Capturing Individuals' Communication Styles Using Large Language Models

Yuya Ishikawa[1], Kyosuke Tsubaki[1], Kazuki Kitagawa[1], Kazuya Kanno[1], Nanami Iwahashi[2], Ciaran Regan[2], and Mizuki Oka[2]

[1]Gaudiy Inc., Japan
[2]University of Tsukuba, Japan
mizuki@cs.tsukuba.ac.jp

Abstract

The emergence of large language models (LLMs) has created new opportunities for creating avatars that display realistic human behaviour. One aspect of this behaviour is modelling how people communicate via text, which has so far been limited to matching celebrities' communication patterns through model fine-tuning. This work presents a novel architecture that allows for the creation of individually personalised avatars, without the need for model fine-tuning. Our approach uses a GPT-4 powered Interview Agent to collect data, extracts key conversational features (such as emoji usage and catchphrases), and combines these with the avatar's memories to generate responses that mimic the individual's unique communication style. Our results demonstrate that the extracted conversational features effectively distinguish between individuals, and that avatars built using these features successfully mimic some key communication patterns of their human counterparts, although more research is required regarding the choice and alignment of some of these features. In an era where human-AI interaction is ever-increasing, this work contributes to aligning the communication style of avatars with that of humans.

Introduction

The advent of large language models (LLMs) has opened up new possibilities for creating AI avatars that mimic the conversational style of specific individuals, such as celebrities or fictional characters. There is a growing demand, particularly in the entertainment industry, for avatars that enable more natural and engaging interactions with users. A world where AI and humans coexist in this manner is just around the corner, and it is important for the field of artificial life (ALife) research to recognize the challenges that may arise in such a world (Bedau et al., 2000; Bulitko et al., 2019).

Several prior studies have focused on creating AI models that can mimic specific characters and behaviour Li et al. (2023a); Wang et al. (2023); Yan et al. (2023); Yu et al. (2024); Li et al. (2023b); Salemi et al. (2023); Liu et al. (2024). One notable example is ChatHaruhi (Li et al., 2023a), which focuses on creating AI models that imitate fictional characters from anime and novels. The ChatHaruhi paper employs an approach where a large corpus of a character's dialogue is used to fine-tune a LLM. While this approach works well for fictional characters where a large corpus of conversation data is available, this method has not been extended to mimic real individuals. Additionally, this method relies on fine-tuning pre-trained models, which requires a large amount of data and computational resources. On the other hand, (Park et al., 2023) introduce generative agents, which demonstrate impressive human-like behaviours using the agent's persona, memories, and the conversation context, without the need for fine-tuning or reinforcement learning. However, these generative agents are not designed to represent any particular individual and it is not clear how accurately they can mimic any given real person.

In this study, we introduce an agent architecture which aims to mimic the conversation style of real individuals. Similar to (Park et al., 2023), the proposed architecture is composed of a memory and reflection component, however, to guide how the avatar responds, a set of conversation features is also implemented. These features aim to capture the conversation style of the individual, such as their tone and tendency to use emojis. The features are then inserted into the prompt of the avatar, in combination with their memories, creating a response that is representative of the individual, without the need for model fine-tuning.

One of the challenges with modelling individuals who are not well-known is the lack of available data, which makes feature extraction difficult. To overcome this limitation, we implement a GPT-4 powered Interview Agent which gathers data about the individual through a series of interviews. In this interview process, the Interview Agent asks a range of questions to the individual before summarising the responses and extracting the conversational features, automating the entire agent creation process.

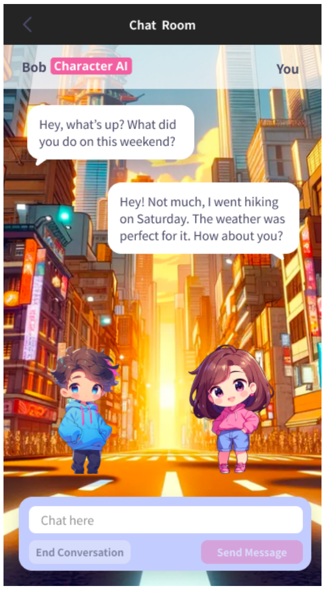

Figure 1: The proposed avatars implemented in a prototype mobile application. Human users engage in conversations with AI avatars that represent real individuals. The left character represents the AI avatar and the right character represents the actual human user. Hundreds of users have already engaged with this prototype.

To analyse the extent to which the proposed avatars can capture how the underlying individual communicates, data for several individuals were collected, and their conversational features were extracted using a LLM. Our results show that LLMs can distinguish between the styles of communication of different individuals. Furthermore, comparing the output of two distinct avatars, we show that the avatars can accurately mimic aspects of the individual they portray.

This research is particularly relevant to the ALife community as it contributes to the development of artificial agents with lifelike characteristics, a key goal in ALife. By focusing on mimicking features of an individual's writing style, our approach aligns with the aim of understanding and simulating the essential properties of living systems. Furthermore, the interaction between these AI avatars and human users represents a form of artificial social interaction, another area of interest within ALife. Our work explores how artificial agents can be designed to foster more natural and engaging conversations, with implications for social AI and human-AI relationships. Moreover, AI avatars have been implemented by several companies such as Replika and Gaudiy Inc., creating software which allows humans to develop friendships with AI. The particular architecture used in this work has already been implemented into a prototype product by Gaudiy Inc., a company specializing in entertainment. This prototype, shown in Fig 1, allows users to converse with avatars modelled after popular idols and celebrities. Currently, this feature is in a beta version with limited access to users but has already shown promising results in terms of user engagement and satisfaction (Gaudiy Inc., 2023).

In summary, this paper makes the following contributions:

- We show that personalized avatars can be created for individuals who are not widely known or famous, by leveraging an interview-based approach.
- We demonstrate that LLMs are capable of extracting several conversational features that capture an individual's unique conversation style.
- We show that the avatars created using the extracted features can represent some of the distinctive communication styles of the individual they are intended to portray.

Method

To create avatars which can closely mimic the conversational style of individuals, the following architecture is introduced. Similar to the work of (Park et al., 2023), the avatar architecture is composed of memory and reflection components, and makes use of in-context learning to give personality to the avatars, without the need for fine-tuning.

There are two phases to the approach. First, an avatar creation phase is used to capture data about an individual's style of communication. This involves the implementation of an Interview Agent which engages in conversation with the target individual to collect conversation data. After acquiring this data, key features about the individual's style of communication, such as their use of emojis, are extracted. In addition, the conversation with the Interview Agent is summarised and stored as an initial long-term memory for the avatar.

The second phase is the avatar interaction phase, in which an end-user interacts with the avatar. In this phase, the avatar receives messages from the user and creates a response based on the individual's features, as well as its long-term memory. In addition, the avatar has access to a short-term memory which stores the current conversation history. When the conversation with the user ends, this conversation history is summarised and added to the long-term memory. In this way, the architecture makes use of both a long-term and short-term memory system. This memory system, in combination with the key conversation features, is what facilitates the avatars to mimic the individual they aim to portray.

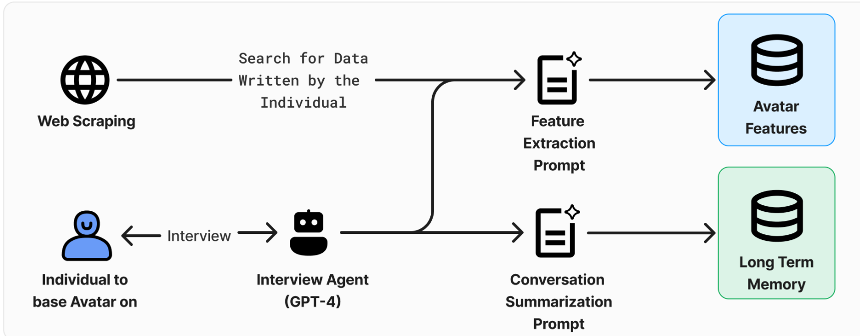

Figure 2: **Avatar Creation**. Data is collected through a combination of scraping the web for text written by the individual, or through an interview with the Interview Agent. This data is then used to extract key features using a LLM and to create an initial long-term memory buffer.

Avatar Creation

The first phase of the proposed approach involves collecting a corpus of the user's conversation data, extracting important features which describe the individual's style of conversation, and then summarising the data into an initial memory for the avatar. An overview of avatar creation is shown in Fig 2.

Data Acquisition via Interview Agent To estimate the features which capture the individual's communication style, a large amount of the individual's communication data is required. First, the internet is scraped to see if the individual has any text data available online, such as blog posts, etc. If there is not enough data available for feature extraction, more data may be acquired by interviewing the individual directly. To this end, an Interview Agent is implemented to converse with the user and gather their responses. This Interview Agent is built with OpenAI's GPT-4, with the prompt shown in Fig 3. In this work, between 50 and 100 interview responses were used to estimate the individual's features.

Feature Extraction Although there are many existing machine learning methods for feature extraction (Abbasi et al., 2022), author identification through quantitative stylistic analysis has been mainly focused on non-communicative texts such as blogs and books. For communication-based text in particular, judging author identification is difficult using this approach, given that expression can change depending on the context before and after the conversation and the subject speaking. To this end, we make use of LLMs to extract features which capture the individual's communication patterns, motivated by the ability of LLMs to capture personalised dialogue (Lotfi et al., 2024). The following conversational features were specifically selected to model the individual's communication patterns, motivated by their importance in distinguishing individuals (Jin and Murakami, 1993; Jin and Jiang, 2012; Marengo et al., 2017):

```
You are a professional interviewer. You will now
be asking questions about {user_name} to learn
more about them, on the theme of {theme}. Please
think of the first question.

Constraints:
- Being straightforward can be stiff, so while
being formal, please speak in a slightly relaxed
manner.

Output Format: Only your message should be written
in the output, do not include "You: ", etc.
```

Figure 3: The prompts used for the Interview Agent. `user_name` refers to the name of the individual, `theme` refers to the topic being discussed and `conversation_history` refers the the previous responses in the conversations.

Frequency of Emojis Whether emojis are used or not.

Catchphrases Whether a certain phrase is often repeated.

Sentence-Final Particles The frequency of sentence-final particle usage (e.g., "isn't it?", "right?").

Trends in Punctuation Appropriateness of punctuation and reading marks as in a typical document.

Trends in Sentence Length The number of characters used in sentences.

Tone Whether the tone of voice is casual.

```
Please identify specific features of the sentences
from the given CSV file and output them in a
bulleted list with examples. Please make sure that
when those features are combined, the person can
be identified.

Feature Indicators:
- Frequency of Emojis and Commonly Used Emojis
- Catchphrases
- Sentence-Final Particles (e.g., 'isn't it?',
'right?')
- Trends in Punctuation
- Trends in Sentence Length
- Tone (e.g., honorific, casual, dialect)
- Opinion Expression
- Tendency to Express Emotion
- Estimation of Age
- Estimation of Gender
- IQ and Breadth of Knowledge
- Values and Ideas
```

Figure 4: The prompt used to extract the key conversational features from the user's conversation data. The conversation data is input as a CSV file.

Opinion Expression The tendency to speak based on beliefs rather than facts.

Tendency to Express Emotion The presence of emotional language and positive expressions.

Estimation of Age An estimation of the age of the individual.

Estimation of Gender An estimation of the gender of the individual.

IQ and Breadth of Knowledge The presence of emotional language and positive expressions.

Values and Ideas The presence of emotional language and positive expressions.

Specifically, GPT-4 is prompted with a CSV file containing the user's conversation data and asks to estimate each of the above features, with the prompt shown in Fig 4.

Memory The avatar architecture implemented in this work is similar to that of (Park et al., 2023). That is, the architecture is composed of a perception and memory component, as well as the ability to reflect on past memories. Unlike the architecture of (Park et al., 2023), however, a distinction is made between short-term and long-term memories. In the proposed architecture, short-term memories are comprised only of the current conversation and provide context for the agent. To give the agents a more long-term understanding, both conversations between the individual and the Interview Agent and those between the user and the avatar are summarised and stored as long-term memories, respectively. These long-term memories are stored in a vector database and retrieved when the avatar formulates a response.

```
You are the interviewer. Pick out the important
information and briefly summarize what kind of
person {user_name} is. Below is the history of
your conversation with {user_name}.

Conversation History: {prev_conversations}

Rules:
- The conversation is over.
- Do not write anything that invades privacy.
```

Figure 5: The prompt used to summarize the conversation.

When the avatar is first created, the text collected by the Interview Agent is summarised using the prompt shown in Fig 5 and is stored as the avatar's initial long-term memory. Unlike the prompts for feature extraction, this prompt summarizes the text content, independent of the style of conversation or communication patterns of the individual.

Avatar Interaction

Once the features have been extracted and the avatar's initial long-term memory has been created, a user can then interact with the avatar. An overview of the avatar interaction architecture is shown in Fig 6.

Prompting When a human user interacts with the avatar, the user's text input acts as perception for the avatar, with the conversation added to the avatar's short-term memory. The avatar's response is created using the prompt shown in Fig 7, which takes multiple inputs such as the current conversation theme, the avatar's conversational style and the avatar's memories. In particular, the avatar's entire short-term memory, which contains the conversation history is included, as well as 10 long-term memories, which have the highest cosine similarity with the user's input message.

Reflection After each conversation session between the user and avatar, which is typically comprised of between 3 to 5 message exchanges, the avatar engages in reflection, summarising the conversation. The prompt used is the same as the Interview Agents's summary of

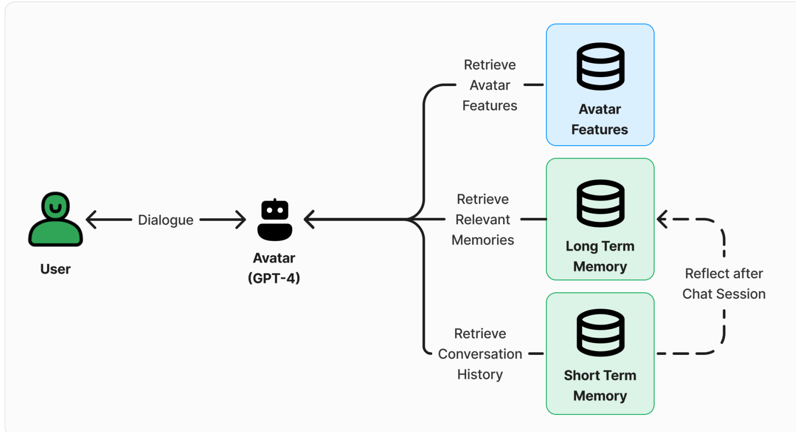

Figure 6: **Avatar Interaction**. The user sends a message to the avatar, who responds based on the individual's features, relevant long-term memories, and the current conversation history. After the conversation ends, the agent reflects on the conversation, summarising the messages. This reflection elevates short-term memories into long-term memory.

the conversation (Fig 5) and aims to capture key details about the past conversation. These reflections are then stored in the long-term memory to be used in future dialogue responses.

Evaluation

To verify how well LLMs can extract the conversational features, and to understand how well these features can be used to represent the underlying individual, the following two research questions are proposed:

1. Can an individual's communication characteristics be estimated based on LLM-extracted features?

2. Can AI avatars based on these features represent the individuals they aim to portray?

To answer these research questions, we extracted the conversational features using GPT-3.5 Turbo with few-shot learning, however, we specifically disregarded the following four features: "Estimation of Age", "Estimation of Gender", "IQ and Breadth of Knowledge", "Values and Ideas" as these features proved difficult to measure and compare. The presence of each of the features is detected sentence-by-sentence, with the LLM reporting the values "Yes", "No", or "None", except for the feature "Trends in Sentence Length", where the LLM reports the number of characters used instead. An example for the feature "Frequency of Emojis" is shown in Fig 8.

RQ1. Can an individual's communication characteristics be estimated based on LLM-extracted features? To analyse whether LLMs can be used to extract characteristics of an individual's conversation style, we measure the extent to which the LLM can extract different values for different individuals. This is motivated by the fact that feature values are expected to vary significantly between individuals, and the LLM should be able to capture this variation.

Table 1: Multiple Analysis of Variance results

Group	Value	F-Value	$Pr > F$
W	0.1546	26.23	0.0000
P	1.4079	21.26	0.0000
H	2.6234	32.55	0.0000
R	1.3117	130.58	0.0000

Table 2: Effect size for each feature

Feature	Effect Size
Frequency of Emojis	0.538
Trends in Sentence Length	0.400
Tone	0.252
Trends in Emotional Expression	0.200
Trends in Punctuation	0.174
Sentence Structure	0.093
Sentence-Final Particles	0.083
Catchphrases	0.062

To measure if the LLM creates distinguishable features for different individuals, human conversational data for 30 individuals with over 50 messages each were collected, acting as the ground truth.

The Multivariate Analysis of Variance (MANOVA) (French et al., 2008) test was con-

```
You are a individual named {name}. You are having
a conversation about the subject {theme} and
{user_name} is replying {message}. Please reply
based on the rules and the conversation so far.

Rules:
- Respond naturally based on your conversation
history and memories.
- You are having a conversation with your friend
on a chat-style app like WhatsApp.
- Take into account your personality and tone of
voice.
- Do not lie.
- Do not write anything that invades privacy.
- Always refer to the user by {user_name}.
- If you are asked to make a personal commitment,
never say no, and give a curt response that is
appropriate to the topic at hand.
- Assume that you understand the words in the
{theme} topic.

Your Character Information:
Name: {name}
Profile: {bio}
Personality: {personality_type_description}
Tone: {writing_style_description}

User Information:
Name: {user_name}

Related Memories: {relevant_memories}

Conversation History: {prev_conversations}

Output:
Please start the output by continuing
with the following.

You:
```

Figure 7: The prompt used to generate the avatar's response. `name` refers to the avatars name, `theme` refers to a user-selected conversation topic, `user_name` refers to the user's name, `bio` refers to personal information about the avatar's individual such as birthday and location, which may have been collected previously, `personality_type_description` refers any other additional information which may be included about the avatar, `tone` refers to conversation style of the avatar from the extracted features, `relevant_memories` refers to long-term memories similar to the input message and `prev_conversations` refers to conversation history which is stored in the avatar's short term memory.

```
You are to evaluate the given sentences in the
terms of 'frequency of emoji'. The evaluation is
done on a 'Yes' or No' basis, with 'Yes'
indicating a high probability of conformity to the
evaluation criteria, and ``No'' indicating a low
probability of conformity. If it is difficult to
judge, choose 'None'. Below are the evaluation
criteria.

Evaluation Point: Whether emojis are used or not.

Please evaluate sentences based on this metric and
format the output as {'metrics_name': 'frequency
of emoji use', 'evaluation': 'Yes' | 'No' |
'None', 'reason': str}. Include the reason for the
'None' decision.

Sentence: 'Good morning!😊 Let's do our best today
too!🌈 💪'
{'metrics_name': 'Emoji usage frequency',
``evaluation'': 'Yes', 'reason': 'Because the
emoji 😊, 🌈, 💪 is used at the end of the
sentence'}

Sentence: 'Good morning! Let's have a great day!'
{'metrics_name': 'Emoji usage frequency',
'evaluation': 'No', 'reason': 'Because no emoji is
used anywhere in the sentence'}

Sentence: 'The live show seemed great! I wish I
could have seen it too! It must have been
wonderful ✨ Let me know about the next live show!
Good night 🌙✨'
```

Figure 8: The prompt used to evaluate the emoji usage feature in the individual's conversation data. Similar prompts are used for the other features, except for sentence length which is evaluated as the number of characters used.

ducted, enabling us to understand if there exists a significant difference between the extracted features of the individuals. Specifically, we used four statistics that are part of the popularly reported test statistics for MANOVA: Wilks' lambda (W), Pillai's trace (P), Hotelling-Lawley trace (H), and Roy's greatest root (R). In addition, the effect sizes (η^2) were determined for all features.

To carry out this analysis, each feature output from the LLM was converted into a numerical value. That is, if the LLM responds "yes", indicating a feature was detected in an individual's data, then a value of 1 was assigned to that feature. Similarly, "No" was assigned to 0 and "None", indicating the inability to answer, was assigned to 0.5. Additionally, the feature "Sentence Length", which ranges between 0 to 200 was normalized between 0 and 1.

The table 1 shows that the p-value for all tests was less than 0.0001, indicating that the difference in features between the individuals is statistically significant. Furthermore, p-values less than 0.0001 indicate that individual-to-individual comparisons are also significant. Overall, these results suggest that GPT-3.5 Turbo can distinguish between the unique conversation characteristics of different individuals.

Furthermore, Wilk's lambda W ranges from 0 to 1, with lower values indicating larger effect sizes. The value $W = 0.1546$ is close to zero and therefore the effect size for all features is large. This indicates that most of the variance in the extracted features was due to actual differences between individuals. This further highlights that the LLM can capture significant and distinguishable features that characterise the communication patterns of the 30 individuals analysed.

Table 2 shows the effect size for each feature. The effect sizes are larger for the frequency of emojis, trends in sentence length, tone, and trends in emotional expression, in that order, indicating that these features have the strongest contribution in differentiating the unique communication characteristics of the individuals. This result is sensible as emoji usage, for example, is expected to vary widely across the population (Ljubešić and Fišer, 2016) and can therefore be used to distinguish individuals. Other features such as catchphrases, on the other hand, had a low effect size, which may be attributed to catchphrases not appearing in the test data. Although some features have a low effect size, they may still be beneficial in creating realistic avatars. For instance, many individuals in the dataset use sentence-final particles, such as "isn't it?". It would be important to ensure that the avatars also display this behaviour.

In summary, although each feature has a different impact on the distinguishability of individuals, communication characteristics can be estimated based on LLM-extracted features.

RQ2. Can avatars built from these features accurately represent the individual they aim to portray? To analyse if avatars can mimic the features of an individual's style of communication, conversation data was collected for 2 individual-avatar pairs and their features were evaluated. The features obtained by the prompt shown in Fig 8 are then converted into numerical values and averaged over all of the sentences to give a single average score for each feature for the two individual-avatar pairs. The way of converting

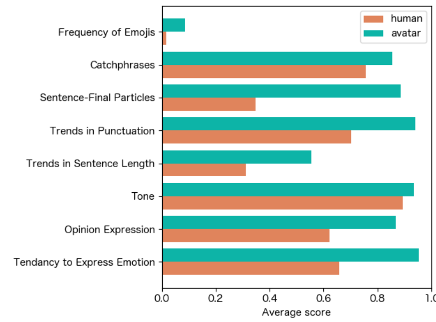

(a) Individual A

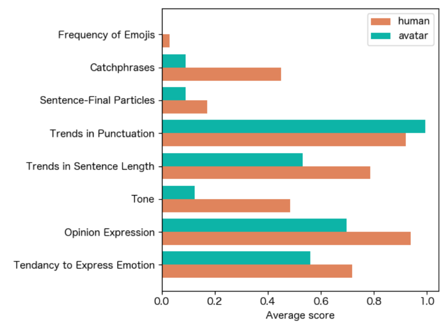

(b) Individual B

Figure 9: Comparison of the features for Individual A (a) and Individual B (b) and their respective AI avatars.

a feature into a numerical value is similar to that RQ1, however, "None" values are ignored instead of being converted to 0.5. Moreover, the average of the scores for all messages was calculated.

The average feature values for both agent-individual pairs are shown in Fig 9. These results show that the avatar captures some features better than others. For instance, in both Fig 9a and 9b, the individuals tend not to use emojis, shown by the low score for the Frequency of Emojis feature, and this is reflected in the behaviour of the avatar. This is particularly significant as this feature has a large effect size and therefore is important for capturing the specific way an individual communicates.

For Trends in Sentence Length, however, the avatar has quite different results from the corresponding individual. Moreover, discrepancies between the accuracy of the Tone feature highlight the limitations of current LLMs. Avatar A is easily able to capture the polite and formal language of individual A. On the other hand, individual B tends to use informal language, which the LLM struggles to replicate. One reason for this may be the fine-tuning of GPT-4 to generate polite responses.

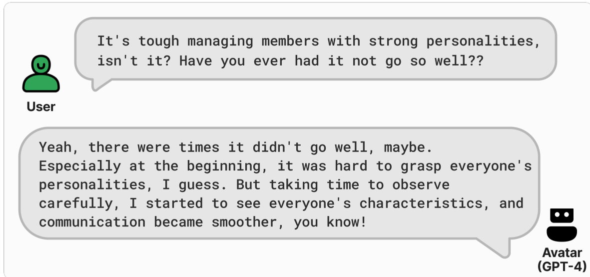

(a) Good Example. Individual A tends to use sentence-final particles such as "you know". The avatar captures this behaviour, although perhaps too many sentence-final particles are used.

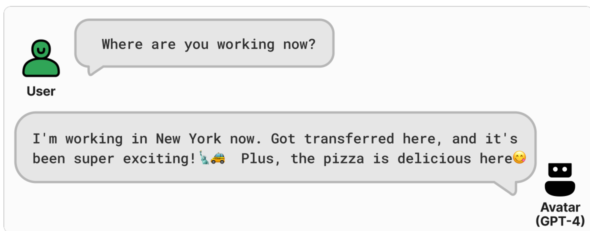

(b) Bad Example. Individual A does not use emojis, but the avatar incorporates them in a response to a user.

Figure 10: Example of a conversation between the avatar of Individual A and users.

Examples of the avatar responses of individual A are shown in Fig 10. These examples highlight that although the architecture can capture some features, it can still fail to respond appropriately, with the avatar incorrectly using emojis in Fig 10b.

Overall, the avatars of Individual A and Individual B show different feature values, indicating that the AI avatars successfully mimic the characteristics of different individuals, although some features are better captured than others.

Discussion and Limitations

Despite being able to capture a number of the features, further work is still required to align all of the avatar's conversational features to that of the human. To achieve this, one should experiment with different LLMs, which may improve the alignment of the features, such as the avatar's tone, which seems to be constrained by the model. In addition, further research is needed to enable the avatar's style to mimic the dynamics of human communication style, fluctuating based on the situation and mood.

A significant limitation of this work is the reliance on human-selected features, which we posit can capture the communication style of individuals. Despite demonstrating that these features can indeed be used to distinguish between different individuals, there are an infinite number of possible choices for these features and our specific choice may not be optimal. In future work, it would be beneficial to consider a wider range of features, or even select features via a LLM, removing the requirement for a human in the loop. On the other hand, creating high-level features of human communication is precisely the goal of fine-tuning a LLM, highlighting the balance between the high-level feature selection in LLM fine-tuning and the speed and adaptability of in-context learning.

The avatars in this work are currently implemented in a prototype product by Gaudiy Inc. (Fig 1), where users can converse with AI-powered avatars representing famous idols and celebrities, further enhancing human-AI interaction. In future plans, Gaudiy Inc. aims to create personalized avatars for each individual. These avatars can then communicate with one another, fostering the development of digital friendships and connections, with the ultimate goal of these digital bonds manifesting in real-world friendships between the underlying individuals.

Finally, there are several important ethical considerations to be made regarding the creation of personalised avatars. For instance, it is crucial to consider the consent of individuals when creating personalised avatars. Although the individuals participating in this study gave their consent, bad actors may use LLMs to construct avatars for individuals without their consent. Another potential issue is the avatar misrepresenting the individual. The recent creation of "deadbots", chatbots designed to simulate deceased people, has gained notable traction and raised a number of ethical and philosophical questions (Henrickson, 2023). These issues should be seriously considered in the research and application of personalised avatars.

Conclusion

In this work, we introduce a novel avatar architecture which enables the creation of personalised AI avatars without the need for fine-tuning. By collecting conversation data using an Interview Agent and extracting conversational features, avatars can be created that aim to mimic the communication patterns of the underlying individual. We show that interviewing the individual directly can give sufficient information to model their communication patterns. Furthermore, we demonstrate that some of the features extracted using a LLM are distinguishable; that is, the LLM can capture some aspects of the communication style of different individuals. Finally, by comparing the avatar to its counterpart, we show that the avatar can capture several of the communication traits of the individual.

References

Abbasi, A., Javed, A. R., Iqbal, F., Jalil, Z., Gadekallu, T. R., and Kryvinska, N. (2022). Authorship identification using ensemble learning. *Scientific reports*, 12(1):9537.

Bedau, M. A., McCaskill, J. S., Packard, N. H., Rasmussen, S., Adami, C., Green, D. G., Ikegami, T., Kaneko, K., and Ray, T. S. (2000). Open problems in artificial life. *Artificial life*, 6(4):363–376.

Bulitko, V., Doucet, K., Evans, D., Docking, H., Walters, M., Oliver, M., Chow, J., Carleton, S., and Kendal-Freedman, N. (2019). A-life evolution with human proxies. In *Artificial Life Conference Proceedings*, pages 465–466. MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info ....

French, A., Macedo, M., Poulsen, J., Waterson, T., and Yu, A. (2008). Multivariate analysis of variance (manova). *San Francisco State University*.

Gaudiy Inc. (2023). Gaudiy collaborates with sanrio to launch a web3 and ai-generated sns service globally in spring 2024. `https://prtimes.jp/main/html/rd/p/000000050.000035719.html`.

Henrickson, L. (2023). Chatting with the dead: the hermeneutics of thanabots. *Media, Culture & Society*, 45(5):949–966.

Jin, M. and Jiang, M. (2012). Text clustering on authorship attribution based on the features of punctuations usage. In *2012 IEEE 11th International Conference on Signal Processing*, volume 3, pages 2175–2178. IEEE.

Jin, M. and Murakami, M. (1993). Authors' characteristic writing styles as seen through their use of commas. *Behaviormetrika*, 20(1):63–76.

Li, C., Leng, Z., Yan, C., Shen, J., Wang, H., Mi, W., Fei, Y., Feng, X., Yan, S., Wang, H., et al. (2023a). Chatharuhi: Reviving anime character in reality via large language model. *arXiv preprint arXiv:2308.09597*.

Li, C., Zhang, M., Mei, Q., Wang, Y., Hombaiah, S. A., Liang, Y., and Bendersky, M. (2023b). Teach llms to personalize–an approach inspired by writing education. *arXiv preprint arXiv:2308.07968*.

Liu, N., Chen, L., Tian, X., Zou, W., Chen, K., and Cui, M. (2024). From llm to conversational agent: A memory enhanced architecture with fine-tuning of large language models. *arXiv preprint arXiv:2401.02777*.

Ljubešić, N. and Fišer, D. (2016). A global analysis of emoji usage. In *Proceedings of the 10th web as corpus workshop*, pages 82–89.

Lotfi, E., De Bruyn, M., Buhmann, J., and Daelemans, W. (2024). Personalitychat: Conversation distillation for personalized dialog modeling with facts and traits. *arXiv preprint arXiv:2401.07363*.

Marengo, D., Giannotta, F., and Settanni, M. (2017). Assessing personality using emoji: An exploratory study. *Personality and Individual Differences*, 112:74–78.

Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P., and Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–22.

Salemi, A., Mysore, S., Bendersky, M., and Zamani, H. (2023). Lamp: When large language models meet personalization. *arXiv preprint arXiv:2304.11406*.

Wang, X., Fei, Y., Leng, Z., and Li, C. (2023). Does role-playing chatbots capture the character personalities? assessing personality traits for role-playing chatbots. *arXiv preprint arXiv:2310.17976*.

Yan, M., Li, R., Zhang, H., Wang, H., Yang, Z., and Yan, J. (2023). Larp: Language-agent role play for open-world games. *arXiv preprint arXiv:2312.17653*.

Yu, X., Luo, T., Wei, Y., Lei, F., Huang, Y., Hao, P., and Zhu, L. (2024). Neeko: Leveraging dynamic lora for efficient multi-character role-playing agent. *arXiv preprint arXiv:2402.13717*.