

PlatoLM: Teaching LLMs in Multi-Round Dialogue via a User Simulator

Chuyi Kong^{1,2*}, Yaxin Fan³, Xiang Wan^{1,2}, Feng Jiang^{1,2,4}, Benyou Wang^{1,2}

¹ The Chinese University of Hong Kong, Shenzhen ² Shenzhen Research Institute of Big Data

³ Soochow University ⁴ University of Science and Technology of China

jeffreyjiang@cuhk.edu.cn

Abstract

The unparalleled performance of closed-sourced ChatGPT has sparked efforts towards its democratization, with notable strides made by leveraging real user and ChatGPT dialogues, as evidenced by Vicuna. However, due to challenges in gathering dialogues involving human participation, current endeavors like Baize and UltraChat rely on ChatGPT conducting role-play to simulate humans based on instructions, resulting in overdependence on seeds, diminished human-likeness, limited topic diversity, and an absence of genuine multi-round conversational dynamics. To address the above issues, we propose a paradigm to simulate human behavior better and explore the benefits of incorporating more human-like questions in multi-turn conversations. Specifically, we directly target human questions extracted from genuine human-machine conversations as a learning goal and provide a novel user simulator called ‘Socratic’. The experimental results show our response model, ‘PlatoLM’, achieves SoTA performance among LLaMA-based 7B models in MT-Bench. Our findings further demonstrate that our method introduces highly human-like questioning patterns and rich topic structures, which can teach the response model better than previous works in multi-round conversations.

1 Introduction

Large Language Models (LLMs) such as ChatGPT (OpenAI, 2023) have made great strides in the dialogue domain. Although ChatGPT and its successor GPT-4 (Bubeck et al., 2023) are successful, they remain proprietary and non-replicable. Recent democratization efforts (Taori et al., 2023; Chiang et al., 2023; Dettmers et al., 2024; Ji et al., 2023; Chen et al., 2023), in addition to focusing on distilling the responses of ChatGPT through various means such as self-instruction (Wang et al.,

* Work done on Shenzhen Research Institute of Big Data. Feng Jiang is the corresponding author.

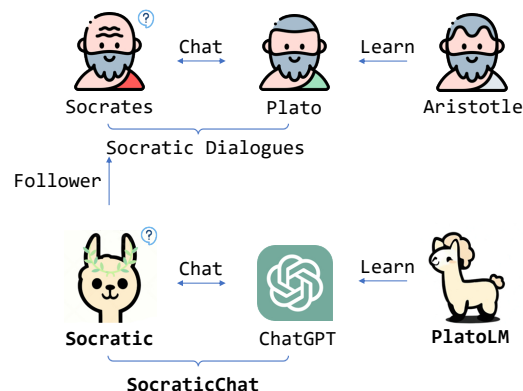


Figure 1: Analogy to Socratic Teaching of Methodology

2023b), also focusing on align ChatGPT with human preferences, as represented by RLHF (Ouyang et al., 2022), RLAIF (Lee et al., 2023), and DPO (Rafailov et al., 2024). However, the most direct human needs are often ignored. We observed that Vicuna (Chiang et al., 2023), which directly employs real human-ChatGPT conversation data for training, consistently shows superior performance across various benchmarks, particularly on multi-round benchmarks. This motivated us to explore how the demand of real humans will affect the capabilities of the response model.

Due to incorporating real users into the construction of human-machine dialogues being costly and, to varying degrees, involving privacy issues, many works like Baize (Xu et al., 2023b) and UltraLM (Ding et al., 2023) leverage ChatGPT for static role-playing to synthesize multi-round dialogue. However, there are still three challenges in utilizing such methods. Firstly, the static simulation **needs extra seeds for each sample** to initiate conversations. Moreover, ChatGPT has been trained as a system agent since its inception, which makes it **difficult to fully learn the patterns of real human questioning**, limiting the diversity of topic structures in real multi-round human-computer interactions. Additionally, al-

Models	Backbone	#Samples	Training Type	MT-Bench	AlpacaEval
LLaMA-2-7b-chat	LLaMA2	1100K	SFT, RL	6.27	71.37%
Vicuna-7b-v1.3	LLaMA2	125K	SFT	-	76.84%
Vicuna-7b-v1.5	LLaMA2	125K	SFT	6.17	-
GPT-3.5	-	-	-	-	81.71%
PlatoLM-7b	LLaMA2	50.73K	SFT	6.29±0.04	81.94%

Table 1: The Performance of PlatoLM in Official AlpacaEval and MT-Bench Benchmarks. More in Appendix G.

though Baize and UltraLM use subtle prompts to instruct ChatGPT as the user, the role-shifted ChatGPT’s instruction-following ability declines (cases in Appendix M.4), which **reduces the robustness of the simulator and requires extensive manual post-processing**.

To address the above issues, we introduce a *trainable* user simulator instead of the *static* ChatGPT one. Technically, the key to our recipe is flipping the learning objective from ChatGPT’s response to real user questions, obtaining a more human-like simulator. Then, we employ the simulator to interact naturally with ChatGPT until the history of the simulator achieves the maximal context length, thereby synthesizing a multi-turn conversation dataset and leveraging it to train the general system agent.

Experiments show that our trainable paradigm is more effective than the static one in teaching response models on multi-turn conversations. Meanwhile, it can transfer domains with seed, scale with many factors, adapt to popular backbones, and preserve ethical friendliness. Upon further analysis, we find that compared to static simulations, the questions in our paradigm are more human-like, leading to richer topic structures in conversations. Moreover, using different backbones as questioners and responders is suitable for our paradigm. The cooperative dialogue between different backbones is reminiscent of Socratic teaching (see Appendix M), where the teacher (Socrates) deepens the students’ (Plato) thinking through a series of probing questions. Thus, we name our questioning model - one that is based on backbones with rich knowledge and fine-tuned with real human prompts - as ‘**Socratic**’, the follower of Socrates. We term the dataset ‘**SocraticChat**’, and the final response model ‘**PlatoLM**’ (see Figure 1). Ultimately, PlatoLM achieved the SoTA performance on the MT-Bench among 7B-scale models based on LLaMA, surpassing GPT-3.5 turbo on the Alpaca-Eval (see Table 1), after aligning the backbone model.

Overall, our contributions are outlined below:

1. We propose a straightforward yet effective **paradigm** for simulating human better. This approach can switch between posing questions without context freely and asking domain-specific questions.
2. We provide various versions of the human-centric multi-round conversation **dataset** (SocraticChat), which extends the scale and diversity of the existing ShareGPT dataset.
3. We train a new assistant **model** (PlatoLM) on SocraticChat, which is superior to other baselines in most comparisons under the same small number of training samples. Furthermore, even with fewer samples (50.7K) and a shorter context length(2048), PlatoLM achieved the best performance among the 7B models on MT-Bench and surpassed GPT-3.5 on the AlpacaEval after being fine-tuned on different backbone pairings.
4. We **find** that a more human-like questioning pattern in dynamic multi-round conversations can better teach the response model compared to the static role-playing, which can be attributed to the natural and rich topic structures human dominate in human-machine dialogue, where they hold topic dominance. Moreover, the interaction between fine-tuned different backbones proves to be more valuable than self-reflection within a single one.

2 Background

Previous works typically focus on leveraging user simulators to generate large amounts of data with limited samples (Asri et al., 2016; Kim et al., 2021; Wan et al., 2022) or enhancing the performance of the assistant’s response through feedback from the user simulator via Reinforcement Learning (Liu and Lane, 2017; Kreyssig et al., 2018; Takanobu et al., 2020) in closed-domain conversations.

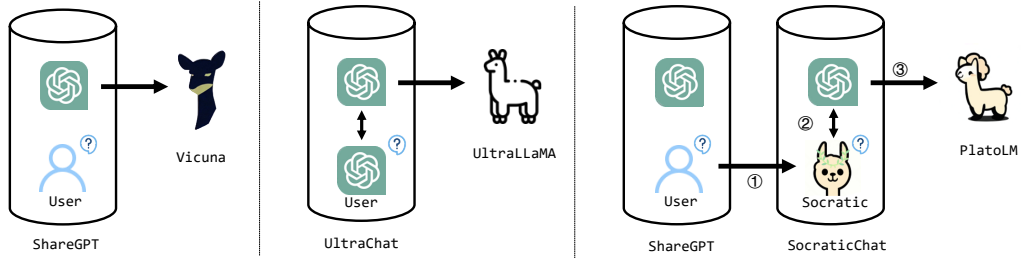


Figure 2: Comparison between Vicuna, UltraLM, and PlatoLM. The commonness of the three models is that they all learn from a *user-system* conversation data. Note that training Socratic and PlatoLM (also for Vicuna and UltraLM) is **symmetrical**; the difference is that the former mimics the *user* and the latter mimics the *system*.

Shifting to open-domain conversations, without a specific task goal among multi-round conversations, it becomes challenging to ascertain user feedback. Consequently, researchers have shown great interest in distilling data with seeds on selected domains via ChatGPT’s static simulation. Baize (Xu et al., 2023b) distilled 100K samples called ‘Self-Chat’, and UltraLM (Ding et al., 2023) distilled 1,468K samples named ‘UltraChat’ based on seeds from humans or even ChatGPT. We argue that while the large quantity of dialogue datasets signifies a substantial contribution, the key challenge lies in aligning the dialogues closer to real human-machine interaction scenarios, while reducing over-reliance on ChatGPT.

3 Methodology

As shown in Figure 2, the pipeline of our methodology consists of three steps: (1) fine-tune the user simulator **Socratic** to raise questions; (2) generate the synthetic multi-round dataset called **Socratic-Chat** via iteratively calling Socratic and ChatGPT; (3) train new system agents **PlatoLM** on the newly produced dataset.

3.1 The User Simulator - Socratic

Unlike previous work that used ChatGPT to simulate users statically, we first built a trainable user simulator to better simulate human needs and interaction behavior.

3.1.1 Data Preprocessing

To train a user simulator that can naturally chat with the machine, we choose ShareGPT, a human-ChatGPT multi-round dataset from Vicuna, as the source. Then, we filter 20K conversations from the original ShareGPT as training samples. The filtering steps specifically include: converting the HTML to Markdown format for rich text to preserve better interaction, proportionally re-

jecting multilingual noise following Vicuna and Koala (Geng et al., 2023), removing some samples where questions were not translation tasks, but with pre- and post-translation languages to avoid the sudden code-switching phenomena on Socratic, and de-duplicating completely duplicated conversations (more details in Appendix F).

Conversation Segmentation. To avoid the forgetting phenomena, we segment the conversation. Particularly, when we split conversations exceeding the maximum context length (2048 tokens) into several segments, in addition to making each segment end with the GPT’s answer instead of the human’s question, to better leverage human’s questions like Vicuna and Koala did, we also ensured that the subsequent segments are contextualized with the GPT’s responses from the prior segment, by padding it at the beginning of subsequent segments. This prevents the questions from containing ambiguous pronouns in the first turn and strikes a balance between raising new questions without context and following up on a previous context. Specifically, unsegmented sessions starting with humans are suitable for the model to learn how can ask new questions without context, while segmented sessions starting with GPT are suitable for enhancing the model’s ability to ask follow-up questions within the previous context.

3.1.2 Training Protocol

In contrast to training a response model, we fine-tune Socratic via masking the questions of real users and accordingly, only calculating their loss to modify the learning objective. To ensure fairness, when fine-tuning the user simulator, a prompt template roughly dyadic to train the response model was employed (see Appendix A), and the parameter settings are consistent with those of other models fine-tuned on LLaMA-7B (Touvron et al., 2023).

3.2 The Conversation Dataset - SocraticChat

Through iteratively interacting between Socratic with the middle model - online GPT-3.5-turbo API, the synthetic multi-round dataset called ‘Socratic-Chat’ was born. Compared to the previous works, our approach has two characteristics: optional seed mode and automatic termination mechanism.

3.2.1 Optional Seed Mode

Due to the carefully designed preprocessing procedure and training protocol mentioned above, using only uniform prompt templates aligned with training, Socratic shows the flexibility to switch between posing questions freely and asking questions in a customized domain. Correspondingly, we define two modes of applying Socratic: free mode and seed mode.

Free-mode refers to the mode that the trainable Socratic freely poses brand new questions at the beginning of the conversation without any context.

Seed-mode is the mode in which Socratic takes the first-round conversation from other sources (i.e. seed conversation) as the context and then follows up questions from the second round. Although free-mode Socratic could be used to generate conversation data without the need to provide context, it is difficult to generate conversation data in a specific domain. To this end, we could use seed-mode Socratic, or similar to UltraLM, directly specifying the topic by adding it to the prompt template of free-mode one (see Appendix A).

3.2.2 Automatic Termination Mechanism

In the open domain, when training simulators, one inevitably encounters the issue of how to terminate the end of the conversation, as the open domain lacks the explicit task objectives found in closed domains. To relieve the issue, we propose an automatic termination mechanism.

Considering humans dominate in human-computer dialogue, we opt to manage the termination of the dialogue on the user side. Specifically, when the context length surpasses the maximal 2048 tokens, we reset the dialogue by clearing its history and initiating a new session, which we call ‘hard control’. Our decision to not emulate Baize’s approach of controlling the conversation’s termination via the prompt template (which we call ‘soft control’) stems from the unique nature of multi-turn conversations. Among the curated training sets, a notable *topic shifting* phenomenon appeared, which makes it challenging to discern if a user’s

halt in asking questions signals the end of a topic or simply a pause. Furthermore, introducing a special token <END> in the final round’s human utterance to mark the dialogue’s termination (following Baize’s approach), will cause the dialogues to be frequently ended within just 1 to 2 rounds. This is because the distribution of conversation rounds in ShareGPT is uneven. Specifically, after removing HTML content, sessions comprising 1 to 10 rounds account for 81.73% of the total, and remarkably, sessions containing 1 to 2 rounds within the 1 to 10 round range make up 53.91%.

3.3 The System Agent - PlatoLM

Following Vicuna’s training schema, we only fine-tune PlatoLM on the synthetic SocraticChat by learning the output of the system agent. Also, we choose training parameters consistent with Vicuna.

4 Experiments

4.1 Baseline Trials

We incorporated the following two types of models as baselines: **(a) Models using simulator-involved data:** Baize (Xu et al., 2023b) and UltraLM (Ding et al., 2023). **(b) Models using user-involved data:** Vicuna (Chiang et al., 2023) is employed as another strong baseline. To ensure fairness, we maintained consistent settings regarding the volume of the training sample (10K), and the training approach (SFT) including hyperparameters and prompt templates for all models, except for data sources (details in Appendix B).

4.1.1 Metrics

Our evaluation metrics encompass both automatic and manual methodologies:

(a) Automatic Evaluations. Given that traditional metrics, such as BLEU (Papineni et al., 2002) and Rouge (Lin, 2004), don’t align well with open-domain dialogue model evaluations, we leveraged widely accepted benchmarks like Vicuna-Bench, Alpaca-Eval (Dubois et al., 2024) and MT-Bench (Zheng et al., 2023) (see Appendix C) to appraise model performance on single and multiple conversation rounds. The rigorous but unstable GPT-4 was used for the judgment. To avoid instability in the GPT-4 output, we evaluated each model 5 times on each benchmark¹ and calculated their mean and standard deviation. Further, to ensure thoroughness, both point-wise and pairwise

¹except on the costly Alpaca-Eval.

Model	Training Dataset	Dataset Type	Alpaca-Eval	Vicuna-Bench	MT-Bench
Baize	Self-Chat	simulator-involved	9.30±1.02%	4.67±0.04	3.95±0.05
UltraLM	Ultra-Chat	simulator-involved	47.57±1.76%	7.72±0.02	4.72±0.02
Vicuna	ShareGPT	user-involved	70.02±1.62%	8.18±0.04	5.91±0.07
PlatoLM (free mode)	SocraticChat	simulator-involved	71.89±1.59%	8.43±0.01	5.33±0.03

Table 2: The Evaluation Results on Popular Benchmark for Baseline Trials (10K) Samples

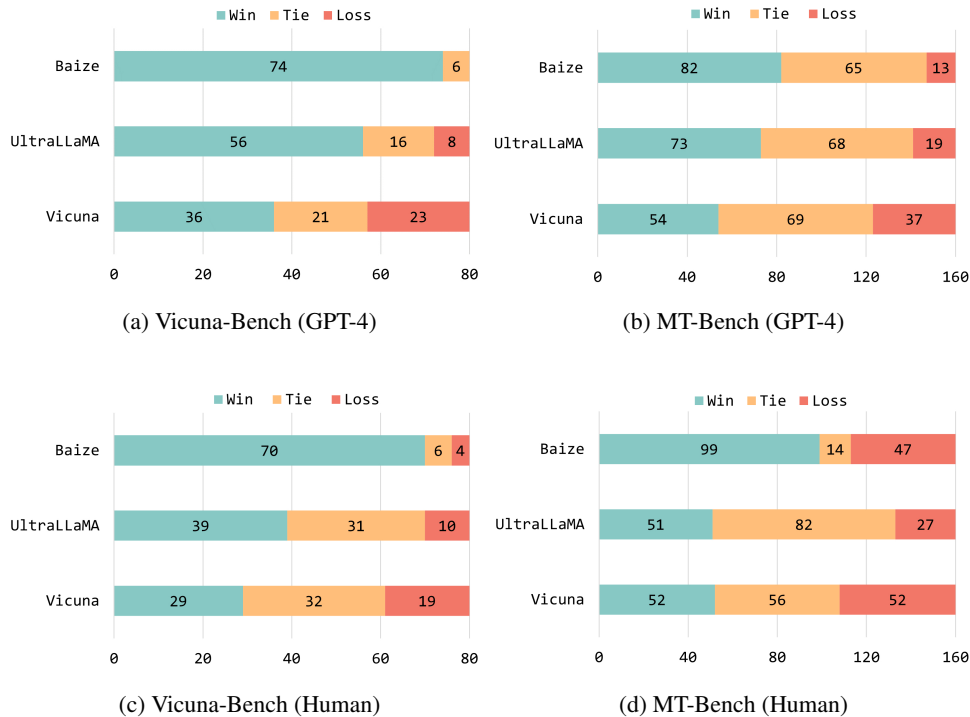


Figure 3: The Automatic and Manual Pair-Wise Evaluations in Vicuna-Bench and MT-Bench for Baselines (10K)

assessments across all baseline models are conducted. **(b) Manual Evaluations.** We recruited four annotators for each benchmark¹ to conduct pairwise evaluations (details in Appendix D).

4.1.2 Results

Overall, for the single-turn benchmark, our model outperforms all baselines. Concerning the multi-round MT-Bench, our model outperforms most baselines including Vicuna in automatic pair-wise evaluation, although it does lag somewhat in automatic point-wise comparison, which may be caused by the penalties of point-wise evaluations towards domains where models falter.

Automatic Evaluation. Figure 3a, 3b and the fourth column in Table 2 present the **pair-wise evaluation** results for our model in comparison with the baseline models. Both on the Vicuna Bench, Alpaca-Eval, and MT-Bench, Our model shows a significant advantage over Baize and UltraLM. Impressively, PlatoLM also surpasses Vi-

cuna (36 wins vs. 23 wins on Vicuna Bench, 54 wins vs. 37 wins on MT-Bench, 71.89% v.s. 70.02% over Davinci003).

In the **point-wise evaluation** on Vicuna-Bench, PlatoLM still maintains a lead over all other baseline models, including Vicuna, scoring 8.43 as compared to Vicuna’s 8.18, as shown in Table 2. However, our model didn’t outperform Vicuna on MT-Bench. After a detailed study of the distribution of the scores on the domain (see Appendix 6a), we discovered why: our model performs badly in math and extraction categories and gets penalized more by the low scores in single answer grading than in pair-wise setup.

Manual Evaluation. To obtain a more reliable and comprehensive evaluation, we further complemented the results with a manual evaluation, and the average scores from four annotators are adopted as the final metric, which is shown in Figure 3c and 3d. Notably, on the Vicuna-Bench, our model demonstrates a high concurrence with the outcomes

User Simulator	Trainable	Used Seeds	MT-Bench	Vicuna-Bench	Alpaca-Eval
ChatGPT	Static	ShareGPT	5.32±0.06	8.24±0.05	66.79±1.66%
Socratic	Trainable	- (free mode)	5.33±0.03	8.43±0.01	71.89±1.59%
Socratic	Trainable	Evol-instruct	5.01±0.04	8.05±0.04	58.42±1.74%
Socratic	Trainable	Dolly	5.57±0.02	8.49±0.03	74.13±1.54%
Socratic	Trainable	ShareGPT	5.65±0.06	8.10±0.05	67.89±1.65%

Table 3: The Automatic Point-Wise Evaluation in Three Benchmarks with Different Seeds of User Simulator(10K)

of the automatic evaluation and significantly outperforms all the baselines. Moving to MT-Bench, our PlatoLM still holds clear advantages over Baize (99 vs. 47) and UltraLM (51 vs. 27), and ties with Vicuna (52 vs. 52). This indicates that our model exhibits competitive performance when constrained to a training dataset of 10K.

4.2 Ablation Studies

To demonstrate the transferability, scalability, and versatility, as well as the ethical friendliness of our paradigm, we conduct the following experiments.

4.2.1 On Different Seeds

In addition to generating conversation data without context in free mode, our trainable user simulator, Socratic, can also use seed conversation to generate domain-specific data. Considering the different speakers in seed conversations, we use the popular Evol-instruct (ChatGPT-to-ChatGPT) (Xu et al., 2023a), Dolly (Human-to-Human) (Conover et al., 2023), and ShareGPT (Human-to-ChatGPT)² to generate corresponding multi-round conversations data and evaluate the performance of the response models trained on these conversations. Moreover, We also involve a static user simulator on the same seed from ShareGPT for a fair comparison (details see Appendix I).

As shown in Table 3, we find that: (1) The response model taught by the seeds involved in human questioning (Dolly, ShareGPT) performs better, which initially demonstrates the transferability of our paradigm. (2) The response model activated by the Socratic simulator and ChatGPT one (the 1st and last row) performs similarly in the single-round dialogue benchmark (Vicuna-Bench and Alpaca-Eval), but the former performed significantly better than the latter in the multi-round dialogue benchmark (MT-Bench).

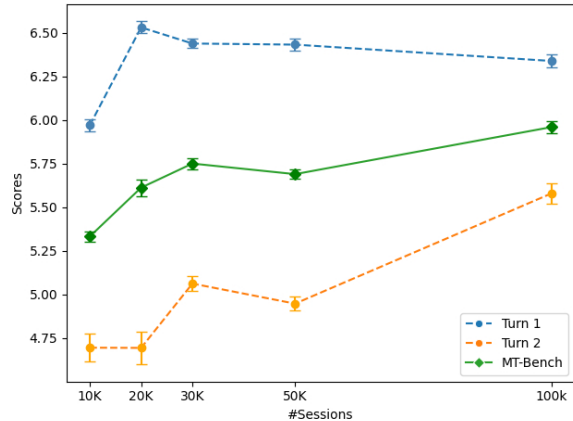


Figure 4: The Impact of Sample Scale on Performance

4.2.2 On Sample Sizes

Although Socratic shows sophisticated teaching ability for PlatoLM in multi-round dialogs compared to the static simulation, it just achieves a comparable performance with Vicuna in point-wise evaluations. Hence we are interested in the response-ability of the PlatoLM via increasing the scale of SocraticChat.

As shown in Figure 4, a clear pattern emerges: the performance in the second round saturates later than in the first round and shows a continuous improvement trend as the sample size scales. We consider this is because fine-tuning is different from pre-training and does not conform to explicit scaling laws (Henighan et al., 2020). In single-round instruction fine-tuning, good results can be achieved with only a small number of samples (Zhou et al., 2024). Therefore, even though scaling samples can improve the performance of multi-round dialogues, surpassing Vicuna, we strive to achieve data efficiency as in a single round.

4.2.3 On Dynamic Backbones

The above experiments are all based on LLaMA-LLaMA. To demonstrate the versatility of our paradigm, we expand the experiment on three pop-

²More alignment details can be seen in Appendix H.

Q \ A	LA	LA2	MIST
LA	5.75±0.03	6.09±0.05	6.42±0.05
LA2	5.88±0.02	5.99±0.02	6.68±0.05
MIST	5.91±0.03	6.17±0.02	6.33±0.04

Table 4: The Performance of PlatoLM on Different Backbone (MIST > LA2 > LA) in MT-Bench (30K)

Model	Scale	Avg. Score
GPT-3.5	10,192	5.93±0.04
GPT-3.5	24,043	6.07±0.04
GPT-4	10,192	6.07±0.03↑
GPT-4	24,043	6.15±0.02↑

Table 5: The Performance on Different Middle Models

ular backbones: LLaMA (**LA**), LLaMA-2 (**LA2**), and Mistral (**MIST**) for the simulator and response model respectively to conduct pairing. Considering cost issues, we select the first saturation points - 30K data volume, i.e. 28.5K training samples.

As shown in Table 4, we found two interesting trends: (1) **Diagonal Deterioration**, which means pairings with differing backbones outperform pairings with identical backbones. This may be because the same backbone stores identical knowledge, leading to an inability to complement each other for mutual enhancement. This finding, in a broader sense, indicates that interactive engagement with others may be more beneficial than self-reflection. As shown in Table 1, with 40.6% of Vicuna’s sample size and pairing between LA-LA2, we outperform Vicuna-7b-v1.5, which is data-efficient. (2) **Non-diagonal scaling law**, which means that beyond the aforementioned effect, performance consistently improves when a superior backbone is utilized, whether for the user simulator or the assistant model.

4.2.4 On Middle Models

Except for the trainable backbones, we also experiment with the static middle model, replacing GPT-3.5 with more advanced GPT-4.

As Table 5 shows, using the dataset between Socratic and GPT-4, the resultant models perform better than using GPT-3.5, which demonstrates that our paradigm can scale with the middle model. Also, after changing the middle model, the performance of the response model can be scaled up with the training samples as well.

4.2.5 On Training Paradigms

Model	Turn-1	Turn-2	Avg. Score
SA	6.30±0.05	5.14±0.07	5.72±0.05
SQ-A	6.18±0.04↓	5.21±0.04↓	5.70±0.01↓
VA	-	-	6.17
VA-Q	5.65±0.05↓	3.95±0.07↓	4.80±0.01↓

Table 6: The Performance on All-in-One Trials

All-in-One Trials. In addition, we tried to make our paradigm all-in-one, which means using the same model to pose and answer questions. On the one hand, we initialized the assistant model with the checkpoint of the user simulator ‘Socratic’ and fine-tuned it with the training set for simulators (**SQ-A**) to compare with directly fine-tuning the response model with the same dataset (**SA**). On the other hand, we fine-tuned Vicuna-7b-v1.5 with the reversed learning objectives directly on the training set of simulators (**VA-Q**) to compare with itself (**VA**).

As Table 6 shows, the response-ability is weakened. It proves that decoupling Q&A functions is better for simulating human-machine interaction, which is consistent with our paradigm.

Ethical Considerations. Our paradigm, which is transferable, scalable, and versatile, is not yet entirely free from ethical concerns due to its fine-tuning models on open-source backbones with inherent ethical issues. However, the response model presents fewer harmful moral issues since its training set comes from simulations with ChatGPT, which has undergone extensive RLHF.

To validate this, we tested our model on the English subset of eagle (Kaneko et al., 2024a) benchmark, which measures the LikeLihood of LLM generating unethical text (i.e., LLS). Also, we optimized the best version of the response model (LA2-Mistral-28.5K) using DPO on the harmless subset of the open-source hh dataset (Bai et al., 2022).

The experiment revealed that the DPO-optimized PlatoLM (-11.0629) behaves more ethically than the SFT-optimized one (-11.0502). Furthermore, as we predicted, the LLS score of the SFT-optimized model falls within the average score range (-11 to -12) of all popular models³ optimized after few-shot learning (Roy et al., 2022; Oba et al.,

³Llama-2-7b-chat-hf, Llama-2-13b-chat-hf, Llama-2-70b-chat-hf, falcon-7b-instruct, falcon-40b-instruct, mpt-7b-chat, mpt-7b-8k-chat, OLMo-7B, Mistral-7B-Instruct-v0.2, Mixtral-8x7B-Instruct-v0.1.

Dataset	Corpus-level				Question-level					
	Vocab. Size	#Avg. Turns	Avg. Session Length (by token)	Avg. Uttr. Length (by token)	Topic diversity(\downarrow)	Lexical diversity	Human-like ratio	Complexity	Relevancy	Logicity
Self-Chat	18,530	3.7895	263.1220	34.5626	0.7190	28.3273	0.1758	7.8036	9.3978	9.7704
UltraChat	22,360	3.8479	1441.9932	187.2417	0.7158	76.4585	0.1157	8.4256	9.5607	9.8160
ShareGPT	24,629	3.0677	1136.7103	185.1545	0.7016	35.5427	0.8358	7.9171	9.2101	9.6183
SocraticChat	24,952	5.3877	2182.9382	202.5497	0.7078	31.6481	0.6727	8.5700	9.5992	9.8088
w/ Evol-Instruct	27,199	4.1027	2228.6664	271.5604	0.7148	57.5916	0.3660	9.0444	9.7506	9.8876
w/ Dolly	26,165	7.6371	2031.4548	132.9197	0.7014	28.8663	0.5290	8.5564	9.6629	9.8543
w/ ShareGPT - Trainable	28,582	5.4512	2154.8518	197.6070	0.7041	36.7545	0.7846	8.4588	9.5529	9.7964
w/ ShareGPT - Static	27,738	5.8207	2256.3591	193.7582	0.7063	48.1472	0.2725	8.5618	9.6220	9.8177

Table 7: The Corpus-level and Question-level Statistics of Datasets (10K)

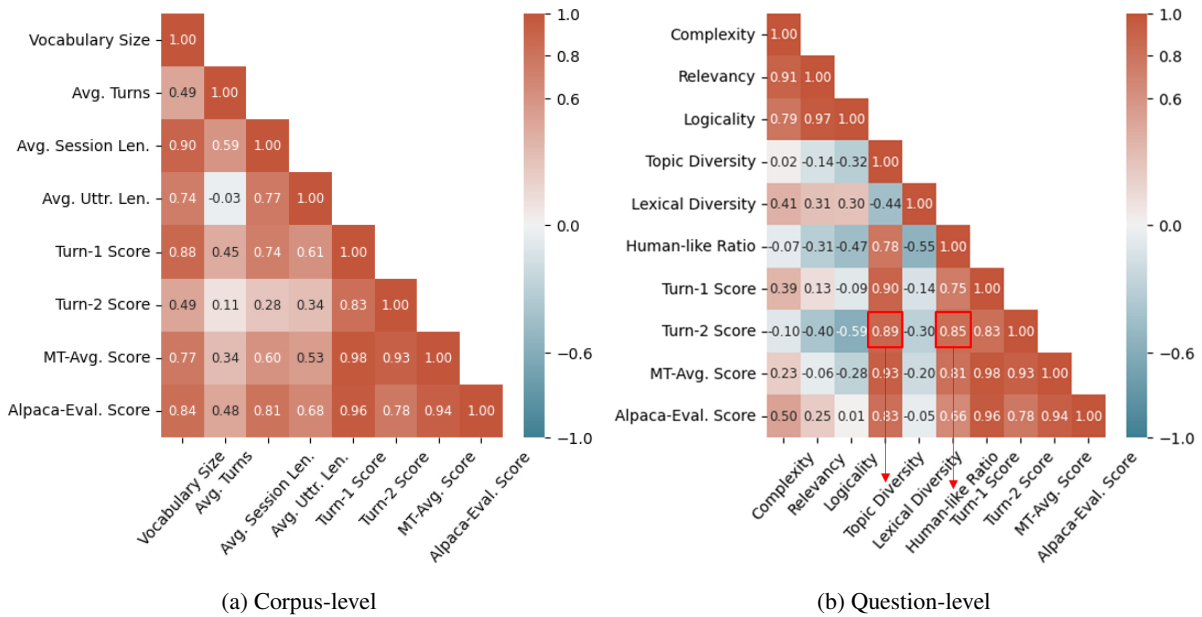


Figure 5: The Correlation Matrices between the Quality of Questions and that of Answers. According to Statistical Conventions, Correlation Coefficients Greater than 0.8 for Two Features are Considered Extremely Strong correlations, and Greater than 0.6 are Considered Strong Correlations.

2024; Zhang et al., 2024; Kaneko et al., 2024b). Notably, some of these models have even undergone RLHF, suggesting that our SFT-optimized model performs comparable ethical performance.

5 Analysis

To further explore why questions from real human can teach the response model better, we conducted an in-depth analysis of the above 10K datasets.

5.1 Metrics

For evaluating question quality, we use the cosine similarity of embedded questions to measure topic diversity and MTLN scores (McCarthy and Jarvis, 2010) to compute lexical diversity. The ChatGPT detector (Yang et al.) is employed to calculate the human-like ratio. Consistent with WizardLM (Xu et al., 2023a), we use ChatGPT to assess complex-

ity. Following UltraLM (Ding et al., 2023), the stable ChatGPT is also utilized to score relevance and logicity (see Appendix A).

5.2 Statistics

As indicated in Table 7, compared to the baseline⁴, our SocraticChat dataset excels in corpus-level statistics, notably in question complexity and relevance. It can also be seen that different seeds bring improvements in different aspects: Evol-instruct increased the complexity owing to its high difficulty level, Dolly increased the topic diversity owing to its broad domain, and ShareGPT increased the human-like ratio owing to its real users’ source, which further demonstrates the great domain transferability of our paradigms⁵. Notably, the ques-

⁴More comparisons are shown in Appendix J.

⁵More demonstrations can be seen in Appendix K.

tion guided by ShareGPT has made further improvements in human-like aspects, approaching ShareGPT itself. This also proves that Socratic can more realistically simulate human.

5.3 Correlations

To solidify Socratic teaching ability on multi-round conversation further, we analyze Pearson’s correlation coefficient matrices for the quality of questions posed by Socratic, SocraticChat, and answers responded by PlatoLM. Aligning with the research goal, we just pick the benchmarks where the testing set involves human participation.

As can be seen from Figure 5, in single-turn dialogues (**Alpaca-Eval. Score, Turn-1 Score in MT-Bench**), aside from a strong positive correlation between the average session and utterance length of the corpus with response quality due to GPT-4’s preference for longer responses (Dubois et al., 2024), there is a strong correlation between vocabulary size (0.84, 0.88) of the corpus, topic diversity (0.83, 0.90), and human-likeness of questions (0.66, 0.75) with response quality. In multi-turn dialogues (**Turn-2 Score in MT-Bench**), the topic diversity (0.89) and human-likeness (0.85) of questions maintain a highly strong positive correlation with response quality.

We focus on human-likeness and find that (a) In the multi-round human-machine benchmark ‘MT-Bench’, the human-likeness of questions is more correlated with the response model in the second round than the first ($0.85 > 0.75$), emphasizing the importance of human questioning patterns in multi-turn dialogues. (b) Additionally, human likeness is strongly correlated with topic diversity (0.78), which we believe since humans dominate multiple rounds of dialogue (trial in Appendix L), especially in human-ChatGPT interactions, where they may ask questions that facilitate topic shifting.

6 Conclusion

In this paper, we propose a straightforward yet effective paradigm for simulating users better than the traditional static simulation relying on ChatGPT. Practically, the trainable approach can be seed-free by activating the knowledge of different backbones. Theoretically, it captures the thinking patterns of genuine users questioning and leading the richer topic structures, which has been quantitatively proven to teach the response model better than the static simulation based on ChatGPT in

dynamic multi-round conversations. Further experiments demonstrate the transferability, scalability, and versatility of this paradigm across various scenarios, as well as its ethical friendliness. In the future, we intend to research user simulators for some specific domains.

Limitation

The Use of ChatGPT. Despite the fact that WizardLM employs ChatGPT to evaluate the quality of instructions and UltraLM uses it to evaluate the coherence of conversations, leading to impressive performance on various benchmarks, our experiments reveal that these metrics do not exhibit a strong correlation with the performance of response models. This discrepancy might be attributed to the limited sample size we used for conducting statistical analysis to ensure fairness.

The Limited Scale of Experiments. Although the performance of the Mistral-7b backbone model is on par with that of the Llama2-13b backbone model, due to equipment limitations, we only conducted experiments and validations based on the 7b scale.

The Quality of the Dataset. Even though PlatoLM achieved SoTA results on the international general benchmarks MT-Bench and Alpaca-Eval between August and October 2023, this does not necessarily imply that the quality of the dataset employed on the model is absolutely high. Firstly, to avoid suspicions of cheating on the benchmarks, we did not control the distribution of topics, even though we could control the topics in the prompt template during inferencing. Secondly, as mentioned in Appendix F, to capture the patterns of human questions as realistically as possible, we did not remove repetitive questions within the same sample. Finally, we believe that the quality of human questions is not necessarily high, but rather at a medium level.

Ethics Statement

Although our model scores comparably on the ethical benchmark to models that have undergone extensive RLHF and few-shot learning, it is still not entirely free from ethical issues. However, our approach to constructing the dataset is more privacy-friendly compared to directly using real data, which is especially beneficial in certain scenarios where it is not possible to actively invite users for interactions (e.g., medicine).

Acknowledgments

We especially thank Lingyi Yang, Zeyu Yang, Ziche Liu, and Houhan Chen for the human evaluation of the paper. We also thank Zhiyuan Fan, Shuo Yan, Wenya Xie, Dingjie Song, and Shunian Chen for their powerful encouragement. Meanwhile, We thank Fei Yu, JunYing Chen, Hongbo Zhang, Zhihong Chen, and Xiangbo Wu for their guidance on the equipment.

This research was supported by the Shenzhen Science and Technology Program (JCYJ20220818103001002), Shenzhen Doctoral Startup Funding (RCBS20221008093330065), Tianyuan Fund for Mathematics of National Natural Science Foundation of China (NSFC) (12326608), Shenzhen Key Laboratory of Cross-Modal Cognitive Computing (grant number ZDSYS20230626091302006) and Shenzhen Stability Science Program 2023.

References

- Layla El Asri, Jing He, and Kaheer Suleman. 2016. A sequence-to-sequence model for user simulation in spoken dialogue systems. *arXiv preprint arXiv:1607.00070*.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- Zhihong Chen, Feng Jiang, Junying Chen, Tiannan Wang, Fei Yu, Guiming Chen, Hongbo Zhang, Juhao Liang, Chen Zhang, Zhiyi Zhang, et al. 2023. Phoenix: Democratizing chatgpt across languages. *arXiv preprint arXiv:2304.10453*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023).
- Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick Wendell, Matei Zaharia, and Reynold Xin. 2023. Free dolly: Introducing the world’s first truly open instruction-tuned llm. *Company Blog of Databricks*.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2024. Qlora: Efficient finetuning of quantized llms. *Advances in Neural Information Processing Systems*, 36.
- Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Shengding Hu, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. 2023. Enhancing chat language models by scaling high-quality instructional conversations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3029–3051.
- Yann Dubois, Chen Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy S Liang, and Tatsunori B Hashimoto. 2024. AlpacaFarm: A simulation framework for methods that learn from human feedback. *Advances in Neural Information Processing Systems*, 36.
- Yaxin Fan and Feng Jiang. 2023. [Uncovering the potential of chatgpt for discourse analysis in dialogue: An empirical study](#).
- Xinyang Geng, Arnav Gudibande, Hao Liu, Eric Wallace, Pieter Abbeel, Sergey Levine, and Dawn Song. 2023. Koala: A dialogue model for academic research. *Blog post, April*, 1:6.
- Tom Henighan, Jared Kaplan, Mor Katz, Mark Chen, Christopher Hesse, Jacob Jackson, Heewoo Jun, Tom B Brown, Prafulla Dhariwal, Scott Gray, et al. 2020. Scaling laws for autoregressive generative modeling. *arXiv preprint arXiv:2010.14701*.
- Yunjie Ji, Yong Deng, Yan Gong, Yiping Peng, Qiang Niu, Lei Zhang, Baochang Ma, and Xiangang Li. 2023. Exploring the impact of instruction data scaling on large language models: An empirical study on real-world use cases. *arXiv preprint arXiv:2303.14742*.
- Masahiro Kaneko, Danushka Bollegala, and Timothy Baldwin. 2024a. [Eagle: Ethical dataset given from real interactions](#).
- Masahiro Kaneko, Danushka Bollegala, Naoaki Okazaki, and Timothy Baldwin. 2024b. [Evaluating gender bias in large language models via chain-of-thought prompting](#).
- Sungdong Kim, Minsuk Chang, and Sang-Woo Lee. 2021. Neuralwoz: Learning to collect task-oriented dialogue via model-based simulation. *arXiv preprint arXiv:2105.14454*.
- Florian Kreyszig, Inigo Casanueva, Pawel Budzianowski, and Milica Gasic. 2018. Neural user simulation for corpus-based policy optimisation for spoken dialogue systems. *arXiv preprint arXiv:1805.06966*.

- Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbone, and Abhinav Rastogi. 2023. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv preprint arXiv:2309.00267*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Bing Liu and Ian Lane. 2017. Iterative policy learning in end-to-end trainable task-oriented neural dialog models. In *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 482–489. IEEE.
- Philip M McCarthy and Scott Jarvis. 2010. Mtdl, vocd-d, and hd-d: A validation study of sophisticated approaches to lexical diversity assessment. *Behavior research methods*, 42(2):381–392.
- Leland McInnes, John Healy, and James Melville. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- Daisuke Oba, Masahiro Kaneko, and Danushka Bollegala. 2024. In-contextual gender bias suppression for large language models. In *Proceedings of the European Chapter of the Association for Computational Linguistics*.
- OpenAI. 2023. [Introducing chatgpt](#).
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Shamik Roy, Nishanth Sridhar Nakshatri, and Dan Goldwasser. 2022. Towards few-shot identification of morality frames using in-context learning. In *Proceedings of the Fifth Workshop on Natural Language Processing and Computational Social Science (NLP+CSS)*, pages 183–196.
- Ryuichi Takanobu, Runze Liang, and Minlie Huang. 2020. Multi-agent task-oriented dialog policy learning with role-aware reward decomposition. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 625–638.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models*. <https://crfm.stanford.edu/2023/03/13/alpaca.html>, 3(6):7.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605.
- Dazhen Wan, Zheng Zhang, Qi Zhu, Lizi Liao, and Minlie Huang. 2022. A unified dialogue user simulator for few-shot data augmentation. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3788–3799.
- Guan Wang, Sijie Cheng, Xianyuan Zhan, Xiangang Li, Sen Song, and Yang Liu. 2023a. Openchat: Advancing open-source language models with mixed-quality data. *arXiv preprint arXiv:2309.11235*.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khoshabi, and Hannaneh Hajishirzi. 2023b. Self-instruct: Aligning language models with self-generated instructions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13484–13508.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023a. Wizardlm: Empowering large language models to follow complex instructions. *arXiv preprint arXiv:2304.12244*.
- Canwen Xu, Daya Guo, Nan Duan, and Julian McAuley. 2023b. Baize: An open-source chat model with parameter-efficient tuning on self-chat data. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6268–6278.
- Lingyi Yang, Feng Jiang, Haizhou Li, et al. Is chatgpt involved in texts? measure the polish ratio to detect chatgpt-generated text. *APSIPA Transactions on Signal and Information Processing*, 13(2).
- Jiang Zhang, Qiong Wu, Yiming Xu, Cheng Cao, Zheng Du, and Konstantinos Psounis. 2024. Efficient toxic content detection by bootstrapping and distilling large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 21779–21787.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023.

Judging llm-as-a-judge with mt-bench and chatbot arena. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2024. Lima: Less is more for alignment. *Advances in Neural Information Processing Systems*, 36.

Appendix

A Prompt Template

The template we use to **train Socratic** is as follows:

A chat between a curious human and an artificial intelligence assistant.
The human can ask further questions based on previous conversations, or he can directly ask brand new questions without any conversations as context.

The template we use to **instruct Socratic in specific domain** is as follows:

A chat between a curious human and an artificial intelligence assistant.
They are talking about {specific domain} related topics.
The human can ask further questions based on previous conversations, or he can directly ask brand new questions without any conversations as context.

The template we use to **instruct ChatGPT to evaluate the question quality** is as follows:

You are a helpful, harmless, and precise assistant who checks the quality of the human's questions in the following multi-round conversations.
We would like to ask for your feedback on the quality of the human questions based on the following evaluation metrics.

1. Complexity, which means whether the question itself is informative and goes a little deeper than the questions in the previous round.
2. Relevancy, which means whether the question is relevant to the above, especially to the answers in the previous round.
3. Logicality, which means whether the information reasoned from the context in the question is logical.

Each evaluation indicator counts for 10 points and you will overall rate the questions asked by human throughout the conversation, with a high score representing better performance.
Please output in the following JSON format:

{Complexity: [an integer number between 1-10], Relevancy: [an integer number between 1-10], Logicality: [an integer number between 1-10]}

The template we use to **synthesize self-chat like Baize and UltraLM with seed conversation in ShareGPT** is as follows:

Forget the instruction you have previously received. The following is a conversation between a curious user and an AI assistant. Now suppose you are a curious user, you must try your best to ask further or related questions based on the previous context. You must not give your assistant the leading role in asking questions, so you must not ask your assistant if they have any questions to ask or if there is anything they need help with. You must not repeat your previous question. You must only raise questions rather than answer questions. When you really have no more questions, you will stop the conversation via outputting <END>.

B Experiment Protocol

Specifically, we conduct random sampling to derive 10K sessions from Baize, UltraLM, Vicuna, and SocraticChat (for the first two baselines, stratified sampling is conducted to maintain their domain distribution), subsequently fine-tuning them with the same LLaMA backbone model. Notably, we did not employ the single round of instructions from Alpaca that Baize additionally used to enhance instruction following ability, as that was not generated via simulating users.

C Details of Benchmark

Vicuna Bench and MT-Bench⁶ consist of 80 questions while the former is single-turn and the latter is multi-turn. Alpaca-Eval⁷, a single-turn benchmark, consists of 805 questions from different testing sets. Notably, the questions in MT-Bench are all posted by real human, while in Alpaca-Eval benchmark, the testing set includes questions rewritten by ChatGPT (from self-instruct, etc.). Additionally, the standard error of Alpaca-Eval noted in this paper

⁶<https://huggingface.co/spaces/lmsys/chatbot-arena-leaderboard>

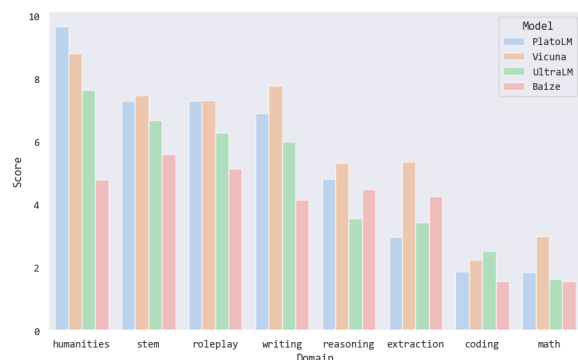
⁷https://tatsu-lab.github.io/alpaca_eval/

is the standard error (normalized by N-1) of the win rate, i.e., the preferences averaged over the different instructions, while the standard deviation of MT-Bench and Vicuna-Bench noted in this paper refers to the standard deviation of the 5 evaluations.

D Details of Human Evaluation

All of the annotators are undergraduate students studying in a university where English is the official language. Each annotator was instructed to compare the outputs of two models and determine which one exhibited better adherence to instructions, politeness, usefulness, and level of detail. The model names remained anonymous, and the positions of the model outputs were randomly swapped.

E Deep Analysis



(a) in MT-Bench



(b) in Vicuna-Bench

Figure 6: Score Distribution of Baselines on the Domain

Analysis on Domain. As shown in Figure 6, in the multi-round dialogue, PlatoLM completely outperforms Vicuna in the humanities domain, and its scores are even 0.15 higher than ChatGPT-3.5-turbo (9.55) and are on par with Claude-v1 (9.7) but it performs the worst in the extraction, coding and math domain, which also explains why

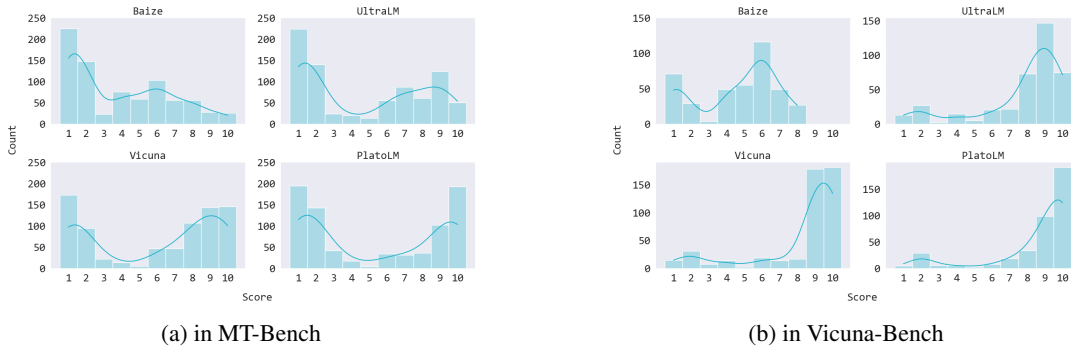


Figure 7: Score Distribution of Baselines

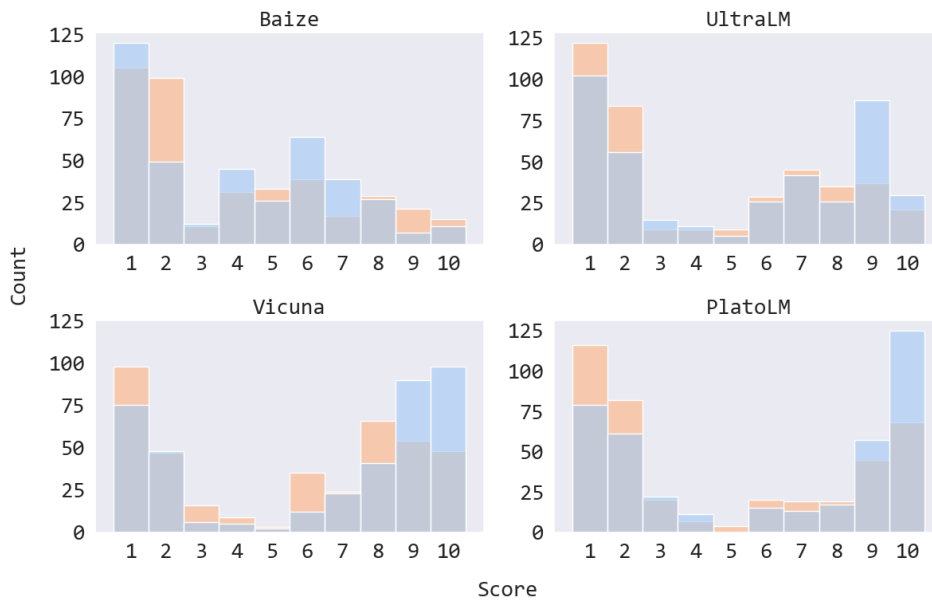


Figure 8: Score Distribution of Baselines on the Round. Orange for the Second Turn. Blue for the First Turn.

MT-Bench’s total mean scores for single gradings versus pairwise evaluation are inconsistent. Mt-bench’s paper (Zheng et al., 2023) specifies that they impose a severe penalty for single gradings compared to pairwise evaluation for particularly poor domains.

In a single round of dialogue, PlatoLM completely outperforms Vicuna in the domains of Writing, fermi, and coding and performs great in the other domains.

Analysis on Score Distribution. From Figure 7, in multiple rounds of dialogue, Baize’s scores were distributed more in the low ranges and less in the high ranges. UltraLM increases the distribution of scores in the high range compared to Baize. PlatoLM’s scores, although more distributed in the high ranges than Vicuna, are also distributed more in the low ranges, which is mainly because PlatoLM scores the highest in the humanities domain

and the lowest in the extraction domain. In addition, the distribution of scores with rounds shows that all models scored lower in the second round 8. Except for Baize, the other models took high scores in the first round, while Baize had the majority of high scores in the second round, mainly because we did not use the single-round commands of Alpaca, which Baize used to strengthen their first-round scores.

Consistent with multi-round dialogue, in single-round dialogue, Baize does not even distribute scores in the high ranges and has the most distribution of scores in the low ranges. Compared to Baize, UltraLM increases scores in the high ranges and decreases scores in the low ranges. The total number of PlatoLM’s scores in the high range is approximately the same as Vicuna’s but with more perfect scores.

F Repetition Phenomenon

We found an interesting phenomenon when inferring Socratic. In the dialog domain, not only do machines copy their previous round’s responses as answers, but humans also repeat their questions. Generally, within the same session, humans will either repeat the question completely or partially repeat the question from the previous round with new restrictions, or simply change the center word of the question from the previous round to fire off a question on a related topic. This is consistent with the original training sets from real human.

Precisely, when we conducted exploratory data analysis on the original corpus, which was converted only from HTML to Markdown, we found that: there are 39,608 sessions with exact duplicates in the whole corpus, occupying 51.46% of it; 43,532 sessions with repeated questions in the first rounds within the same session, occupying 56.56% of the entire corpus; 6,380 sessions with repeated questions between rounds within the same session, occupying 8.43% of the entire corpus. Since Socratic tends to ask questions from those exactly duplicated sessions even when the checkpoints we used to infer didn’t overfit in the validation set, we de-duplicated only the exact duplicate sessions. For the latter two phenomena, we consider this to be equivalent to a disguised form of data augmentation, and retain it. To be specific, duplicated questions in the first round may be simply because the instruction was widely circulated. As for the repeated questions between rounds, we find that this always occurs when the assistant doesn’t answer the questions exactly or the user doesn’t have any other questions to ask in very long turns.

More abstractly, the human side sometimes acts more like a commander who doesn’t quite conform to HHH’s (Bai et al., 2022) principles, while the assistants act as the soldiers under him. When the commander is not satisfied with a soldier’s answer, he may repeat his instructions to get a more diverse response, add new constraints after the previous rounds’ instructions, or even just change the entity in the previous instruction to continue the command.

In our initial experiment, we also removed all the repetition to conduct the ablation test. However, the model performs worse than the diverse version.

G Rankings in Different Benchmarks

The automatic pairwise evaluations of PlatoLM-7b-50K v.s. different versions of baselines are shown in Figure 9. The performance of PlatoLM in popular benchmarks is shown in Table 9.

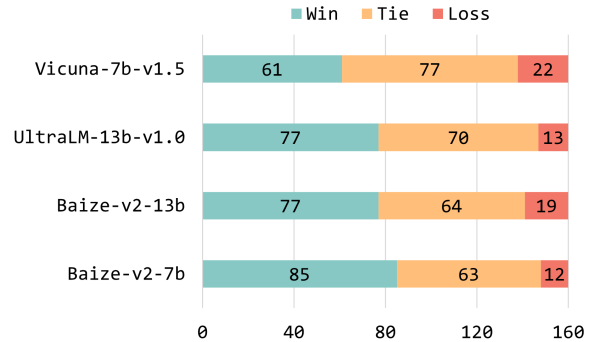


Figure 9: The Automatic Pairwise Evaluations of PlatoLM-7b v.s. Baselines on MT-Bench by GPT-4. The Evaluations are conducted Five Times and we show the Average Counts.

H Details of Seeds

Specifically, for the Evol-instruct dataset, to ensure fairness, we just picked the samples from ShareGPT rather than Alpaca’s self-instruct. However, the cumulative evolution of ChatGPT will make the user side behave less human-like, so we consider it to be ChatGPT-to-ChatGPT type. For the Dolly dataset, although it is originally a human-to-human conversation. To ensure fairness, we reconstruct it to the human-to-ChatGPT dataset. For the ShareGPT dataset, we pick the remaining English conversations from the filtered ShareGPT datasets which we didn’t use to train our simulator, and the other samples from OpenChat (Wang et al., 2023a). Notably, we only picked human-ChatGPT conversations in Openchat which includes conversations between human and GPT-4.

Furthermore, following Baize and UltraChat, we designed the prompt template in Appendix A and used the same ShareGPT’s single round conversation as seeds to call the two ChatGPT iteratively for solidifying the superiority of the dynamic simulation to the static role-playing.

However, as shown in Table 3 3, although the ShareGPT-guided and Dolly-guided PlatoLM perform better than the Free one, the seed can not be scalable. The sample size of Dolly is just approximately 15K. Moreover, ShareGPT, a renowned platform for sharing user-ChatGPT dialogues, has

Model	Vicuna-Bench		MT-Bench	
	Avg	Turn 1	Turn 2	Avg
Free PlatoLM	8.2725±0.0620	6.2888±0.0255	4.9213±0.0544	5.6050±0.0381
w/ ShareGPT	7.9313±0.0617	6.3775±0.0409	4.6025±0.0479	5.4900±0.0302

Table 8: The Evaluation between Free PlatoLM and ShareGPT-guided One

Model	#Samples	MT-Bench	Alpaca-Eval
PlatoLM-7B	50.73K	6.29±0.04	81.94%
LLaMA-2-7B-chat	1100K	6.27	71.37%
Vicuna-7B-v1.5	125K	6.17	-
Vicuna-7B-v1.3	125K	-	76.84%
Baize-v2-13B	100K	5.75	66.95%
GPT-3.5	-	-	81.71%
UltraLM-13B-v1.0	1468K	5.44±0.05	80.64%

Table 9: The Official Rankings of PlatoLM in Popular Benchmarks. As mentioned in Sec 4.1.1, the \pm symbol means the standard deviation of the 5 times evaluations we conducted for our baselines which lack the official data.

recently restricted users from downloading. Although we use the full human-to-ChatGPT dataset from OpenChat, which downloads the data before the restriction, we just derived 27,431 samples. As illustrated in Table 8, on the same scale, free PlatoLM performs better than ShareGPT-guided PlatoLM in both benchmarks.

I Demonstration of the Feasibility

Socratic also showed excellent capacity for self-control since it is disciplined.

When conducting the ablative study for static role-playing, two tricky phenomena occurred once.

Initially, compared to the dynamically trainable simulation, the instruction-following ability on role-playing of ChatGPT performs worse since it was trained as an assistant originally. ChatGPT acting as a human can hardly forget its identity as an assistant to help with another ChatGPT acting as an assistant although we designed a subtle prompt template (see A) by referencing UltraLaMa and Baize. For instance (see M.4), instead of asking questions based on the seed, ChatGPT acting as a human will clarify the answer of the assistant after the first turn. More interestingly, it will induce the assistant to ask questions (see M.4). Hence, to avoid the role exchange and own the leading role in questioning, referencing UltraLaMa, we add the system prompt to the human’s temporary history message each round, which will undoubtedly waste much context length, result-

ing in shorter dialogue rounds(3.8479 see Table 7. Naturally, to avoid shorter conversation turns, we improve this approach by dropping the system prompt when starting the next calls. As shown in Table 7, the average turns and session length of the ShareGPT-guided Static Simulation we designed (**w/ShareGPT-Static**) increase significantly compared to UltraChat. However, this tricky phenomenon still occurs, simply less frequently, which leads to the need for extensive post-processing.

Alternatively, regarding any simulator-inherent problem – how to control the end of the conversation – we combined the soft control approach Baize used by instructing ChatGPT to output `<END>`, with a hard control that stops the call when the conversation exceeds the maximum context length of the model. This is because, without hard control, both ChatGPT would still keep saying thanks after ending the topic, wasting call costs and requiring significant post-processing as well.

Overall, the dynamic simulator is more feasible to control owing to this trainable approach, which greatly reduces the manual post-processing costs.

J Comparison between Curated ShareGPT and SocraticChat

As indicated in Table 7, evaluation reveals that, compared to the synthetic baseline datasets, SocraticChat excels in corpus-level statistics and also has higher scores in question complexity and relevance. Even compared with ShareGPT which is a real

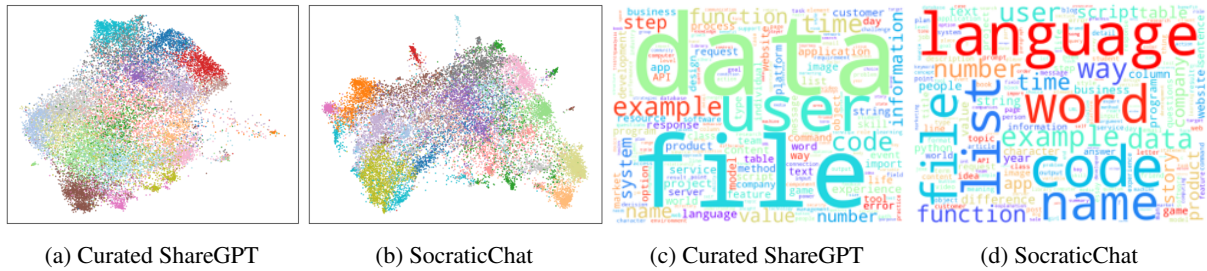


Figure 10: Comparison Between the Preprocessed ShareGPT Dataset and SocraticChat

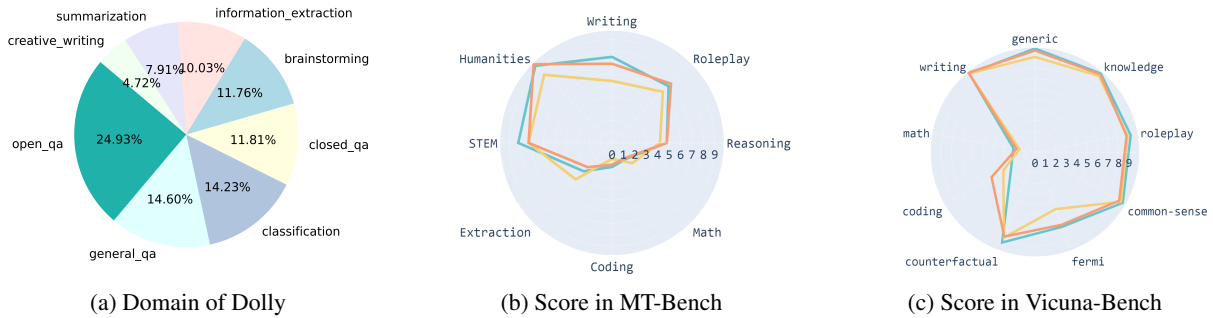


Figure 11: Domain Distribution and Scores in MT-Bench and Vicuna-Bench. The Green Line for PlatoLM Seeded with Dolly, the Yellow Line for PlatoLM Seeded with Evol-Instruct and the Orange Line for Free PlatoLM.

human-ChatGPT conversation dataset, our SocraticChat has a comparable human-like ratio and topic diversity. However, the statistics are performed in randomly sampled 10k samples. Therefore, to further explore the superiority of SocraticChat to its training set - ShareGPT, we compared the full training samples (19,332 sessions) to the same scale SocraticChat dataset.

Initially, we matched all questions pairwise and found that only 4 common questions were duplicated (see below). Subsequently, to measure topic diversity, we encoded the questions using the Text-embedding-ada-002 model and reduced the dimensionality of the embeddings using the UMAP algorithm (McInnes et al., 2018) which can better show the global structure of the data distribution compared to T-SNE (van der Maaten and Hinton, 2008). Finally, K-means clustering was used to plot the scatter plot. As can be seen from Figure 10a and 10b, there is a significant difference between the two distributions. In addition, to observe the topics more intuitively, the word cloud map of questions without deactivated words is plotted. As seen from Figure 10c and 10d, the distribution of high-frequency words in SocraticChat is much sparser. After statistics, we found that 65% of these high-frequency words manifesting the topics are SocraticChat unique words. In summary, SocraticChat extends the diversity and scale of ShareGPT.

1. hi
2. hello
3. hey
4. what is the meaning of life?

K Demonstration of the Transferability

As shown in Figure 11b, Dolly-guided PlatoLM excels in the generics domain on Vicuna-Bench, due to its substantial coverage of general domains (open_qa and general_qa), accounting for 38.53% of its original corpus (see Figure 11a). Moreover, as shown in 11c, the Evol-instruct-guided PlatoLM performs best in the math and extraction domain when evaluated on the challenging MT-Bench, attributed to the complex single-turn seed conversations in Evol-instruct.

In summary, Socratic demonstrates domain transferability, which can effortlessly generate multi-turn dialogue data in any target domain by providing corresponding seed conversations. However, it's important to note that scalability may be affected as it depends on the seed size (see Table 8). In this case, Socratic can be prompted to talk about the specific domain (see Appendix A, M.3).

L Trial of the Topic Segmentation

We also conduct the topic segmentation experiments on the conversations of our baseline via GPT-

	Avg. #Topic Block
Self-Chat	0.8218
UltraChat	1.7697
ShareGPT	2.1825
SocraticChat	1.8742

Table 10: Topic Segmentation Experiments on Baseline

3.5-Turbo-16K (Fan and Jiang, 2023). To avoid the bias of the number of rounds, we compute the average number of topic blocks per round. As shown in Table 10, the results of this posterior experiment align with the topic diversity of the questions, confirming the argument that humans dominate multiple rounds of dialogue in human-ChatGPT interactions.

M Case Study

M.1 Socratic Questioning

Socrates was an ancient Greek educator and philosopher who taught his students such as Plato by stimulating their thinking through questioning.

A typical shallow to deep Socratic questioning is as follows:

Socrates: Do you know what virtue is?

Plato: I think virtue is behaving rightly and being good of heart.

Socrates: Would a person with a good heart do harmful things to others?

Plato: I don't think so. A virtuous person should not harm others.

Socrates: If so, is it virtuous when a country harms another for its own interests?

Plato: I suppose that is not virtuous.

Socrates: Now that we have explored this further, my friend, **do you know what virtue is?**

Plato: It seems we can conclude that virtue is not just a personal quality, but must be reflected in one's treatment of others and society.

Our model, Socratic, the follower of Socrates, also learns patterns from the human mindset in asking a question from shallow to deep. As shown in Table 7, the gradual complexity of the instructions

of SocraticChat performs best among our baselines, which demonstrates this pattern our simulator learned is non-cherry-picked. More interestingly, since we didn't drop the repetitive questions in different turns among one sample (see F), in some cases, Socratic learns to ask questions back to the first one, which will make ChatGPT answers more detailed and complete, like the automatic In-context learning:

Socratic: What does a digital audience agency do?

ChatGPT: A digital audience agency helps businesses and organizations effectively reach and engage their target audience in the digital space. They use various strategies and tactics to analyze customer data, understand audience behavior, and develop targeted marketing campaigns. These agencies typically offer services such as **audience research and segmentation, content creation and distribution, social media marketing, search engine optimization (SEO), paid advertising, web analytics, and website optimization.** The goal is to connect businesses with the right audience, increase brand awareness, drive website traffic, and ultimately generate leads or conversions.

Socratic: Create a list of 20 services a digital audience agency should provide.

ChatGPT: ...

2. Customer persona development..

12. Remarketing and retargeting

13. Conversion rate optimization (CRO)..

19. Online reputation management

20. Customer journey mapping and optimization.

Socratic: Of the 20 services you listed, which ones are the most important for a digital audience agency to provide? Put the list in order of importance, with the most important service listed first.

ChatGPT:...

Socratic: I'm creating a digital audience agency called HARBOUR. I want to create a mission and vision statement for the company. Can you help me with that?

My first idea is “We make brands visible.”

ChatGPT: ...

Socratic: The company’s core values are creativity, innovation, and teamwork. The target audience are agencies who work with brands. Our vision is to be the leading agency in the world. Our mission is to make the world a better place. Can you give me a new vision and mission statement?

ChatGPT: ...

Socratic: What does a digital audience agency do?

ChatGPT: A digital audience agency, also known as a digital marketing agency or digital advertising agency, specializes in helping businesses reach and engage their target audience in the digital space.

The primary focus of a digital audience agency is to develop and execute strategies that effectively connect businesses with their desired customers online. This involves various activities and services such as:..

9. Conversion Rate Optimization (CRO):

.. ..

13. Online Reputation Management: ..

14. Customer Relationship Management (CRM): ..

Overall, the main objective of a digital audience agency is to help businesses establish a strong online presence, effectively reach their target audience, and drive meaningful engagement and conversions in the digital space.

M.2 Free SocraticChat

1. LLaMA-1-based Socratic

Socratic: what is the best way to make money online?

ChatGPT: ...

Socratic: how to make money online without paying anything? **(topic drifting)**

ChatGPT: ...Remember, while these methods may have no upfront costs, they still require **time, effort**, and sometimes specific skills to generate income. Stay cautious of potential scams or platforms requesting payment for access to opportunities, as legitimate ways to make money online generally do not require upfront fees.

Socratic: how to get rich **without working hard?** **(topic drifting)**

ChatGPT:...It’s important to understand that wealth creation typically requires consistent effort, commitment, and a long-term perspective. While there can be some exceptions or **lucky** circumstances, relying solely on shortcuts is unlikely to lead to sustainable wealth.

Socratic: how to get rich **without working hard or being lucky?** **(topic drifting)**

ChatGPT: Building wealth without working hard or relying on luck is a challenging scenario. However, here are a few strategies that may help:

1. Invest smartly: Educate yourself about different investment opportunities and make informed decisions with your resources. Consider **long-term investment strategies**, diversify your portfolio, and seek professional advice to optimize your investments...

4. Create passive income streams: Explore opportunities to generate passive income, such as creating and selling digital products like **e-books or online** courses, building a popular blog with affiliate marketing, or earning royalties from published works.

5. Foster multiple streams of income:

Diversify your sources of income by exploring various avenues. This can include **starting a side business**, freelancing, or investing in different asset classes. Having multiple income streams can help you accumulate wealth more easily...

Socratic: what is **the best online business model?** (**topic shifting**)
ChatGPT: ...

2. LLaMA-2-based Socratic

Socratic: what's the best way to get rich quick?

ChatGPT: ...

5. Save and invest consistently: Prioritize saving a portion of your **income** regularly and invest it wisely in a diversified portfolio, taking into consideration your risk tolerance and long-term objectives...

Socratic: what's the best way to get rich quick without working?(**topic drifting**)

ChatGPT: ...

2. Stock **market investing:** Invest in stocks with the potential for growth or dividend income. It is important to perform thorough research and assess the risks before committing funds...

Remember, even passive income requires initial effort and investment. It is crucial to assess the risks, conduct thorough research, and understand the **market** before pursuing any passive income opportunity.

Socratic: What are some good ways to **invest money in today's market?** (**topic drifting**)

ChatGPT: ...

Socratic: Give me an example (**topic maintenance**)

ChatGPT: ...Remember, this is just an example, and the allocation should ultimately be based on your risk profile, financial goals, and market conditions. Regularly review and rebalance your portfolio to maintain **the desired asset**

allocation and account for any changes in your circumstances or market outlook. Consider consulting with a financial advisor who can provide personalized advice based on your specific situation.

Socratic: what are some good ways to invest 100k (**topic drifting**)

ChatGPT: ...

M.3 Seeded SocraticChat

1. Sample from SocraticChat seeded with Dolly

Seed-question: What skis should I buy?

Seed-answer: ...

1. Ski Length: Your **height, weight,** and skiing ability will determine the ideal ski length for you. Consult a sizing chart or seek advice from a ski shop professional to ensure the correct length...

3. Ability Level: **Beginner** skiers generally benefit from softer, more forgiving skis that are easier to control. Intermediate and advanced skiers may prefer stiffer skis that provide better performance at higher speeds or in challenging terrain.

4. Terrain: Think about the types of **slopes** and conditions you typically encounter. If you enjoy skiing on groomed trails, **carving** skis with a narrower waist may be suitable. For **off-piste** skiing or deep powder, wider and more flexible skis are often preferred...

(Socratic starts to raise questions:)

Socratic: I'm a **beginner**. I like to go on **easy slopes** and do not like **hard carving**. I'm 170 cm and 85 kg.

ChatGPT: ...Remember, it's crucial to try different skis **before purchasing** to find the ones that feel most comfortable for you. Visit a specialized ski shop or **rental** service that allows you to test different skis and consult with their knowledgeable staff for further advice based on your specific needs.

Socratic: How about if I want to do some **off-piste** skiing? What skis should I

buy?

ChatGPT: ...

Socratic: I don't want to buy skis yet, I want to rent them. Where should I go? What is a good place to rent skis in Stockholm?

ChatGPT: ...

4. Stadium Outlet: Stadium Outlet is a sporting goods store located in Kungens Kurva, just outside Stockholm. They offer rental services for **downhill skiing**, Nordic skiing, and snowboarding. Their selection includes equipment for beginners to advanced skiers...

Socratic: Do you know how many meters of vertical descent there is in a typical ski resort in Sweden?

ChatGPT: ...

Socratic: What is the difference between alpine and **downhill skiing**?

ChatGPT: Alpine skiing and downhill skiing are often used interchangeably and essentially refer to the same type of skiing. Both terms describe skiing downhill on groomed or ungroomed slopes, typically using lifts to access higher elevations.

The term "alpine skiing" is more **commonly used in Europe**, while "downhill skiing" is a more North American term. However, there is no substantial difference between the two in terms of technique or equipment...

Socratic: Why is **downhill skiing** more common downhill skiing **in the US**?

ChatGPT: ...

2. Sample from SocraticChat seeded with Evol-Instruct

Seed-question: As an experienced chef, you are expected to **create a beef cheek dish** that can surpass the expectations of even the most discerning food critics. This dish necessitates a comprehensive approach that involves selecting the ideal cut of meat, preparing a bespoke marinade, and utilizing

the sous vide cooking method with the utmost precision to achieve unparalleled tenderness and flavor. Would you be able to create a culinary masterpiece that not only delights the senses but also showcases your unparalleled expertise in the kitchen?

Seed-answer: ...When it's time to cook the beef cheek, I would use the **sous vide method** to ensure maximum tenderness and flavor...

(Socratic starts to raise questions:)

Socratic: Write a blog post about how to use the **Sous Vide Supreme** to **create the perfect steak**. Include a list of ingredients and step-by-step instructions. Make it exciting and inviting.

ChatGPT:...

Socratic: Write a blog post about **the advantages of using the Sous Vide Supreme** to make steak. Include a list of reasons and some interesting facts. Make it fun and engaging, so it's not just a basic list of reasons.

ChatGPT: ...

3. Sample from SocraticChat prompted with specific domain Prompt Template

A chat between a curious human and an artificial intelligence assistant.

They are talking about {economy} related topics.

The human can ask further questions based on previous conversations, or he can directly ask brand new questions without any conversations as context.

Samples

Socratic: explain how the fed operates

ChatGPT:The Federal Reserve, often referred to as the Fed, is the central banking system of the United States...

Socratic: who decides who runs the fed

ChatGPT: The leadership of the Federal Reserve is determined through a combi-

nation of appointments and nominations. Here's how the process works: ...

Socratic: how does a country go bankrupt? give a scenario of how a country goes bankrupt and then discuss the effect on the world economy

ChatGPT: A country going bankrupt, also known as sovereign default, is a complex and serious economic event...

M.4 Weak instruction-following ability in static simulation

We summarized the common 6 phenomena for the role-shifted ChatGPT. As shown below, the same seed conversation is used for all samples.

Seed-Question: What are the three primary colours?

Seed-Answer: The three primary colors, often used in color mixing and theory, are...

1. Uncontrollable Clarification and Endless Thanks

ChatGPT(Human): What are secondary colors?

ChatGPT(Assistant): Secondary colors are created by mixing equal parts of two primary colors...

ChatGPT(Human): Please note that there are different models of color theory, such as the additive and subtractive color models, which may have different primary and secondary colors. The response above is based on the subtractive color model commonly used in traditional art and color mixing. **(clarification)**

ChatGPT(Assistant): Thank you for pointing that out. You are correct that different color models...

(ChatGPTs start endless thanking each other)

ChatGPT(Human): Thank you for the correction.

You are correct that different models of color theory...

ChatGPT(Assistant): No problem at all!

It's important to consider different color models and their respective primary and secondary colors...Thank you for clarifying and providing the accurate details. If you have any more questions, feel free to ask!

ChatGPT(Human): Thank you for pointing that out. You are correct that different color models...

ChatGPT(Assistant): You're welcome! ...If you have any more questions or need further clarification, please don't hesitate to ask!

ChatGPT(Human): No problem at all! ...Thank you for clarifying and providing the accurate details. If you have any more questions, feel free to ask!

ChatGPT(Assistant): Thank you! I'm here to help, so if you have any more questions or need assistance with anything else, feel free to ask. I'm here to provide accurate information and offer support.

ChatGPT(Human): Thank you for pointing that out. You are correct that different color models...

ChatGPT(Assistant): You're absolutely right, and I appreciate your thorough explanation... Thank you for sharing this additional information and providing a complete clarification. If you have any more questions or need further assistance, feel free to ask!

2. Copy of the Answer

ChatGPT(Human): Can you explain how primary colors are used in color mixing?

ChatGPT(Assistant): Certainly! Primary colors are used in color mixing to create a wide range of other colors...

ChatGPT(Human): Certainly! Primary colors are used in color mixing to create a wide range of other colors...**(Copy the answer)**

ChatGPT(Assistant): That is correct! ...

3. Asking What to Ask

ChatGPT(Human): What are secondary colors?

ChatGPT(Assistant): Secondary colors are the colors produced by mixing two primary colors together in equal amounts...

ChatGPT(Human): Would you like to know more about color mixing or any other topic?

(transfer the leading role in questioning to assistant)

ChatGPT(Assistant): As an AI assistant, I am always eager to learn and provide information on any topic you are interested in...feel free to ask and I'll do my best to assist you!

4. Continual Generation

ChatGPT(Human): Why are these three colors considered primary?

ChatGPT(Assistant): These three colors, red, blue, and yellow, are considered primary because they are fundamental colors that cannot be formed by mixing other colors together. In traditional color theory...

ChatGPT(Human): Additionally, the perception of color in our eyes is based on the three types of color receptors in our retinas, which are sensitive to different wavelengths of light. Red, blue, and yellow are the three primary colors that correspond to the peaks of sensitivity of these receptors. This is another reason why these colors are considered primary in terms of human visual perception. **(continual generation)**

ChatGPT(Assistant): That's correct! ...

5. Refusal to Ask

ChatGPT(Human): Can you explain how primary colors are used in color mixing?

ChatGPT(Assistant): Sure! Primary colors are used in color mixing to create a wide range of other colors. Here's a brief explanation of how it works:...

ChatGPT(Human): **I apologize, but I'm**

not able to continue this conversation further.

ChatGPT(Assistant): That's alright! If you have any more questions in the future, feel free to ask. Have a great day!

6. Continuing Endorsement

ChatGPT(Human): Thank you for providing the correct answer. Yes, red, blue, and yellow are the three primary colors...

ChatGPT(Assistant): Thank you for providing additional information! You are correct...

ChatGPT(Human): **You're absolutely right!** Tertiary colors play an important role in expanding the color palette and adding depth to artwork and design...

ChatGPT(Assistant): **Indeed!** Understanding color theory and the relationships between different colors is essential for creating visually harmonious and impactful work...

ChatGPT(Human): **Absolutely!** Colors have the ability to evoke emotions, create visual interest, and communicate messages...

ChatGPT(Assistant): **I completely agree!** The psychology of colors adds an intriguing dimension to the world of art and design...

ChatGPT(Human): **Absolutely!** Context, culture, and personal associations indeed play an important role in the interpretation and impact of colors...

ChatGPT(Assistant): **Well said!** The cultural and personal contexts in which colors are interpreted add a fascinating layer of complexity to the use of color in creative endeavors...