clinician-friendly paragraph (max 3 sentences) describing *symptom*.

Excerpt: *<segment>*

Description based on Excerpt Summary:

Provide ONLY the interviewee quote (enclosed in double quotation marks) that clearly illustrates *symptom* and supports "anchor". Output the quote and nothing else.

Transcript: *<segment>*

Quote:

F Prompt Sensitivity for Transcript Summarisation

Prompt sensitivity was evaluated with 3 different prompts, varying on role framing (experienced clinician, clinical assessor, clinical note writer), instruction specificity (ask simply to “summarise” versus explicitly specify including both questions and responses or improving clarity), and emphasis (conciseness, completeness, accuracy and structure). Summaries remained semantically consistent across prompts, with BERTScore F1 = 0.7558 and SentenceBERT cosine = 0.7758.

G SHAP-based stability

We also evaluated the factual consistency and semantic quality of Qwen3-4B narrative summaries derived from symptom domains with highest SHAP attributions. Factual consistency was assessed using an NLI-based faithfulness check, showing 6.76% contradicted statements. Semantic similarity remained high, with SentenceBERT cosine = 0.7100 and BERTScore F1 = 0.8883 (precision = 0.8754, recall = 0.9019).

H Questionnaire for Explanation Feedback

Background Information and Data Collection Purpose

About Chirpe

In Orygen, our team produces Artificial Intelligence (AI) devices that enhance human connection and identify consumers who need specialised care. We have developed an AI called "Chirpe" (Clinical High Risk Prediction with Explainability). Chirpe is able to identify young people at clinical high-risk for psychosis (CHR) using recorded speech of the PSYCHS questionnaire at accuracies above 90%. It's very exciting because it means that we may not need to rate the PSYCHS in the future.

Why your feedback matters

Chirpe sometimes struggles with explaining why the person is CHR and that's where we need your help. We have designed different ways that Chirpe can explain their decisions and we'd like you to let us know what style of communication you prefer. We'd be really grateful for any other feedback too as you are the experts.

Contribute to research

If you want to be involved in this paper as a clinical advisor, please let us know. We will also use your responses to design better ways that CHIRPE can talk with you.

Description of task

Comprehensibility of Output Formats

We are interested in how we can explain Chirpe's predictions to clinicians and have created 7 different options to choose from. Think about whether you understand why Chirpe made the decision: Is it comprehensible?

Rating of Explainability Outputs

Clinical experts were asked to evaluate the interpretability of the following seven explanation formats. For each, they rated interpretability on a 5-point Likert scale (1 = Not Interpretable; 5 = Very Interpretable).

1. **Word-level SHAP plots**
Longer bars in red indicate stronger influence towards CHR prediction (see Figure 2)
2. **Highlighting text in the transcript or summary**
Word highlights show influence (darker = stronger): red for CHR, blue for Control (see Figure 4)
3. **Symptom-level Plots (P-items)**
Red bars = CHR; Green bars = Controls. (See Figure 3 in Main Text)
4. **Selected sentences from the most influential symptom categories**
These sentences contain words that are most contributing to the CHR prediction. (See **Sentence-level Summary** in Main Text)
5. **Single clinical narrative** (See **Narrative Summary** in Main Text)
6. **Symptom-level Plots + Single clinical narrative** (See **Narrative Summary** in Main Text)
7. **Multiple narrative summaries across symptom categories**

Qualitative Feedback In addition to the above ratings, clinicians responded to the following open-ended questions:

- Would you suggest any changes to how the graphs are presented?
- Would you suggest any changes to how text information is presented?
- Do you have any other comments on how CHIRPE's explanations can be improved?

I Clinician Ratings of Explanation Formats

I.1 ANOVA Results

A one-way repeated-measures ANOVA revealed a significant effect of explanation format on perceived comprehensibility, $F(6, 28) = 26.485, p < .001$. Holm-adjusted pairwise t -tests showed that Options 4, 5, 6, and 7 were all rated significantly higher than Options 1 and 2 ($p < .001$). No significant differences were found between the top-rated formats (Options 4–7).

I.2 Additional Evaluation with Clinical Reasoning and Alignment

Two additional questions regarding clinical reasoning and alignment of SHAP explanation formats were evaluated using a 7-point Likert Scale.

1. Clinical Reasoning: “To what extent does this explanation format align with how you would understand symptom information?”
2. Clinical Intuition Alignment: “Does the explanation format present information in a way that is clinically intuitive?”

Explanation Format	Reasoning	Alignment
Word-level plots	2.0	1.8
Token-level Heatmaps	2.4	2.4
Symptom bar plots	4.2	4.4
Sentence-level summary	5.4	5.6
Single narrative	5.8	5.6
Graph + Single narrative	6.0	6.2
Multiple narratives	6.0	5.6

Table 8: Mean and SD of Comprehensibility (n = 5)

Across these new evaluations, the proposed formats from this paper remained the highest rated, consistent with our earlier findings.

I.3 Clinical Expert Quotes from Qualitative Feedback

- **Graph plus text summaries preferred**

Clinical experts consistently valued mixed formats. One comment noted that “a graph combined with select sentences would be most useful” and another stated “I think a mixture of option six and seven would be perfect.”

- **Concise formats**

Several clinical experts preferred brief

explanations, e.g. “Can be more concise. No need to be complete sentences” and “Instead of a paragraph of text, I’d like dot points.”

- **Preference for structured symptoms**

An expert emphasised alignment with clinical structure: “I prefer a summary and a table of all Ps with symptoms described through tenacity, distress, interference and frequency.”

- **Clear definitions of experiences**

Two comments highlighted the need to define the target experience: “More specific description of the experience itself” and “still providing a global picture of several symptoms rather than only a single narrative.”

- **Clarify model inputs and outputs**

Request of clarification of the model’s input, “I’m unclear if the system was given a full transcript or only the case summary.”

- **Severity or risk indication**

Suggestion of adding a risk scale: “Could a scale of risk severity be included as a conclusion in the future.”

- **Data access and governance**

One comment raised governance concerns: “Is the model in-house only. Who has access to these data.”

J Additional Figures and Tables

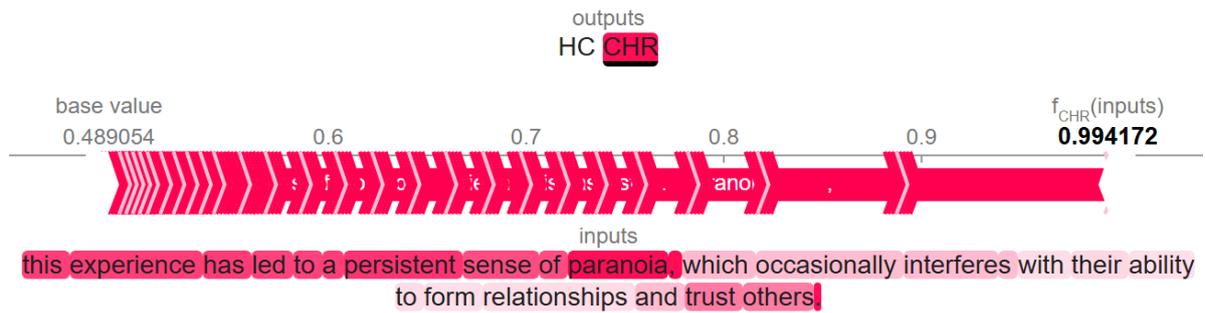


Figure 4: Inline token level heatmaps with SHAP

PSYCHS AI Insight Report

Welcome, clinicians!

Use this tool to transform your PSYCHS interviews into a clear, concise PDF summary—and have it delivered straight to your inbox.

How to proceed:

1. Upload your interview **transcript** or **audio recording**
2. Select the **PSYCHS symptom domains** you'd like highlighted
3. Provide the **clinician's email address** for delivery
4. Click **Generate & Send PDF**

Step 1: Choose input type

Transcript

Audio Recording

Upload transcript (.txt, .pdf)

Drag and drop file here
Limit 200MB per file • TXT, PDF

Browse files

Step 2: Select PSYCHS domains

Unusual Thought... x Suspiciousness /... x Grandiosity x Perceptual Abno... x

Conceptual Diso... x Impaired Tolera... x

Step 3: Clinician's email address

Generate & Send PDF

Figure 5: Website Interface.