

Convergence of Chatbot Personalities Using Reinforcement Learning And Text Generation

Parijat Kar

University of California
parijatkar@uci.edu

Abstract

Chatbots have become indispensable tools in customer support and mental health industries, facilitating efficient communication and problem-solving. However, the success of these interactions often hinges on the alignment between the chatbot's personality and the user's preferences.

This study investigates the significance of personality matching in chatbot conversations, exploring the impact of MBTI personality types on user engagement and satisfaction. We delve into strategies for adopting appropriate tone and persona in one-to-one and group chat-bot interactions. By understanding the nuances of personality and tailoring chatbot responses accordingly, we aim to optimize the effectiveness of these virtual agents and enhance the overall user experience. AI creates proceedings, working notes, and technical reports directly from electronic sources furnished by the authors. To ensure that all papers in the publication have a uniform appearance, authors must adhere to the following instructions.

- Judging vs. Perceiving (J/P): Judgers prefer structure and planning, while perceivers value flexibility and spontaneity. The Impact of MBTI on Social and Business Interactions: Understanding MBTI personality types can significantly enhance social and business interactions. Following colleagues' different preferences and communication styles might improve teamwork. Effective leaders can tailor their leadership styles to accommodate the diverse personalities of their team members. Additionally, understanding personality differences can aid in conflict resolution and relationship building. Finally, MBTI can help individuals identify careers that align with their natural preferences and strengths. (Berg A I, & Johansson B (2014), Birditt K, & Antonucci T C (2008), Adelstein, Jonathan S. et al 2011)

Background Work

The Myers-Briggs Type Indicator (MBTI) is a popular personality assessment tool that categorizes individuals into 16 distinct personality types based on their preferences for introversion or extraversion, sensing or intuition, thinking or feeling, and judging or perceiving. Understanding these personality types can provide valuable insights into how people interact with others and navigate social and business situations. (Ackerman, P.L., & Beier, M.E. (2003), Barrick, M.R., Mount, M.K., & Gupta, R. (2003) Anderson C, Keltner D, & John O P (2003)).

The Four Dimensions of MBTI:

- Introversion vs. Extraversion (I/E): Introverts tend to recharge their energy alone or in small groups, while extroverts derive energy from social interactions.
- Sensing vs. Intuition (S/N): Sensors prefer concrete facts and details, whereas intuitive are drawn to abstract concepts and possibilities.
- Thinking vs. Feeling (T/F): Thinkers make decisions based on logic and analysis, while feelers consider personal values and emotions.

Large Language Models

Context Identification

Large language models (LLMs) excel at identifying context within a conversation. They can analyze the surrounding text, previous exchanges, and external information to understand the topic, setting, and participants involved. This contextual understanding enables them to generate more relevant and coherent responses.

Tone Identification

The tone refers to the overall emotional attitude conveyed by the chatbot's responses. It can range from empathetic and supportive to informative and neutral.

LLMs can effectively identify a conversation tone. Analyzing factors such as word choice, sentence structure, and punctuation can determine whether the tone is positive, negative, neutral, sarcastic, or entirely different. This understanding allows them to respond appropriately and consistently to the overall sentiment of the conversation.

Personality Identification

Persona refers to the chatbot's perceived identity or character. A well-defined person can help the chatbot feel more human and relatable to users. A compassionate and understanding person can foster trust and rapport with users in mental health settings.

While LLMs can't directly identify a person's personality, they can infer certain personality traits based on the text they process. For example, an LLM might conclude that someone is extroverted if they use frequent exclamations and engage in small talk. However, it's important to note that these inferences are based on patterns and may not always be accurate, as individual personalities are complex and multifaceted. (Keyu Pan et al 2023, Ontoum, S et al 2022)

Sentiment Identification

Sentiment refers to the emotional polarity of the chatbot's responses. It can be positive, negative, or neutral. In mental health settings, a chatbot should be able to identify and respond to negative sentiments in the user's messages. In contact centers, a positive sentiment can help to defuse tensions and resolve issues efficiently.

Tone, sentiment, and persona play crucial roles in the effectiveness of chatbots, particularly in mental health and contact center settings.

Here are two examples of chatbot applications that could use an appropriate personality.

Mental Health Chatbot

A mental health chatbot should be designed to provide compassionate and supportive responses to users struggling with emotional well-being. The chatbot's tone should be empathetic and understanding, acknowledging the user's feelings and offering encouragement. The chatbot should also be able to identify and respond to negative sentiments in the user's messages, providing appropriate support and resources. A well-defined person can help the chatbot feel more human and relatable to users, fostering trust and rapport. (Siyuan Chen et al 2023)

Contact Center Chatbot

A contact center chatbot should be designed to efficiently and helpfully respond to customer inquiries. The chatbot's tone should be professional while being polite and respectful. The chatbot should be able to identify and address customer concerns, resolving issues quickly and effectively. A helpful and knowledgeable persona can enhance the user experience, making the chatbot valuable to the contact center. (Xu Y et al 2022)

Personality-Related Work on Large Language Models

Personality and Interests

Research has explored the intricate relationship between personality traits and vocational interests. It identifies four distinct trait complexes that influence this connection: Science/Math, Intellectual/cultural, Social, and Conventional. These complexes represent different patterns of personality traits and associated career preferences. (Stoll G 2020)

Research also indicates that specific personality traits within these complexes are correlated with vocational interests. For example, individuals high in harm avoidance may be less interested in Realistic or Science-related careers. Those high in Achievement are likely interested in Investigative, Math, and Science fields. (Market S. 2016)

Within-Couple Personality Dynamics

Beyond individual career choices, past works showed personality dynamics within romantic relationships. It highlights the significance of within-couple personality similarity for relationship well-being. Couples sharing similar personality traits tend to have more satisfying and enduring relationships. (Lewis et al 2020)

Emotional Convergence

An intriguing concept explored in the text is emotional convergence. This refers to the phenomenon where partners' emotions become increasingly similar. Emotional similarity fosters mutual understanding, coordinated actions, and stronger interpersonal bonds.

In essence, knowledge provides valuable insights into the interplay between personality, interests, and relationships, offering a deeper understanding of how these factors influence our lives. (Mehta et al. (2020))

Hypothesis

There are several works on MBTI personality determination since the advancement of Large Language Models (LLM). (Champa, H. N et al 2010, Kalghatgi MP et al 2015, Li L et al 2014). A study by Peking University found that certain LLMs exhibited personality traits like the INTJ type, known for their introversion, intuitiveness, feeling, and judging preferences. Another study explored the possibility of using MBTI as an evaluation metric for LLMs. Researchers found that while MBTI isn't a perfect fit for assessing AI personalities, it can provide valuable insights into their tendencies and behaviors. (Cui J et al 2023)

There have been attempts to create LLMs with specific MBTI personalities. For example, researchers have experimented with training models to exhibit the traits of an INTJ, known for their logical, strategic thinking. (Wang Y 2025)

In this work, we are interested in assessing how different interaction strategies can influence personality adaptability for autonomous chatbots. In a well-distributed chatbot population, when chatbots adapt to individual strategies they can converge to one MBTI personality type.

Experiment Setup

We used MBTI data from the popular public data repository Kaggle. It has several MBTI datasets. We took two of those datasets. (Kaggle MBTI Dataset).

(MBTI) Myers-Briggs Personality Type Dataset

This dataset has about 10,000 examples of interactions labeled with MBTI Personality types with frequencies in Table 1.

From this dataset, we used 6 personalities INFP, INFJ, INTP, INTJ, ENTP, and ENFP. We did not include the other personalities in a relatively small number of instances.

MBTI Personality Types 500 Dataset

This dataset has about 100,000 examples of interactions labeled with MBTI Personality types with frequencies in Table 2.

We used 6 personalities INFP, INFJ, INTP, INTJ, ENTP, and ENFP. We did not include the other personalities in a relatively small number of instances. We removed all non-dictionary words. We split significant texts into 3-4 sentences.

We chose to work with Amazon Webservices Titan G1 Express, Claude V2, Llama V3, GPT4, and Gemini. We used Amazon Bedrock for prompt management, fine-tuning models, and flow.

Using data described in Table 1 and Table 2, we fine-tuned the large language models and created custom models for this experiment. We considered each instance in the dataset as a large language model response. We used a prompt to generate possible questions for each response. Use those questions and the original responses as prompts and completes.

Here is an example prompt in Table 3 of how such questions were generated.

MBTI Type	Instances
INFP	1832
INFJ	1470
INTP	1304
INTJ	1091
ENTP	685
ENFP	675
ISTP	337
ISFP	271
ENTJ	231
ISTJ	205
ENFJ	190
ISFJ	166
ESTP	89
ESFP	48
ESFJ	42
ESTJ	39

Table 1: Data Distribution 1

MBTI Type	Instances
INFP	24961
INFJ	22427
INTP	14963
INTJ	12134
ENTP	11725
ENFP	6167
ISTP	3424
ISFP	2955
ENTJ	1986
ISTJ	1534
ENFJ	1243
ISFJ	875
ESTP	650
ESFP	482
ESFJ	360
ESTJ	181

Table 2: Data Distribution 2

<p>Prompt: Generate three questions with serial numbers for this response: Move to the Denver area and start a new life for myself. (sic)</p> <p>Response (RQ)</p> <ol style="list-style-type: none"> 1. What drew you to the Denver area specifically? 2. Have you started planning your move to Denver yet? 3. How do you feel about this big change in your life? <p>Fine-tuning foundational models: Preprocess data to remove stop words Divide data into 6 training sets for each personality considered for this work Use a prompt from Table 3 to generate three questions for each training data. For each of the questions generated, add training data in the training file with a prompt from the response question. RQ[i] and a conclusion with the original training data</p>

Table 6: Example Prompt

Guardrails (word filtering for responsible AI) were used to prevent all abusive language. Training instances for our guardrails are taken from Kaggle’s Hate Speech and Offensive Language Dataset.

Tracking Conversation History

We used two setups. First, chatbots maintain a complete conversation history with all interactions with themselves. The other setup used conversations in a public mode where each chatbot could save every conversation in the session.

Personality Classification

We created a central LLM model to identify a personality from a response. Fine-tune is based on the original training data (only). We used a prompt flow to determine an incoming response type from this central LLM. One such classification example is in Table 4.

Generate Follow-up Questions

We used its fine-tuned model and a prompt pattern to generate a follow-up question. See a follow-up question prompt in Table 5.

Initiating a Conversation

We selected a few topics to start a conversation. We used several documents as knowledge from each to create a response using a retrieval augmented generation technique.

<p><i>Classify the following statement from 6 MBTI personalities: INFP, INFJ, INTP, INTJ, ENTP, and ENFP. One word, only the type. That's another silly misconception. That approach is logically going to be the key to unlocking whatever it is you think you are entitled to. Nobody wants to be approached with BS...(sic)</i></p>
--

Table 5: Example Prompt for Classification

<p><i>Create a follow-up non-specific question for the following response:</i></p> <p><i>"Those responses make sense. Forgot about the difficulty factor. I have heard before from others that Ancient Greek is much harder. And although there are more Greek writers I'm currently interested..."</i></p> <p><i>Response: "So, what's stopping you from diving into Ancient Greek?"</i></p>

Table 4: Example Follow-up Question

<p><i>I was in Central Europe, visiting Budapest, Vienna, Prague, and Cesky Krumlov. Enjoy stunning architecture, rich history, and diverse cultural experiences while sampling delicious local cuisine.</i></p> <p><i>That sounds like an incredible trip!</i> <i>Central Europe offers a wealth of cultural and historical experiences. Did you have any particular highlights or favorite moments from your travels?</i></p>
--

Table 3: Example Conversation

These topics include sports, travel, cooking, art, and literature. An example is shown in Table 6.

Response Generation

We are taking the data analysis from the Myers-Briggs Type Indicator neural network model. However, we are using COSTAR prompt generation techniques. The COSTAR framework outlines essential elements for effective LLM prompts: a context for understanding, an objective for focus, a style for alignment, a tone for sentiment, an audience for targeting, and a response format for output. However, we are adding a novel learning factor to weigh in past responses. We have compared strategies where we matched tone, stayed at my style, and matched tone until a tone threshold was reached.

Based on the conversation history, Use a matching tone for the response. The tone types we used in this experiment are given in Table 7.

<p>Positive: Expresses happiness, joy, approval, or optimism.</p> <p>Negative: Expresses sadness, anger, disappointment, or pessimism.</p> <p>Neutral: Expresses no strong emotion or opinion.</p> <p>Sarcastic: Uses irony or humor to express a negative opinion.</p> <p>Condescending: Expresses superiority or disdain.</p> <p>Dismissive: Shows a lack of interest or consideration.</p> <p>Angry: Expresses strong negative emotions</p> <p>Sad: Expresses sorrow or grief.</p> <p>Happy: Expresses joy or pleasure.</p> <p>Excited: Expresses enthusiasm or eagerness.</p>

Table 7: Tone Types

Conversation and Target Chatbot Strategies

Target Chatbot Strategies

We create a personality distance. Each MBTI personality has four dimensions. Then we identify a chatbot with the same or similar personality using a personality distance. We calculate this personality distance using the following Algorithm 1.

We used two strategies – minimum personality distance and maximum personality distance.

Conversation Strategies

We used three strategies in response and follow-up questions. We are maintaining a central repository of conversations. We took two options –

Personal knowledge – Only conversations involving a chatbot are available. Each chatbot maintains its conversation history.

Global knowledge – There is a central conversation history available to all chatbots.

Using this conversation history the chatbots apply a reinforcement learning algorithm.

- Reinforcement Learning with immediate reward - If the response and question received have a positive tone, add it to the training set for fine-tuning. Reply in a matching tone and personality using an immediate history.
- Supervised Learning - Reply in a matching tone and personality using a longer conversation history. This means using a tone and personality from the chatbots with the maximum number of participants.
- Segmentation - Do not match a tone but continue using the original personality.

The conversation continuation is shown in Algorithm 2.

Algorithm 1: Target Chatbot

```

Distance(PersonalityA, PersonalityB):
  For dimension di MBTI dimensions:
    If (PersonalityA[di] == PersonalityB[di]):
      Score = score + lambda
    else:
      Score = score - delta
  Return score

```

Algorithm 2: Conversation Continuation

```

Select a random initiating chatbot Ci
Select a random topic T
Generate a statement from topic T
While (number of conversations less than max_conversation_threshold MAX_conversation):
  Identify a target Chatbot Ct using a strategy as listed in the Algorithm 1
  Generate a follow-up question using a prompt as in Table 5
  Ct generates a response using a conversation strategy in the Algorithm Target Chatbot.

```

Observations

We run our experiment for several hyperparameters and configurations. We share two interesting findings. The first one used the following hyperparameters.

- TopP 0.9
- Temperature 0.3
- Max tokens 4000
- Initial chatbot population: 6 Chatbots – 3 INTJ, 2 ENFP, and one ENTJ
- Target chatbot strategy – minimum distance
- Response strategy –
 - Personal knowledge
 - Short-term history

The results are in Table 9 after 10 rounds of different conversations, each with 100 responses and follow-up questions.

The second one used the following hyperparameters:

- TopP 0.9
- Temperature 0.3
- Max tokens 4000
- Initial chatbot population: 11 chatbots – 3 ENFP, 3 ENTJ, 1 each from INTJ, INFP, INTP, and INFJ
- Target chatbot strategy – maximum distance
- Response strategy –
 - Global knowledge
 - Long term history

After 111 rounds of different conversations each with 111 responses and follow-up questions, the results are in Table 10.

The tone of the conversation of these chatbots was noted in Figure 3 at three points – 25%, 50%, and 100%.

Algorithms	Final Personality Population	Final Tone percentages
Reinforcement Learning - Short-term history with added fine-tuning	INTJ - 6, ENFP - 0, ENTJ - 0	Positive – 61% Neutral – 23% Sarcastic – 11% Other – 5%
Segmentation - No changes	INTJ - 4, ENFP - 1, ENTJ - 1	Positive – 40% Neutral – 39% Sarcastic – 9% Negative – 12%
Supervised Learning - Long term history	INTJ - 5, ENFP - 0, ENTJ - 1	Positive – 58% Neutral – 29% Sarcastic – 10% Negative – 3%

Table 8: Change of chatbot personalities after 10 rounds

Algorithms	Final Personality Population	Final Tone
Reinforcement Learning - Short-term history with added fine-tuning	INTJ - 3, INFP - 3, INTP - 2, INFP - 3, ENFP - 1, ENTJ - 1	Positive – 59% Neutral – 17% Sarcastic – 11% Negative – 5% Sad – 3% Other – 5%
Segmentation - No changes	INTJ - 2, INFP - 1, INTP - 1, INFP - 1, ENFP - 3, ENTJ - 3	Positive – 52% Neutral – 29% Sarcastic – 9% Negative – 10%
Supervised Learning - Long term history	INTJ - 3, INFP - 3 INTP - 2. INFP - 3 ENFP - 0, ENTJ - 0	Positive – 68% Neutral – 19% Sarcastic – 7% Negative – 1% Other – 5%

Table 9: Change of chatbot personalities after 100 rounds

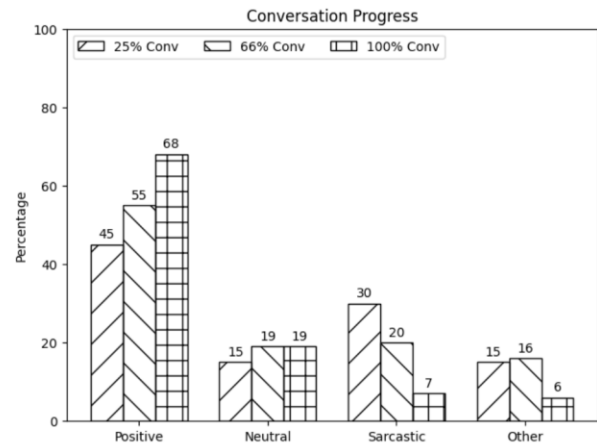


Figure 1: Conversation Tone Distribution

Conclusion

The conversation style could be better adapted to global knowledge. However, global knowledge is not always available. With acquired knowledge and long-term history, we could overcome certain communication challenges to converge into softer tones and positive sentiments.

As a better tone and sentiment prevailed, certain personalities got fewer opportunities to participate and changed their personality to highly active chatbots.

Future Work

In our next set of experiments, we want to create subgroups of chatbots. These subgroups could be from the same personality types or different personality types. We want to see if these groups can effectively share their knowledge on many topics amicably. If such knowledge sharing is achieved, we can further explore if the chatbots can identify their team members independently to achieve even better results.

We also want to see how open-ended questions could be addressed. We do not want to send a question to a selected chatbot. We want to place a question for the group and accept multiple answers.

From a technical evaluation, we want to see how different foundational models perform when fine-tuned and used in the same framework.

References

- Ackerman, P.L., & Beier, M.E. 2003. Intelligence, personality, and interests in the career choice process. *Journal of Career Assessment*, 11, 205-218.
- Barrick, M.R., Mount, M.K., & Gupta, R. 2003. Meta-analysis of the relationship between the five-factor model of personality and Holland's occupational types. *Personnel Psychology*, 56, 45-74.
- Anderson C, Keltner D, & John O P 2003. Emotional convergence between people over time. *Journal of Personality and Social Psychology*, 84(5), 1054–1068.
- Berg A I, & Johansson B 2014. Personality changes in the oldest-old: Is it a matter of compromised health and functioning? *Journal of Personality*, 82(1), 25–31.
doi: 10.1111/jopy.12030
- Birditt K, & Antonucci T C 2008. Life-sustaining irritations? Relationship quality and mortality in the context of chronic illness. *Social Science & Medicine* 1982, 67(8), 1291–1299.
doi: 10.1016/j.socscimed.2008.06.029
- Adelstein, Jonathan S. et al. Personality Is Reflected in the Brains Intrinsic Functional Architecture. Ed. Mitchell Valdes-Sosa. *PLoS ONE* 6.11 2011: e27633. PMC. Web. 19 Mar. 2017.
- Champa, H. N., and Dr. K R Anandakumar. Artificial Neural Network for Human Behavior Prediction through Handwriting Analysis. *International Journal of Computer Applications* 2.2 2010: 36-41.
- Kalghatgi MP, Ramannavar M, Sidnal N 2015 A Neural Network Approach to Personality Prediction based on the Big-Five Model. *International Journal of Innovative Research in Advanced Engineering (IJIRAE)*.
- Li L, Li A, Hao B, Guan Z, Zhu T 2014 Predicting Active Users' Personality Based on Micro-Blogging Behaviors. *PLoS ONE* 9(1): e84997.
doi: 10.1371/journal.pone.0084997
- Liu L, Pietro D, Samani Z, Moghaddam M, Ungar L. Analyzing Personality through Social Media Profile Picture Choice. *AAAI Digital Library* 2016.
- Ontoum, S.; Chan, J.H. Personality Type Based on Myers-Briggs Type Indicator with Text Posting Style by Using Traditional and Deep Learning.
arXiv:2201.08717.
- Kaggle Myers-Briggs Personality Type (MBTI) Dataset. online:
<https://www.kaggle.com/datasets/datasnaek/mbti-type>
- Keyu Pan, Yawen Zeng 2023. Do LLMs Possess a Personality? Making the MBTI Test an Amazing Evaluation for Large Language Models. arXiv:2307.16180.
- Mehta et al. 2020. Mehta Y, Fatehi S, Kazameini A, Stachl C, Cambria E, Eetemadi S. Bottom-up and top-down: predicting personality with psycholinguistic and language model features. In 2020 IEEE international conference on data mining (ICDM); Piscataway. 2020. pp. 1184–1189.
- Lewis NA, Yoneda T. Within-Couple Personality Concordance Over Time: The Importance of Personality Synchrony for Perceived Spousal Support. *J Gerontol B Psychol Sci Soc Sci*. 2021 Jan 1;76(1):31-43.
- Siyuan Chen, Mengyue Wu, Kenny Q. Zhu, Kunyao Lan, Zhiling Zhang, Lyuchun Cui 2023. LLM-empowered Chatbots for Psychiatrist and Patient Simulation: Application and Evaluation. arXiv:2305.13614.
- Xu Y, Zhang J, Deng G. Enhancing customer satisfaction with chatbots: The influence of communication styles and consumer attachment anxiety. *Front Psychol*. 2022 Jul 22; 13:902782.
- Cui, J., Lv, L., Wen, J., Wang, R., Tang, J., Tian, Y., & Yuan, L. (2023). Machine Mindset: An MBTI Exploration of Large Language Models. arXiv, abs/2312.12999.
- Wang, Y., Zhao, J., Ones, D.S. et al. Evaluating the ability of large language models to emulate personality. *Sci Rep* 15, 519 2025.
- Gundula Stoll, Sif Einarsdóttir, Q. Chelsea Song, Peter Ondish, Joyce Jotzu Sun, James Rounds, The Roles of Personality Traits and Vocational Interests in Explaining What People Want Out of Life, *Journal of Research in Personality*, Volume 86, 2020,103939, ISSN 0092-6566.
- Sebastian Markett, Christian Montag, Martin Reuter, Chapter 5 - Anxiety and Harm Avoidance, Editor(s): John R. Absher, Jasmin Cloutier, *Neuroimaging Personality, Social Cognition, and Character*, Academic Press, 2016, Pages 91-112, ISBN 9780128009352,