Modeling Nonlinear Change in Psychotherapy: Toward an AI Decision-Support System with Synthetic Client Data

Oskar Šonc†
os05793@student.uni-lj.si
Pedagoška fakulteta, Kognitivna znanost
Ljubljana, Slovenia

Rok Smodiš
rs68734@student.uni-lj.si
Pedagoška fakulteta, Kognitivna znanost
Ljubljana, Slovenia

Tine Kolenik
tine.kolenik@ccsys.de
Institute of Synergetics and Psychotherapy Research, Paracelsus Medical University
Salzburg, Austria

Günter Schiepek
guenter.schiepek@ccsys.de
Institute of Synergetics and Psychotherapy Research, Paracelsus Medical University
Salzburg, Austria

Wolfgang Aichhorn
w.aichhorn@salk.at
University Hospital of Psychiatry, Psychotherapy, and Psychosomatics
Salzburg, Austria

Abstract

Psychotherapists typically choose interventions based on limited, session-bound information. We present a partial viability study of an AI decision support system for psychotherapy, which addresses this issue. The system forecasts next-day changes in five synergetic process variables: problem severity (P), therapeutic success (S), motivation (M), emotions (E), and insight (I), and combines these forecasts with phase transition detection to support anticipatory guidance. We created synthetic client personas and simulated daily trajectories for eighty to one hundred days with weekly sessions. Each day included a diary entry that aligns with the simulated state. We extracted features from diaries and session evaluations, including sentiment, readability, syntactic complexity, lexical richness, agreement, and discrepancy between client and therapist ratings. We evaluated Random Forest as the main model, along with Gradient Boosting and Ridge baselines, using splits by client. We also added a Pattern Transition Detection Algorithm (PTDA), which identifies critical fluctuations and potential transitions. Across dimensions, our preliminary results indicate that diary sentiment is the strongest predictor of next-day change. The pipeline demonstrates feasibility and provides a path to interpretable, real-time recommendations. Next steps include clinical validation on real data.

Keywords

decision support, psychotherapy, phase transitions, diary text, synthetic data

1 Introduction

Psychotherapeutic change is nonlinear, often marked by discontinuous shifts rather than steady improvement [1, 2]. Capturing these dynamics requires intensive monitoring, as daily diaries, high-frequency questionnaires like the Therapy Process Questionnaire, and brief session ratings yield time series suitable for detecting transitions [3, 4, 5]. Such data support both retrospective analyses and anticipatory detection, enabling real-time feedback for clinical decisions [4, 5]. Computational decision-support systems (DSS) have been proposed to operationalize this potential, integrating multimodal data to forecast therapeutic shifts and recommend personalized interventions [6]. Forecasts in this context mean short-horizon predictions of the five nonlinear synergetic state variables, which are problem severity (P), therapeutic success (S), motivation to change (M), emotions (E), and insight (I). They evolve daily and influence one another through nonlinear functions [1, 7]. By combining machine learning with interpretable, synergetic modeling, these systems aim to improve the timing and precision of interventions [4, 6]. The present study contributes a partial viability test of such a DSS, focusing on synthetic client data and evaluating a forecasting pipeline across these five dimensions. Our goal is not clinical validation, but methodological feasibility and guidance for future evaluation with real clinical data [1, 4].

This study is part of a broader project carried out by the Institute of Synergetics and Psychotherapy Research at the Paracelsus Medical University Salzburg. The project aims to develop an application that supports psychotherapists by suggesting and explaining personalized interventions across the five state variables (P, S, M, E, I).

Information Society 2025, Ljubljana, Slovenia

https://doi.org/10.70314/is.2025.cogni.9

2 Related work

2.1 Nonlinear change and intensive time-series in psychotherapy

The synergetic model represents change through five interacting state variables: problem severity (P), therapeutic success (S), motivation to change (M), emotions (E), and insight (I). Their nonlinear coupling produces instabilities and discontinuous transitions. Simulations show positive largest Lyapunov exponents, which imply restricted predictability, and daily self-report with the Therapy Process Questionnaire has been validated for intensive monitoring [7, 2, 4].

2.2 Phase-transition detection and forecasting

Transition-sensitive methods (e.g., PTDA-inspired indicators) detect impending shifts in trajectories; given chaotic dynamics, long-range prediction is infeasible and only short-horizon forecasts are appropriate for applied decision support [5].

2.3 Computational psychotherapy and decision support

Computational DSS integrate multimodal process data (e.g., questionnaires, diaries) with interpretable modeling and therapist-aligned explanations to generate actionable recommendations; our approach anchors these recommendations explicitly in the synergetic five-variable model [6].

2.4 Synthetic data for psychotherapy NLP pipelines

Because psychotherapy data are sensitive, synthetic corpora have been explored; zero-shot generations are often shallow, while few-shot/taxonomy-guided prompting and human-in-the-loop filtering improve fidelity [8, 9].

2.5 Empathy and therapeutic language modeling

Large language model generated therapy dialogues can train empathy detectors: augmenting a Reddit dataset with 420 synthetic dialogues improved F1 by up to 0.10 (exploration 0.48→0.53, interpretation 0.32→0.48, emotional reaction 0.58→0.59), and replacing 50% of the data raised interpretation accuracy from 0.50 to 0.57 while other metrics remained comparable; the study generated 10,464 synthetic dialogues and evaluated on 579 real dialogue pairs [10].

3 Research Objectives

This paper has two aims:

(1) **Partial viability study.** Demonstrate a working pipeline—data schema, feature extraction, forecasting, recommendation and explanation—that operates on synthetic clients mirroring our planned real-world data collection (five synergetic state variables + diary + pre/post session ratings).
(2) **Bridging detection and forecasting.** Outline how phase-transition detection (e.g., PTDA and related convergence-validated methods) can be integrated with short-horizon forecasting of the five dimensions to generate anticipatory intervention suggestions (e.g., focus more on S and I when a transition is imminent).

We position this as a pre-study on synthetic data, designed to de-risk methodological choices and inform the design of a pilot with genuine clinical time series.

4 Methodology

4.1 Synthetic Dataset Generation

We first generated a set of client personas with demographic and diagnostic diversity (e.g., gender, age, primary complaint). Personas were created using LLMs guided by structured prompts. For each persona, we initialized the five synergetic state variables—problem severity (P), therapeutic success (S), motivation to change (M), emotions (E), and insight/new perspectives (I)—from the persona profile, defaulting to near-neutral when unspecified. Daily diary entries complemented the numerical scores and were produced with a fixed prompt using GPT-4o-mini (temperature 0.7), conditioning on the current day's state and a brief progress note to ensure narrative coherence across days.

For each persona, we simulated daily trajectories of the five synergetic state variables defined in the nonlinear change model (Schiepek et al., 2016; Schiepek et al., 2017). The dynamics evolved with small fixed linear couplings and mild damping plus additive Gaussian noise, with values clipped to the range [-3, 6]. We simulated 80–100 days per client and designated every seventh day as a therapy session, which included structured pre- and post-session evaluations completed by both client and therapist. This procedure produced time series that exhibit variability, occasional instabilities, and realistic recovery trajectories; diary generation was conditioned on the simulated states to align text with day-level changes. Random seeds were fixed for reproducibility.

All outputs were stored in structured JSON files that contained raw trajectories, diary texts, session-day flags, and evaluation ratings. These were subsequently enriched with feature representations to support model training and interpretability.

4.2 Feature Extraction

Features were derived from both diary entries and session evaluations, capturing textual signals and structured ratings that inform downstream forecasting. Diary texts were processed using standard natural language processing pipelines, extracting sentiment scores (VADER, TextBlob), readability indices, syntactic complexity measures, lexical richness and word counts.

Session-day evaluations were transformed into quantitative descriptors by computing mean, variance, and maximum differences across therapist pre-, therapist post-, and client post-session ratings for each of the five synergetic state variables (P, S, M, E, I). Additionally, similarity metrics (cosine similarity, Euclidean distance) were calculated to assess alignment and discrepancy between client and therapist perspectives.

4.3 Model training

Forecasting models were developed to predict next-day scores for each dimension. We trained Random Forest regressors as the primary model due to their robustness and ability to provide interpretable feature importances. In addition, we fit Gradient Boosting regressors as a complementary tree-ensemble with a different bias–variance profile, tuning the number of estimators and learning rate on validation folds. Finally, we included a regularized linear comparator via Ridge Regression; features were standardized before fitting, providing a strong high-bias baseline and an interpretable contrast to the tree models. All models were trained and evaluated under the same grouped-by-client splits and metrics to enable direct comparison.

4.4 Phase-Transition Detection

To assess critical fluctuations in therapeutic trajectories, we implemented a phase-transition detection layer. Peaks in dynamic complexity were flagged as candidate transitions. A PTDA-inspired algorithm was then applied to combine these signals with additional markers, yielding annotated trajectories with transition indicators.

Together, these components operationalize an end-to-end pipeline that simulates client trajectories, extracts features, forecasts next-day changes, and detects phase transitions, producing interpretable outputs suitable for therapist review. The pipeline can be seen in Figure 1.

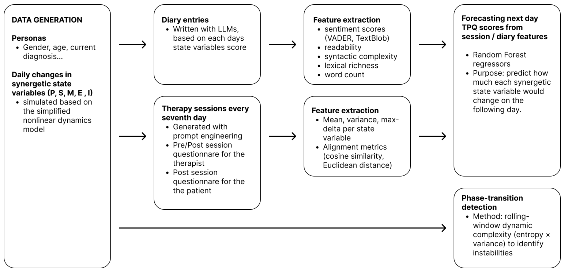

Figure 1: Project pipeline.

5 Results

Permutation importance analyses revealed that sentiment was consistently the strongest predictor across all dimensions. Figure 2 illustrates this pattern for the Emotions dimension, where sentiment clearly dominated over other text-based features such as readability, syntactic complexity, or lexical diversity. This finding indicates that the emotional valence expressed in daily diary entries was the most informative signal for forecasting short-term changes in therapy process variables.

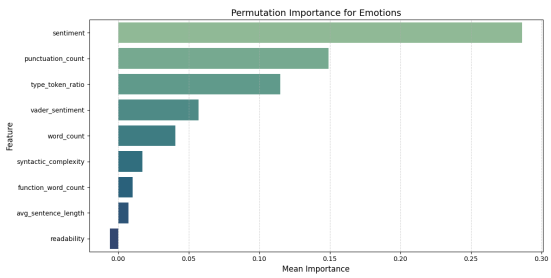

Figure 2: Predictive power of different diary entry characteristics (example: Emotions dimension).

On a cohort of synthetic clients, the pipeline trains stably and yields face-valid recommendations, exhibiting behavior that matches qualitative expectations across the five dimensions.Random Forest models produced smooth short-horizon predictions for Therapeutic success(S) and Emotion(E), with appropriately higher variance on Motivation and Insight when diaries include inconsistent motivational/insight signals.

We trained Random Forest, Gradient Boosting, and Ridge Regression models to forecast next-day values of the five synergetic state variables on synthetic client data generated from our nonlinear change model. Table 1 summarizes the performance in terms of mean squared error (MSE) and coefficient of determination (R^2), reported under the same evaluation protocol for direct comparability. Ridge obtained the lowest MSE (0.188) and the only positive R^2 (0.01), indicating a small but consistent gain over trivial predictors. Random Forest and Gradient Boosting achieved comparable errors (MSE $\approx$ 0.22) with negative R^2, indicating worse performance than a mean (null) baseline on this short-horizon task. For reference, the mean baseline (the R^2 anchor) has MSE $\approx$ 0.190, implying that Ridge reduces error by about 1% in aggregate.

Table 1: Therapy data predictions

Model	MSE	R^2
Random Forest	0.2208	-0.138
Gradient Boosting	0.2176	-0.151
Ridge Regression	0.1878	0.013

6 Discussion

6.1 What would likely change with real data.

Our synthetic-only runs mainly validated the pipeline, but they also favored a carry-forward baseline due to piecewise plateaus and low noise. On real synergetic state variables + diary + session series, we would expect less baseline advantage because of non-stationarity and therapist actions, modest but consistent MAE gains for short-horizon forecasts on P/S/E (with Motivation and Insight remaining more variable), and transition warnings with clinically useful lead times once thresholds are tuned to real fluctuations rather than generator quirks [7, 4, 5, 6].

6.2 Evidence from similar synthetic-vs-real comparisons.

Prior work has shown that synthetic corpora can help, but only when evaluated on real test sets. Cabrera Lozoya, Hernandez Lua, Barajas Perches, Conway, and D'Alfonso [10] generated 10,464 synthetic dialogues and found that augmenting Reddit data improved empathy F1 by up to 0.10, while replacing up to 50% of the organic data preserved or improved performance (e.g., interpretation accuracy 0.57 vs. 0.50) on a 579-pair clinical test set (MOST+ and Alexander Street). This pattern supports our claim that real synergetic state variables + diary data are necessary to calibrate feature weights (e.g., diary sentiment) and transition thresholds reliably.

6.3 Limitations and next steps

Our evaluation used only synthetic labels, a single-client showcase, and no ground-truth transitions, so we could not report precision/recall or lead time. Moreover, synthetic text may encode generator biases, inflating the apparent weight of some features [8, 9]. Next steps are therefore clear: collect real data, define transitions and compare to stronger baselines [5, 6].

7 Conclusion and Future Work

We presented a DSS that organizes session planning around five nonlinear change dimensions and provides explainable, forecast-driven intervention focus suggestions. This partial viability study shows the full pipeline operating on synthetic clients and specifies the next steps: (i) collect pilot synergetic state variables + diary + session micro-data; (ii) integrate a phase-transition layer; (iii) quantitatively evaluate forecasting and recommendation usefulness on real data; (iv) add a block-diagram figure and finalize a web prototype for therapist feedback. Ultimately, our goal is a hybrid system where nonlinear modelling and interpretable ML jointly inform what to focus on next in a given session.

8 Ethical note

This DSS is meant to be assistive only and does not automate clinical decisions or crisis response, emphasizing clinician oversight at all times. All results are based on synthetic data and make no claims of clinical efficacy. Recommendations are meant to be shown to clinicians for judgment only. Data are minimised and access-controlled; if ever in use, model inputs/outputs will be logged for audit.

References

[1] G. Schiepek, B. Aas, and K. Viol, "The mathematics of psychotherapy: A nonlinear model of change dynamics," *Nonlinear Dynamics, Psychology, and Life Sciences*, vol. 20, no. 3, pp. 369–399, 2016.

[2] H. Schöller, K. Viol, W. Aichhorn, M.-T. Hütt, and G. Schiepek, "Personality development in psychotherapy: A synergetic model of state-trait dynamics," *Cognitive Neurodynamics*, vol. 12, no. 5, pp. 441–459, Jun. 2018, ISSN: 1871-4099. DOI: 10.1007/s11571-018-9488-y [Online]. Available: http://dx.doi.org/10.1007/s11571-018-9488-y

[3] A. M. Hayes, J.-P. Laurenceau, G. Feldman, J. L. Strauss, and L. Cardaciotto, "Change is not always linear: The study of nonlinear and discontinuous patterns of change in psychotherapy," *Clinical Psychology Review*, vol. 27, no. 6, pp. 715–723, Jul. 2007, ISSN: 0272-7358. DOI: 10.1016/j.cpr.2007.01.008 [Online]. Available: http://dx.doi.org/10.1016/j.cpr.2007.01.008

[4] G. Schiepek et al., "The therapy process questionnaire - factor analysis and psychometric properties of a multidimensional self-rating scale for high-frequency monitoring of psychotherapeutic processes," *Clinical Psychology amp; Psychotherapy*, vol. 26, no. 5, pp. 586–602, Jul. 2019, ISSN: 1099-0879. DOI: 10.1002/cpp.2384 [Online]. Available: http://dx.doi.org/10.1002/cpp.2384

[5] G. Schiepek et al., "Convergent validation of methods for the identification of psychotherapeutic phase transitions in time series of empirical and model systems," *Frontiers in Psychology*, vol. 11, Aug. 2020, ISSN: 1664-1078. DOI: 10.3389/fpsyg.2020.01970 [Online]. Available: http://dx.doi.org/10.3389/fpsyg.2020.01970

[6] T. Kolenik, G. Schiepek, and M. Gams, "Computational psychotherapy system for mental health prediction and behavior change with a conversational agent," *Neuropsychiatric Disease and Treatment*, vol. Volume 20, pp. 2465–2498, Dec. 2024, ISSN: 1178-2021. DOI: 10.2147/ndt.s417695 [Online]. Available: http://dx.doi.org/10.2147/NDT.S417695

[7] G. K. Schiepek et al., "Psychotherapy is chaotic—(not only) in a computational world," *Frontiers in Psychology*, vol. 8, Apr. 2017, ISSN: 1664-1078. DOI: 10.3389/fpsyg.2017.00379 [Online]. Available: http://dx.doi.org/10.3389/fpsyg.2017.00379

[8] Z. Li, H. Zhu, Z. Lu, and M. Yin, "Synthetic data generation with large language models for text classification: Potential and limitations," in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, 2023. DOI: 10.18653/v1/2023.emnlp-main.647 [Online]. Available: http://dx.doi.org/10.18653/v1/2023.emnlp-main.647

[9] V. Veselovsky, M. H. Ribeiro, A. Arora, M. Josifoski, A. Anderson, and R. West, *Generating faithful synthetic data with large language models: A case study in computational social science*, 2023. DOI: 10.48550/ARXIV.2305.15041 [Online]. Available: https://arxiv.org/abs/2305.15041

[10] D. Cabrera Lozoya, E. Hernandez Lua, J. A. Barajas Perches, M. Conway, and S. D'Alfonso, "Synthetic empathy: Generating and evaluating artificial psychotherapy dialogues to detect empathy in counseling sessions," in *Proceedings of the 10th Workshop on Computational Linguistics and Clinical Psychology (CLPsych 2025)*, Association for Computational Linguistics, 2025, pp. 157–171. DOI: 10.18653/v1/2025.clpsych-1.13 [Online]. Available: http://dx.doi.org/10.18653/v1/2025.clpsych-1.13