

**Figure 2:** Analytical Framework of SupervisorBot: (A) R2D2 model flowchart. (B) Reinforcement learning framework of the psychotherapy recommendation system problem.

two agents in conversation, including the agreement on the goals to be achieved and the tasks to be carried out, and the bond, trust and respect to be established over the course of the dialogue [6]. In [7], we developed a natural language processing (NLP) approach to infer this quantity in real-time as ratings. Here we propose the Reinforced Recommendation model for Dialogue topics in psychiatric Disorders (R2D2), a the first ever recommendation system of dialogue topics proposed for the psychotherapy setting. It transcribes the session in real-time, predicts the therapeutic outcome as a turn-level rating, and recommends treatment strategy that is best for the current context and state of the psychotherapy. It is the first step to solving the global issue of mental health by augmenting the treatment and education of clinical practitioners with a recommendation system of therapeutic strategy.

## 2. Methods

Fig 1 is an outline of the analytic framework. The continuous audio stream is fed into the system. First, we perform the speaker diarization (e.g. using real-time solutions such as [8, 9, 10]) which separates audio into dyads of doctor-patient, which are then transcribed into natural language turns for real-time downstream analyses.

**Therapeutic quality ratings.** After we obtain a relatively well diarization result, we can configure the quality assessment setting by specifying a proper inventory. In this system, we use the Working Alliance Inventory

(WAI), a set of self-report measurement questionnaire that quantifies the therapeutic bond, task agreement, and goal agreement [11, 12, 13]. Operationally, our goal is to derive from these 36 items three alliance scales: the task scale, the bond scale and the goal scale. They measure the three major themes of psychotherapy outcomes: (1) the collaborative nature of the dialogue participants' relationship; (2) the affective bond between them, and (3) their capabilities to agree on treatment-related short-term tasks and long-term goals. The score corresponding to the three scales comes from a key table which specifies the positivity or the sign weight to be applied on the questionnaire answer when summing in the end.

**Transcription and real-time rating assessment.** Now we are ready for real-time quality annotation. Given the audio stream for a given user, we first transcribe the diarized audio stream with standard automatic speech recognition module [14]. Following the approach proposed in [7, 15, 16, 17], we embed both the dialogue turns and WAI items with deep sentence or paragraph embeddings (in this case, Doc2Vec [18]), and then compute the cosine similarity between the embedding vectors of the turn and its corresponding inventory vectors. With that, for each turn (either by patient or by therapist), we obtain a 36-dimension working alliance score, which we may save in a bidirectional relational database as in [19].

**Topic modeling as recommendation items.** First, we define the "items", "users", "contents" and "ratings" in our recommendation system. Here, the "items" the system recommends are treatment strategies. In this ex-

ample, we represent these strategies as a topic that the therapist should initiate or continue for the next turn. Given a large text corpus of many psychotherapy sessions, as in [20] we can first perform topic modeling to extract the main concepts discussed in the psychotherapy. We use the Embedded Topic Model (ETM) [21] in this work because it was shown to create the most diverse concepts in psychological corpus [20]. In this study, we use annotate each turn with their most likely topic and identifies seven unique topics (Topic 0 is about figuring out, self-discovery and reminiscence; Topic 1 is about play. Topic 2 is about anger, scare and sadness. Topic 3 is about counts. Topic 6 is about explicit ways to deal with stress, such as keep busy and reaching out for help. Topic 7 is about numbers. Topic 8 is about continuation.)

**Recommendation system setting.** Then, we pair these “items” with the “users” and “contents”, which in our case, would be the patientID, his or her previous turns, their aggregated formats and other meta data. For instance, we know that within each sessions, there exists many pairs of turns, and they would belong to the same “user”. However, one can also assign all turns belong to one clinical label, or all turns related to a certain topic as one “user”. In this example, we choose the session ids as users. And lastly, the “ratings” would be patient’s inferred alliance scores predictive of the therapeutic outcomes. Creating this database from historical data, we can train our system. Since we have defined our users, items, contents and ratings, the recommendation engine can be easily crafted with content-based [22, 23, 24] and collaborative filtering [25, 26, 27, 28]. Since our session turns are sequential and can specify a state or timestamp, it might be suitable for RL [29, 30, 31] and session-based methods [32, 33, 34], which can be neuroscience or psychiatry-inspired [35, 36, 37, 38] to provide further interpretable clinical insights. During the deployment, our system registers our session as a new “user” if we adopt a session-based item, providing punctuated rater evaluations as inference anchors [39]. Next steps include predicting these inference anchors as states (like [20? ]) and training chatbots as reinforcement learning agents given these states.

**Deep reinforcement learning recommendation approaches.** Reinforcement learning approaches are effectively applied in language and speech tasks (as reviewed in [40]), among which recommendation is an important use case. As shown in Figure 2, the reinforcement learning environment is formulated such that the recommendation agent takes an action by recommending a strategy (say, a discussion topic). And the therapist will interact with the patient taking that suggestion into account. The dialogue interaction, in turn, has a quality evaluation of some sort (say, the therapeutic working alliance score). This serves as a reward to the recommendation agent to update its weights. In the meanwhile, the

state is progressed to the next therapeutic states. As a first step, we evaluate three popular deep RL algorithms. Based on the deterministic policy gradient in an actor-critic architecture, the Deep Deterministic Policy Gradients (DDPG) [41] is a model-free algorithm for continuous action spaces, and one of the first successful algorithms to learn policies end-to-end. Building upon the Double Q-Learning [42], Twin Delayed DDPG (TD3) [43] is a similar solution is proposed to correct for the overestimated value issue, and yields more competitive results in various game settings. As the online data collection of RL models are usually time consuming, in real world industrial setting, these models are usually trained using previously collected data. As a result, there is a growing popularity of offline reinforcement learning methods [44]. Among them, Batch Constrained Q-Learning (BCQ) [45] is the first continuous control deep RL algorithm with competitive results in off policy evaluations by restricting the agent’s exploration in action space.

**Reinforced Recommendation model for Dialogue topics in psychiatric Disorders (R2D2).** Combining all the elements, we have our R2D2 model (Figure 2A). Here we identify each session as a user, and the states are frames of dialogues which can be labelled their topics in real time and their ratings with a working alliance (WA) inference module. The reinforcement learning core, powered by deep RL, predicts the best action represented by an embedding for the items (topics). This embedding can be pre-computed, for instance, using dimension reduction techniques to find clusters of different topics in a low-dimensional space. We use the Doc2Vec embedding of the original dialogue turns, averaged by their topic labels, such that each action (i.e. the topic id) have an averaged representation in the sentence embedding space. This action representation can be translated into a topic label with nearest neighbor, and a given dialogue response will be selected from the historical dialogue data to continue the conversation. The reward can then be computed using the working alliance rate.

### 3. Empirical results

**Experimental setting.** To evaluate the recommendation systems, we preprocess the Alex Street psychotherapy dataset<sup>1</sup>, which consists of transcribed recordings of over 950 therapy sessions between multiple anonymized therapists and patients, into a recommendation system format (219,999 recommendation actions) and then split it into 95/5 train-test sets. The dataset consists of four types of psychiatric conditions: anxiety, depression, schizophrenia and suicidal cases. We train R2D2 on each of the text corpus, as well as on all four together. To set up the batch

<sup>1</sup><https://alexanderstreet.com/products/counseling-and-psychotherapy-transcripts-series>

training for reinforcement learning, we cut the turns into frames of 10 turn pairs and use a batch size of 32. We represent the action spaces (the topics to recommend) in three candidate embedding spaces: the averaged 300-dimension Doc2Vec embedding for each topic, the averaged 36-dimension principal component analysis (PCA) embedding, and the averaged 2-dimension Uniform Manifold Approximation and Projection (UMAP) embedding. Due to the space limit, we only present the results for the first embedding, but leave the others in appendix. We train the R2D2 with three reinforcement learning agents (DDPG, TD3 and BCQ) each for 50 epochs, where their losses consistently drop and converge in a stable way. Based on the loss curve, there are no overfitting in all model training processes.

**Empirical results.** To evaluate the performance of the three recommendation agents, we compute the Pearson’s  $r$  of the recommended actions with their corresponding ground truth actions in the test set (Table 1). Since we are the first system in this application problem, there are no state-of-the-art or baseline so far. Instead, we compare among variants of R2D2. Other than testing on different subset of the datasets and reinforcement learning algorithms, we also use three different scales of working alliance as our ratings: task, bond and goal, which measures different aspects of emotional alignments in psychotherapy. We observe that the best performing model for four disorders are: R2D2-DDPG-TASK for depression sessions with a correlation of 0.3796, R2D2-BCQ-TASK for depression session (0.4042), R2D2-TD3-GOAL for schizophrenia sessions (0.4599) and R2D2-BCQ-BOND for suicidal sessions (0.4152). If we consider all four classes together, R2D2-TD3-GOAL appears to be the best performing models (0.3765). We notice that the DDPG and TD3 bases of R2D2 yields similar rankings among using three working alliance scales as their ratings, while the BCQ tends not to. For instance, in schizophrenia cases, the alignment in the goal scale appear to provide a far more advantageous recommendation prediction than the other two implicit feedbacks (task and bond alignments), while in R2D2-BCQ, the effect is less pronounced. For specific disorders, R2D2-DDPG is the recommender winner for anxiety, depression and schizophrenia, and R2D2-TD3 is the winner for suicidal cases (which should be taken with a grain of salt considering the small amount of data we have on them). When pooling the sessions of four disorders together, the recommender winner appears to be R2D2-TD3, which may suggest that R2D2-TD3, given its twin delayed mechanism to correct for value overestimation, are better suited for heterogeneous rather than homogeneous cases. It was a surprise that R2D2-BCQ doesn’t demonstrate in our dataset, an advantage to constrain the possible extrapolation errors by the non-offline methods. This evaluation provides a proof of concept. Future work will focus on systematically comparing a larger spectrum

**Table 1**

Pearson’s  $r$  of the actual actions taken in the test set with their predicted actions

	Anxi	Depr	Schi	Suic	All
R2D2-DDPG-TASK	<b>0.3796</b>	0.3376	0.1556	0.3292	0.0578
R2D2-DDPG-BOND	0.2417	0.3838	0.1539	0.0873	0.1455
R2D2-DDPG-GOAL	0.0761	0.3682	0.4589	-0.0210	0.2243
R2D2-TD3-TASK	0.0707	0.1310	0.0443	0.3188	0.3357
R2D2-TD3-BOND	0.2018	0.3363	0.0908	0.3070	0.1101
R2D2-TD3-GOAL	0.0984	0.2222	<b>0.4599</b>	0.2044	<b>0.3765</b>
R2D2-BCQ-TASK	0.1128	<b>0.4042</b>	0.1401	0.1422	0.0825
R2D2-BCQ-BOND	0.0778	0.0876	0.0987	<b>0.4152</b>	0.0885
R2D2-BCQ-GOAL	0.0810	0.1231	0.0833	0.0788	0.0780

of deep reinforcement learning and model architectures.

**Ethical considerations.** Following the ethical guidelines in [46, 47] and the operational suggestions in [48], we make sure that all training examples that we evaluate on are properly anonymized with pre- and post-processing techniques, and disclaim that these investigations are proof of concept and require extensive caution to prevent from the pitfall of over-interpretation.

## 4. Web-Based System: SupervisorBot

“SupervisorBot” is an interactive web-based system (Fig 3 [39]). We first give users the instructions on how to use the system. Then they are lead to input their own inventory used to evaluate the dialogue quality. In this case, we put in a default one, using the working alliance inventory. They are guided to input the score scale corresponding to each inventory item and click on “Analyze” to finalize. In the speaker diarization part, we compute and visualize the Mel Frequency Cepstral Coefficients (MFCC) in a sliding window fashion given the real-time audio input from microphone, with the MFCC bands color coded in the page. Finishing these two steps as the preparation, the system is now running, and the therapist can sit back and go on with the session. The app now moves to the annotation panel, where the therapist can see that a transcript is displayed, along with who is speaking. The computed alliance score in the three scales are also dynamically displayed in real-time according to the content of the dialogue turn. This is helpful information to assist the therapist. And in our last panel, we have our recommendation guidance. The topics to choose from are ranked and top N are displayed. The therapist can use it as a hint and initiate his response given a top recommendation. The system will transcribe his response and highlight the topic he most likely ended up choosing in the last round, and save that information as part of historical data. The system refreshes its parameters at the end of each session to fit new data.




Figure 3: State screenshots of SupervisorBot web app: inventory inputs, diarization, annotation, strategy recommendation.

## 5. Conclusions and Future Directions

In this work, we provide a practical example of how a real-time recommendation system can help therapists better treat their patients in psychotherapy sessions with informative clinical annotations and recommendations of treatment strategies with deep reinforcement learning. Although in this example, the strategies are the topics for the therapist to initiate or continue, the same approach can be extended to more complex and nuanced treatment suggestions. For instance, in the ABC approach of cognitive behavioral therapy (CBT), our system can suggest a belief (B) to guide the patients to better understand the causality between the activating event (A) and its consequence (C).

Before we conclude, another interesting perspective to view this line of research is hidden in Figure 1: while the recommendation agent is driven by reinforcement learning, the therapist (and even patient) have their agency which updates under the reinforcement learning principles. For instance, the patient can directly offer feedback to the therapists. And given the feedback, the therapist may adjust his or her internal model to weigh on the quality of the suggestions by the recommendation agent. Next steps include modeling these theory of minds and confidence levels in this multi-participant human computer interaction setting. providing punctuated rater evaluations as inference anchors. Next steps include predicting these inference anchors as states (like [20, 49]) and training chatbots as reinforcement learning agents given these states (like [36, 37, 38]).

## References

- [1] V. Patel, S. Saxena, C. Lund, G. Thornicroft, F. Bain-gana, P. Bolton, D. Chisholm, P. Y. Collins, J. L. Cooper, J. Eaton, et al., The lancet commission on global mental health and sustainable development, *The lancet* 392 (2018) 1553–1598.
- [2] A. Satiani, J. Niedermier, B. Satiani, D. P. Svendsen, Projected workforce of psychiatrists in the united states: a population analysis, *Psychiatric Services* 69 (2018) 710–713.
- [3] X. Wang, S. Hegde, C. Son, B. Keller, A. Smith, F. Sasangohar, et al., Investigating mental health of us college students during the covid-19 pandemic: cross-sectional survey study, *Journal of medical Internet research* 22 (2020) e22817.
- [4] C. E. Watkins Jr, Being and becoming a psychother-apy supervisor: The crucial triad of learning dif-ficulties, *American Journal of Psychotherapy* 67 (2013) 134–150.
- [5] B. E. Wampold, How important are the common factors in psychotherapy? an update, *World Psych* 14 (2015) 270–277.
- [6] E. S. Bordin, The generalizability of the psychoana-lytic concept of the working alliance., *Psychother-apy: Theory, research & practice* (1979) 252.
- [7] B. Lin, G. Cecchi, D. Bouneffouf, Deep annotation of therapeutic working alliance in psychotherapy, arXiv preprint arXiv:2204.05522 (2022).
- [8] B. Lin, X. Zhang, Speaker diarization as a fully on-line bandit learning problem in minivox, in: *ACML*, 2021, pp. 1660–1674.
- [9] B. Lin, X. Zhang, Voiceid on the fly: A speaker recognition system that learns from scratch, in:



- INTERSPEECH, 2020.
- [10] B. Lin, X. Zhang, Speaker diarization as a fully online learning problem in minivox, arXiv preprint arXiv:2006.04376 (2020).
- [11] A. O. Horvath, An exploratory study of the working alliance: Its measurement and relationship to therapy outcome (1981).
- [12] T. J. Tracey, A. M. Kokotovic, Factor structure of the working alliance inventory., *Psychological Assessment: A journal of consulting and clinical psychology* 1 (1989) 207.
- [13] D. J. Martin, J. P. Garske, M. K. Davis, Relation of the therapeutic alliance with outcome and other variables: a meta-analytic review., *Journal of consulting and clinical psychology* 68 (2000) 438.
- [14] J. Adorf, Web speech api, KTH Royal Institute of Technology (2013) 1–11.
- [15] B. Lin, G. Cecchi, D. Bouneffouf, Working alliance transformer for psychotherapy dialogue classification, arXiv preprint arXiv:2210.15603 (2022).
- [16] B. Lin, Personality effect on psychotherapy outcome: A predictive natural language processing framework, arXiv preprint (2022).
- [17] B. Lin, Voice2Alliance: automatic speaker diarization and quality assurance of conversational alignment, in: INTERSPEECH, 2022.
- [18] Q. Le, T. Mikolov, Distributed representations of sentences and documents, in: *International conference on machine learning*, PMLR, 2014, pp. 1188–1196.
- [19] B. Lin, Knowledge management system with nlp-assisted annotations: A brief survey and outlook, in: *CIKM Workshops*, 2022.
- [20] B. Lin, D. Bouneffouf, G. Cecchi, R. Tejwani, Neural topic modeling of psychotherapy sessions, arXiv preprint arXiv:2204.10189 (2022).
- [21] R. Wang, X. Hu, D. Zhou, Y. He, Y. Xiong, C. Ye, H. Xu, Neural topic modeling with bidirectional adversarial training, in: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 340–350.
- [22] M. J. Pazzani, D. Billsus, Content-based recommendation systems, in: *The adaptive web*, Springer, 2007, pp. 325–341.
- [23] C. Basu, H. Hirsh, W. Cohen, et al., Recommendation as classification: Using social and content-based information in recommendation, in: *Aaai/iaai*, 1998, pp. 714–720.
- [24] C. C. Aggarwal, et al., *Recommender systems*, volume 1, Springer, 2016.
- [25] B. Sarwar, G. Karypis, J. Konstan, J. Riedl, Item-based collaborative filtering recommendation algorithms, in: *Proceedings of international conference on World Wide Web*, 2001, pp. 285–295.
- [26] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, T.-S. Chua, Neural collaborative filtering, in: *Proceedings of the 26th international conference on world wide web*, 2017, pp. 173–182.
- [27] Y. Koren, S. Rendle, R. Bell, Advances in collaborative filtering, *Recommender systems handbook* (2022) 91–142.
- [28] X. Su, T. M. Khoshgoftaar, A survey of collaborative filtering techniques, *Advances in artificial intelligence* 2009 (2009).
- [29] G. Zheng, F. Zhang, Z. Zheng, Y. Xiang, N. J. Yuan, X. Xie, Z. Li, Drn: A deep reinforcement learning framework for news recommendation, in: *Proceedings of the 2018 world wide web conference*, 2018, pp. 167–176.
- [30] X. Wang, Y. Wang, D. Hsu, Y. Wang, Exploration in interactive personalized music recommendation: a reinforcement learning approach, *ACM TOMM* 11 (2014) 1–22.
- [31] L. Zou, L. Xia, P. Du, Z. Zhang, T. Bai, W. Liu, J.-Y. Nie, D. Yin, Pseudo dyna-q: A reinforcement learning framework for interactive recommendation, in: *Proceedings of the 13th International Conference on Web Search and Data Mining*, 2020, pp. 816–824.
- [32] J. Li, P. Ren, Z. Chen, Z. Ren, T. Lian, J. Ma, Neural attentive session-based recommendation, in: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, 2017, pp. 1419–1428.
- [33] S. Wu, Y. Tang, Y. Zhu, L. Wang, X. Xie, T. Tan, Session-based recommendation with graph neural networks, in: *Proceedings of AAAI*, volume 33, 2019, pp. 346–353.
- [34] M. Ludewig, D. Jannach, Evaluation of session-based recommendation algorithms, *User Modeling and User-Adapted Interaction* 28 (2018) 331–390.
- [35] B. Lin, D. Bouneffouf, G. Cecchi, Split Q Learning: Reinforcement Learning with Two-Stream Rewards, in: *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19, International Joint Conferences on Artificial Intelligence Organization*, 2019, pp. 6448–6449. URL: <https://doi.org/10.24963/ijcai.2019/913>. doi:10.24963/ijcai.2019/913.
- [36] B. Lin, G. Cecchi, D. Bouneffouf, J. Reinen, I. Rish, A story of two streams: Reinforcement learning models from human behavior and neuropsychiatry, in: *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, 2020, pp. 744–752.
- [37] B. Lin, G. Cecchi, D. Bouneffouf, J. Reinen, I. Rish, Models of human behavioral agents in bandits, contextual bandits and rl, in: *International Workshop on Human Brain and Artificial Intelligence*, Springer, 2021, pp. 14–33.
- [38] B. Lin, G. Cecchi, D. Bouneffouf, J. Reinen, I. Rish,

- Unified models of human behavioral agents in bandits, contextual bandits and rl, *arXiv preprint arXiv:2005.04544* (2020).
- [39] B. Lin, G. Cecchi, D. Bouneffouf, Supervisorbot: Nlp-annotated real-time recommendations of psychotherapy treatment strategies with deep reinforcement learning, *arXiv preprint arXiv:2208.13077* (2022).
- [40] B. Lin, Reinforcement learning and bandits for speech and language processing: Tutorial, review and outlook, *arXiv preprint arXiv:2210.13623* (2022).
- [41] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, D. Wierstra, Continuous control with deep reinforcement learning, *arXiv preprint arXiv:1509.02971* (2015).
- [42] H. Hasselt, Double q-learning, *Advances in neural information processing systems* 23 (2010).
- [43] S. Fujimoto, H. Hoof, D. Meger, Addressing function approximation error in actor-critic methods, in: *International conference on machine learning*, PMLR, 2018, pp. 1587–1596.
- [44] S. Levine, A. Kumar, G. Tucker, J. Fu, Offline reinforcement learning: Tutorial, review, and perspectives on open problems, *arXiv preprint arXiv:2005.01643* (2020).
- [45] S. Fujimoto, D. Meger, D. Precup, Off-policy deep reinforcement learning without exploration, in: *International conference on machine learning*, PMLR, 2019, pp. 2052–2062.
- [46] T. Matthews, K. O’Leary, A. Turner, M. Sleeper, J. P. Woelfer, M. Shelton, C. Manthorne, E. F. Churchill, S. Consolvo, Stories from survivors: Privacy & security practices when coping with intimate partner abuse, in: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2017, pp. 2189–2201.
- [47] S. Graham, C. Depp, E. E. Lee, C. Nebeker, X. Tu, H.-C. Kim, D. V. Jeste, Artificial intelligence for mental health and mental illnesses: an overview, *Current psychiatry reports* 21 (2019) 1–18.
- [48] B. Lin, Computational inference in cognitive science: Operational, societal and ethical considerations, *arXiv preprint arXiv:2210.13526* (2022).
- [49] B. Lin, D. Bouneffouf, G. Cecchi, Predicting human decision making in psychological tasks with recurrent neural networks, *PloS one* (2022).