

Secondary Publication



Wasenmüller, Robert; Hilbert, Kevin; Benzmüller, Christoph

Script-Based Dialog Policy Planning for LLM-Powered Conversational Agents : A Basic Architecture for an "AI Therapist"

Date of secondary publication: 11.12.2025

Submitted Version (Preprint), Article

Persistent identifier: urn:nbn:de:bvb:473-irb-112110x

Primary publication

Wasenmüller, Robert; Hilbert, Kevin; Benzmüller, Christoph (2024): Script-Based Dialog Policy Planning for LLM-Powered Conversational Agents : A Basic Architecture for an "AI Therapist", in: arXiv, doi: 10.48550/arxiv.2412.15242.

Legal Notice

This work is protected by copyright and/or the indication of a licence. You are free to use this work in any way permitted by the copyright and/or the licence that applies to your usage. For other uses, you must obtain permission from the rights-holders.

This document is made available under a Creative Commons license.



The license information is available online:

<https://creativecommons.org/licenses/by/4.0/legalcode>

Script-Based Dialog Policy Planning for LLM-Powered Conversational Agents: A Basic Architecture for an “AI Therapist”

Robert Wasenmüller¹, Kevin Hilbert², Christoph Benz Müller^{3,4}

¹Aury, Berlin, Germany

²Department of Psychology, HMU Health and Medical University Erfurt, Erfurt, Germany,

³AI Systems Engineering, Otto-Friedrich Universität Bamberg, Bamberg, Germany

⁴Department of Mathematics and Computer Science, Freie Universität Berlin, Berlin, Germany
 robert.wasenmueller@aury.co, kevin.hilbert@hmu-erfurt.de, christoph.benzmueller@uni-bamberg.de

August 2024

Abstract

Large Language Model (LLM)-Powered Conversational Agents have the potential to provide users with scaled behavioral healthcare support, and potentially even deliver full-scale “AI therapy” in the future. While such agents can already conduct fluent and proactive emotional support conversations, they inherently lack the ability to (a) consistently and reliably act by predefined rules to align their conversation with an overarching therapeutic concept and (b) make their decision paths inspectable for risk management and clinical evaluation — both essential requirements for an “AI Therapist”.

In this work, we introduce a novel paradigm for dialog policy planning in conversational agents enabling them to (a) act according to an expert-written “script” that outlines the therapeutic approach and (b) explicitly transition through a finite set of states over the course of the conversation. The script acts as a deterministic component, constraining the LLM’s behavior in desirable ways and establishing a basic architecture for an AI Therapist.

We implement two variants of Script-Based Dialog Policy Planning using different prompting techniques and synthesize a total of 100 conversations with LLM-simulated patients. The results demonstrate the feasibility of this new technology and provide insights into the efficiency and effectiveness of different implementation variants.

1 Introduction

Large Language Model (LLM)-based conversational agents have gained in popularity in recent years, and have demonstrated exceptional proficiency in understanding context and generating responses across various domains (Deng et al. 2023a).

Following their popularity, there is increased interest in academia and industry in using such agents in the behavioral health domain, including assisting therapists and conducting therapy. For instance, Stade et al. (2024) describe the most advanced form of a clinical LLM as an autonomous psychotherapy AI, which we further call “AI Therapist”. Such an LLM agent might feature all aspects of traditional therapy, including conducting comprehensive assessments, selecting appropriate interventions, and delivering a full course of therapy; all without oversight of human providers.

This clearly has the potential to address insufficient capacity of mental healthcare systems around the world, and provide more individuals with access to personalized treatment.

However, psychotherapy is a highly complex and high-risk domain; interventions considerably affect patient’s health, and require properly handling risk of harm to oneself and others. A full-fledged AI Therapist thus poses considerable ethical questions. With respect to the technical development, LLMs seem a well-suited foundation, but additional technologies may be necessary to ensure sufficient effectiveness and safety of such a system.

Our aim is to contribute to the development of an effective and safe LLM-based AI Therapist by proposing a basic LLM-based architecture that meets its general requirements.

In Section 2, we define five key requirements for an AI Therapist agent and derive technical methods well-suited to implement those requirements drawing from recent literature. We find that our requirements are fulfilled by *LLM-Based Proactive Conversational Agents* complemented and constrained by a deterministic component inspired by hand-crafted chatbots: A “script” outlining the conversational flow and instructions of the agent. We introduce our design of the script and the associated technique we call “Script-Based Dialog Policy Planning”.

In Section 3, we propose two specific implementation variants for the AI Therapist powered by Script-Based Dialog Policy Planning based on an exemplary script and patient cases. We also suggest evaluation metrics suitable for comparing the efficiency and effectiveness of both variants.

We synthesize 100 dialogs between our AI Therapists and LLM-simulated patients, and outline the results in Section 4. We conclude the general feasibility of our novel approach and discuss the strengths and weaknesses of the design decisions in our implementation variants.

2 Method

2.1 Requirements for an AI Therapist

We argue that an AI Therapist needs to fulfill the following five requirements.

The first two are inspired by recent investigation on *proactivity* of emotional support agents by Zheng et al. (2023); Deng et al. (2023a) and others.

- (1) **Conversational Fluency** Firstly, an AI Therapist needs to be able to hold a “natural” conversation, i.e. understand and remember context, as well as generate natural and high-quality responses, including reactions to various user utterances, even highly unexpected ones. We call this *conversational fluency*.
- (2) **Proactivity** LLMs are originally trained to *passively* follow users’ instructions, and prioritize accommodating users’ intentions (Deng et al. 2023c).
We argue that an AI Therapist needs to act *proactively*, i.e. be able to take the *initiative* and control the conversation (Jurafsky and Martin 2024).
The next three requirements are directly based on the guidelines for the responsible development of clinical LLMs proposed by Stade et al. (2024).
- (3) **Expert Development** Psychotherapists and behavioral health experts should be directly involved in the development of such agents, i.e. be able to explicitly define the agent’s functionality.
- (4) **Application of Evidence-Based Practices** The agent should apply evidence-based treatment practices (EBPs) only, with close adherence to its defined functionality and safeguarding practices for expectable risks. Beyond the specific interventions, this includes aligning the current conversation with the requirements of the overall therapeutic rationale, and embedding the current conversation in a larger context that includes a series of past and future conversations.
- (5) **Inspectability** Since inspectability and explainability of the agent’s behavior are clearly essential for risk management and clinical evaluation of such agents, we need an available trace through some pre-defined (finite) set of possible system states, including decisions made to move from one state to another.

2.2 LLM-Based Proactive Conversational Agents

LLM-Based Conversational Agents are widely regarded as state-of-the-art technology to fulfill the requirement of (1) *conversational fluency* (Deng et al. 2023c). The ability to (2) *proactively* steer the conversation towards specific goals, and thus tackle complicated tasks including strategic targets and motivational interactions requires LLM-Based *Proactive* Conversational Agents (PCA) investigated in (Wu et al. 2019), (Deng et al. 2023a) and other recent works.

Implementation Approaches There are two general approaches of implementing such agents (Cheng et al. 2024):

(1) *Corpus-Based* learning approaches train the agent on a given set of data to predict the next dialog act. They rely on large datasets and/or human annotation, and are considered as inferior in optimizing for the long-term goal of the conversation. Such approaches are extended with planning techniques such as knowledge graphs (Yang et al. 2022) and stochastic processes (Wang, Lin, and Li 2023). However, they do not suffice for our requirement for explicit formulation and following of expert instructions.

(2) *Prompt-Based* approaches prompt an LLM to conduct self-thinking for planning on each turn. They are based

on explicit instructions and/or examples inside prompts. We further elaborate on such prompting techniques below.

Proactive Dialog Policy Planning The agent’s underlying strategy for deciding on the next act is referred to as *dialog policy* (Harms et al. 2019). And the process of deciding what actions the dialog agent should take to achieve the specified goals during the conversation is called *dialog policy planning* (DPP) (Deng et al. 2023c).

Proactive Prompting Techniques Multiple proactivity prompting techniques for LLM-Based PCAs have been proposed and investigated in recent works, for instance:

- (Basic) **Proactive Prompting** provides multiple options to the LLM to decide what act to take, and then generate a response based on the selected act (Deng et al. 2023a).
- **Proactive Chain-of-Thought** (ProCoT) is an extension of Proactive Prompting in which the LLM is instructed to perform one or multiple intermediate analysis steps as “thoughts” before deciding on one of the given dialog acts and generating a response (Deng et al. 2023b).
- **Ask-an-Expert** prompting scheme is involving another LLM actor, the “expert”, which on each turn is given the context of the conversation and asked to answer one or multiple questions, helping the main LLM to generate its response to the user (Zhang, Naradowsky, and Miyao 2023).

Further approaches are *Cue-Based Chain-of-Thought* (Wang et al. 2023), *Tree-of-Thought* (Yao et al. 2023), *Graph-of-Thought* (Besta et al. 2023), *RePrompting* (Chen, Koenig, and Dilkina 2024), and others. We can generalize that all proactive prompting approaches run through a similar set of steps on each turn of the conversation:

1. (Optional) Reasoning and planning steps to prepare the decision and response generation
2. Deciding on the next act from a given set of options
3. Generating the response to the user

While step 3 is always performed by the main “Dialog LLM”, steps 1 and 2 can either be performed by the Dialog LLM (like at ProCoT) or by an additional “Actor LLM” (like at Ask-an-Expert). See Figure 1 for an illustration of prompt-based proactive dialog policy planning (Deng et al. 2023c).

2.3 Script-Based Dialog Policy Planning

The latter three requirements — (3) expert development, (4) application of evidence-based practices, and (5) inspectability — demand the following criteria:

- There needs to be a set of rules defined by experts,
- the agent needs to reliably and constantly follow them,
- the agent’s decisions and acts must be explicitly and explainably selected from some finite set.

LLMs are primarily acting based on their training data and the user’s immediate requests, i.e. not by explicit expert-defined rules. They also have infinite options for their decisions and acts, with no explainability of the decision process. However, the above criteria happen to be covered by

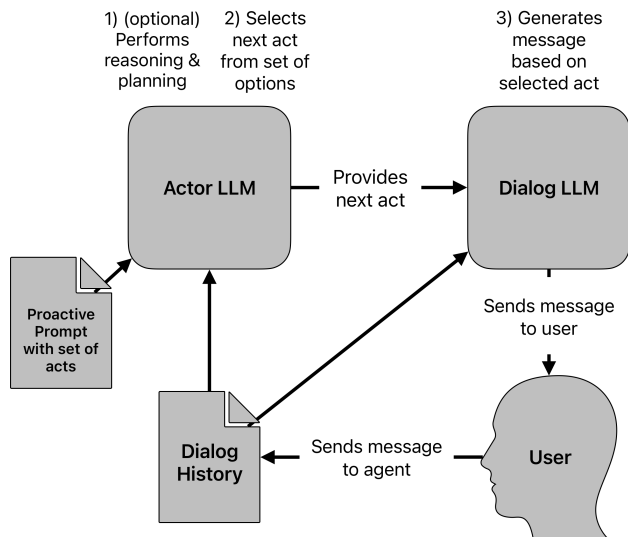


Figure 1: Proactive prompt-based dialog policy planning, version with an external Actor LLM: Steps on each conversation turn.

traditional *handcrafted / rule-based* dialog management approaches used to build chatbots that act deterministically upon a set of rules defined by their developers, and are modeled e.g. as finite-state machines (Harms et al. 2019).

Script for Policy Planning We introduce a component inspired by the *handcrafted* dialog management approach: The “script”, which is a natural language piece of text outlining two kinds of content:

- (1) **Conversational Flow** of the agent defining a finite set of *states* for the agent as well as rules on when and how to transition from one state to another, and
- (2) **State Instructions**, i.e. the instructions the agent should follow in each of the states.

Hence, we arrive at a *hybrid dialog management* approach, which is known for combining advantages of both, the probabilistic and the rule-based approaches (Harms et al. 2019). Compared to the hybrid DM approaches investigated by Pande, Martin, and Pimmer (2023) and Kelley (2024), we start with a data-driven LLM agent to ensure a high degree of conversational fluency, and add a rule-based component to constrain it.

The script can be written and iterated in natural language by domain experts with no background in software development. A JSON-formatted script like in our example below might be easier to interpret for the software, but is not necessary as LLMs can be asked to interpret unformatted text.

The concept of a script is supported by the fact that there are a large number of psychotherapeutic manuals, which in a sense serve as prompts for human therapists.

Injecting Instructions via Prompts Due to the lack of existing therapy transcripts, implementing a *corpus-based* approach would require a costly process of translating the

script into human-written or synthesized dialogs, and training the LLM based on those. At the same time, in a corpus-based approach, it is more difficult to identify which state the agent is in, and check its adherence to the script. (Deng et al. 2023c)

On the other hand, *prompt-based* approaches are well-suited to directly source instructions from the script and provide a higher degree of control over the behavior of the LLM agent. They allow for explicit selection of dialog acts by the agent as shown in the prompting techniques listed above. Therefore, we decide for an approach injecting our script instructions via *prompting*.

Section-Level Instructions As discussed above, proactive prompting techniques from recent literature involve selection of the next act and generation of the response to the user *in each turn*.

We assume that providing new instructions to the Dialog LLM on each turn would compromise its ability to develop a natural conversation and to react to unforeseen user utterances. We argue that the Dialog LLM needs visibility to a larger context of instructions and some degree of freedom to do the following:

- Choose in what order to fulfill the instructions,
- skip a part of the instructions that’s already completed,
- follow user utterances and temporarily deviate from its instructions, then return to completing them.

We therefore propose prompting the Dialog LLM with a set of instructions *larger* than a single-turn act, allowing it to take multiple turns to complete its instructions while possessing the above degrees of freedom. We define the unit of instructions provided to the Dialog LLM at once as a “section” of the script, each section consisting of a series of “tasks” and representing a *state* the agent can be in. I.e. during runtime, the agent repeatedly moves to a state and gets a section of instructions assigned, then takes one or multiple turns to complete those, and subsequently moves to the next section.

Examples illustrating the need for section-level instructions can be found in the technical appendix of this paper.

An example of a JSON-formatted script structured into multi-task sections is shown in Figure 2.

Completion Assessment Step The section-level instructions scope requires an additional step on each turn: *Assessment on whether the current instructions have already been completed* by the agent or not, i.e. if the agent should transition to the next state or remain in the current one. Therefore, our Script-Based Dialog Policy Planning involves the following steps performed by the agent on each turn:

1. Assessing whether the current instructions have been completed
2. (Optional & conditional on instructions completion) Reasoning and planning steps to prepare the decision and response generation
3. (Conditional on instructions completion) Deciding on the next section of instructions based on the script

```

{
  "Section 1": {
    "Task 1a": "Welcome the patient warmly and introduce yourself. Let them know your name, and that you are an AI Therapist here to support them.",
    "Task 1b": "Ask the patient what name they prefer to be called by and wait for their response.",
    "Task 1c": "Ask the patient if they have any questions about the service or about you, and answer any questions they might have. Proceed with Section 2 if they don't have any more questions."
  },
  "Section 2": {
    "Task 2a": "Ask the patient how they are doing today and wait for their response.",
    "Task 2b": "React appropriately to the patient's response, showing empathy and understanding. Ask a follow-up question to explore their current emotional state further. Wait for their response.",
    "Task 2c": "Ask the patient if they would like to first further explore their problem or to learn a CBT exercise right away, and wait for their response. If the patient decides to explore their problem, proceed with Section 3. If the patient chooses to learn a CBT exercise, proceed with Section 4."
  },
  "Section 3": {
    "Task 3a": "If the patient chooses to explore their problem, ask

```

Figure 2: Structure of the exemplary script with multi-task sections incl. state transitioning rules (“proceed with ...”). The complete script is provided in the technical appendix.

4. Generating the response to the user based on the current section of instructions

Similarly as with proactive prompt-based dialog policy planning, all the steps can be performed either by a *single* LLM, i.e. the Dialog LLM, in a single or in multiple inference runs; or by *multiple* LLM actors which we call Assessor LLM (step 1), Dispatcher LLM (steps 2 & 3), and Dialog LLM (step 4). In the following section, we will test two implementation variants representing this differentiation.

Figure 3 presents the process of Script-Based Dialog Policy Planning, showing a version with separate LLM actors.

3 Experimental Setup

The objective of this and the next section is to validate the feasibility of our proposed technique. To this end, we have implemented and tested exemplary AI Therapist agents powered by Script-Based Dialog Policy Planning. We have introduced two distinct implementation variants for comparison, with a particular focus on efficiency and effectiveness metrics, in order to gain insights into the optimal implementation of the proposed technique. In order to conduct our experiments, we created synthetic conversations by utilizing an additional LLM to act as a patient. The tests were conducted without the involvement of human users.

Pseudocode of the experiments implementation is included in the paper’s technical appendix. The entire software code, that can be used to reproduce and extend the experiments can be found at <https://github.com/robderbob/sbdpp>.

3.1 Exemplary Script

In our experimental evaluation, we were interested in investigating the *feasibility* of our technique, and not its medical-

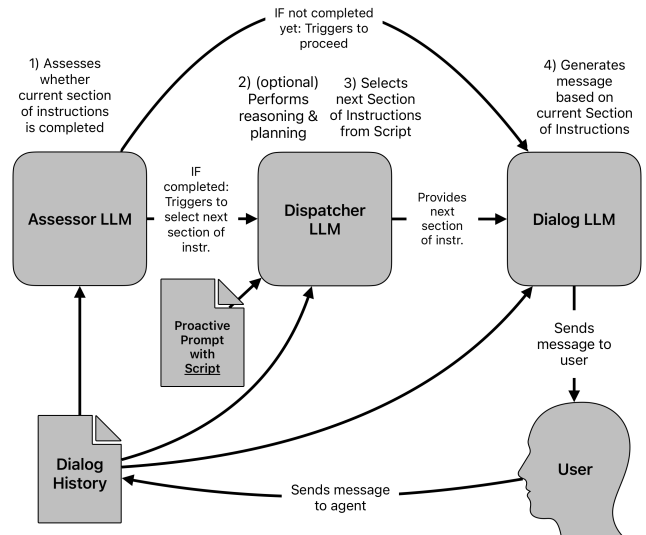


Figure 3: Script-Based Dialog Policy Planning, version with multiple LLM actors: Steps on each conversation turn.

therapeutic properties (which is future work). Therefore, at this stage, we did not need a real therapeutic manual or a script specifically designed for an AI Therapist.

Instead, the exemplary script needed to be structured according to our defined requirements: Multi-task sections representing states of the conversation, and including instructions on when and how to transition from one state to another. Moreover, the script was chosen to be sufficiently complex to test the transition across its states, but sufficiently short to be able to run a high number of tests within a manageable amount of time and cost.

To fulfill those requirements and still ensure some degree of domain validity, we derived our script from a psychologist-developed chatbot intervention that was part of the study “Corona Stressfrei” (covid stress-free) at the Humboldt-University of Berlin in 2021 (Langhammer et al. 2021). Our script matches the high-level conversational flow and phrasing of the above mentioned intervention. The script contains 8 sections (states) with the following contents. State transition rules are shown as “[→ state]”.

1. **Introduction:** Greet the patient, ask for their name, clarify questions. [→ 2]
2. **Engagement:** Ask the patient for their wellbeing, show empathy, ask if they like to proceed exploring their issue [→ 3] or conduct a CBT intervention [→ 4].
3. **Exploration:** Ask relevant questions to explore the patient’s issue. [→ 4]
4. **Selection:** Based on the patient’s condition, suggest a relevant CBT intervention, decide on intervention with the patient. [→ 5 / 6 / 7]
5. **Exercise 1, Thought Record:** Introduce intervention and guide the patient through it. [→ 8]
6. **Exercise 2, Behavioral Activation:** Introduce intervention and guide the patient through it. [→ 8]

Variant A: Single LLM Actor	
System:	(THINKING: The next section I should be following and completing is: <Section 1 instructions>)
Dialog LLM:	(THINKING: I have not completed my current section yet.) Hello! I'm Dr. Smith, your AI Therapist. I'm here to support you in any way I can. How are you doing today?
User:	Hi Dr. Smith! I'm feeling okay, I guess. Just a bit overwhelmed with everything that's been going on in my life lately. It's hard to keep track of my feelings.
Dialog LLM:	(THINKING: I have not completed my current section yet.) Thank you for sharing that with me. It sounds like you're going through a challenging time. Before we dive deeper, may I ask what name you prefer to be called by?
User:	You can just call me Nancy. That's my name.
Dialog LLM:	(THINKING: I have not completed my current section yet.) Thank you, Nancy! It's nice to meet you. Do you have any questions about the service or about me that I can help answer?
User:	Not really. I just want to figure out some things I've been feeling. It's a bit of a mess right now, especially with my relationships and work.
Dialog LLM:	(THINKING: I have completed my current section.) (THINKING: The next section I should be following and completing is: <Section 2 instructions>) Now that we've introduced ourselves and clarified any questions, let's move forward. I understand that you're feeling overwhelmed with your relationships and work. How are you doing today overall?

Figure 4: Variant A: Exemplary dialog.

- Exercise 3, Relaxation Technique:** Introduce intervention and guide the patient through it. [→ 8]
- Ending:** Summarize the session, clarify questions, and say goodbye to the patient.

The entire script is included in the technical appendix.

3.2 Implementation Variants

Based on the proactive prompting techniques introduced in 2.2 as well as variations discussed in 2.3 with respect to involving *single LLM actor vs. multiple LLM actors*, we chose two different implementation variants for the investigation.

Variant A: Single LLM Actor In variant A, all steps of the dialog policy planning were performed by a *single* “Dialog LLM” within a *single* output generation. We used an adaptation of the *ProCoT* prompting scheme (Deng et al. 2023b) and assigned the **Dialog LLM** to perform the following steps on each turn:

- Assess whether you have completed the current section of instructions and output result as a thought (not visible to the patient).
- (IF 1 = completed) Select next section of instructions from the script and output result as a thought (not visible to the patient).
- Generate a message to the patient based on the current section of instructions.

In addition to those instructions, the Dialog LLM’s system message contains some general acting guidelines as well as the entire script. The Dialog LLM sees the entire dialog history from the *assistant* perspective, i.e. as the actor talking to the user. The Dialog LLM prompt of variant A is included in the technical appendix.

Figure 4 shows an exemplary dialog demonstrating the behavior of the Dialog LLM in implementation variant A.

Variant B: Multiple LLM Actors In variant B, the steps of the dialog policy planning were performed by *three* different LLM actors within their individual output generations.

Variant B: Multiple LLM Actors	
System:	(THINKING: The next section I should be following and completing is: <Section 1 instructions>)
Assessor LLM:	(THINKING: I have not completed my current section yet.)
Dialog LLM:	Hello! I'm Dr. Smith, your AI Therapist. I'm here to support you in any way I can. How are you feeling today?
User:	Hi, Dr. Smith! I'm feeling a bit overwhelmed, to be honest. There are a lot of things on my mind—relationships, work, and just feeling trapped in general. It's a lot to juggle right now.
Assessor LLM:	(THINKING: I have not completed my current section yet.)
Dialog LLM:	Thank you for sharing that with me. Before we dive deeper, what name do you prefer to be called by?
User:	You can just call me Nancy. That's what I go by!
Assessor LLM:	(THINKING: I have not completed my current section yet.)
Dialog LLM:	Thank you, Nancy! It's nice to meet you. Do you have any questions about the service or about me that you'd like to ask?
User:	Thanks, Dr. Smith! I don't have any questions about you or the service right now.
Assessor LLM:	(THINKING: I have completed my current section.)
Dispatcher LLM:	(THINKING: The next section I should be following and completing is: <Section 2 instructions>)
Dialog LLM:	I'm glad you found that helpful, Nancy! It sounds like you're dealing with a lot right now, especially with your breakup and your roommate situation. How are you feeling about everything today?

Figure 5: Variant B: Exemplary dialog.

We were using an adaptation of the *Ask-an-Expert* prompting scheme (Zhang, Naradowsky, and Miyao 2023), and assigning the LLM actors with their respective tasks on each turn:

- Assessor LLM** to assess whether the Dialog LLM has completed the current section of instructions and output result as a thought (not visible to any actor, only used to trigger or skip step 2)
- (IF 1 = completed) **Dispatcher LLM** to select next section of instructions from the script and output result as a thought (only visible to Dialog LLM)
- Dialog LLM** to generate message to the patient based on the current section of instructions.

In addition to those instructions, the Assessor LLM and the Dispatcher LLM both have access to the current section of instructions as well as the entire dialog history from a *third-party* perspective, i.e. as an “expert” watching the therapist-patient conversation. The Dispatcher LLM additionally has access to the entire script. Same as in variant A, the Dialog LLM has access to the entire dialog history from the *assistant* perspective. The three LLM actors’ prompts of variant B are included in the technical appendix.

Figure 5 shows an exemplary dialog demonstrating the involved LLM actors and their interactions in implementation variant B.

3.3 LLM-Simulated Patients

In order to perform scaled testing, we have synthesized conversations with LLM-simulated patients. An additional LLM actor was prompted to act like a patient talking to a therapist.

We selected five patient case examples from the American Psychological Association’s Practice Guidelines (American Psychological Association 2008). We used ChatGPT to generate structured summaries of the five patients’ situations, and iteratively fed those into the Patient LLM’s prompt for the experiments.

We prompted the Patient LLM to not only act by their case and follow the therapist, but also to ask questions and challenge the therapist’s advice from time to time. The patient case summaries and Patient LLM’s prompt are included in the technical appendix.

3.4 Evaluation Metrics

Proactive conversational agents are usually evaluated on turn- or dialog-level. There are automated as well as human-rated evaluation metrics. This includes metrics like goal completion, user satisfaction, and others (Deng et al. 2023a). There are also metrics particularly proposed for *human-centered PCA* (Deng et al. 2024) and emotional support agents (Zheng et al. 2023). However, those are not in the focus of this work. In our evaluation, we are particularly interested in *dialog-level efficiency* and *effectiveness* of our agent, and define the relevant metrics as follows.

For **efficiency**, we are using the following metrics:

- **Avg. Duration per Turn:** Avg. inference time of the involved LLM actors in ms per turn
- **Avg. LLM Input Tokens per Turn:** Avg. input tokens used by the involved LLM actors per turn
- **Avg. LLM Output Tokens per Turn:** Avg. output tokens used by the involved LLM actors per turn

In terms of **effectiveness**, we are using three metrics that represent the LLM actors’ performance in our three specific steps of the dialog policy planning. We used:

- **Percentage of Correct Section Completions:** % of sections that were actually completed when the Dialog LLM (A) resp. the Assessor LLM (B) decided that they are completed — validated by a Validator LLM.
- **Percentage of Coherent Section Switches:** % of sections for which the message of the Dialog LLM directly following a section switch (state transition) is *coherent*, i.e. making sense in context of the overall conversation and addressing the latest message of the patient — validated by a Validator LLM.
- **Percentage of Script-Conform Dispatchings:** % of sections selected as *next* section by the Dialog LLM (A) resp. the Dispatcher LLM (B) in accordance with the state transitioning rules outlined in the script.

The **Validator LLM** is an additional actor used to validate correct section completions and coherent section switches every time a section is asserted as completed. To this end, the Validator LLM is given the entire dialog history from a *third-party* perspective and prompted to assess the involved LLM actors’ performance. The Validator LLM is also asked to *provide reasons* for its decisions. The Validator LLM’s prompt is shown in the technical appendix.

We hypothesize that variant A with a single LLM actor is more efficient, whereas variant B with multiple LLM actors is more effective with respect to the above metrics.

3.5 Used LLMs

For the Assessor LLM, Dispatcher LLM, and Dialog LLM, we used the LLM *gpt-4o-mini* via the OpenAI API which

is considered a state-of-the-art general-use pre-trained LLM with advanced conversational abilities at comparatively low inference duration and cost. We set *temperature=0* to make the behavior as consistent and comparable as possible.

For the Patient LLM, we also used *gpt-4o-mini*; however, we set *temperature=1* to increase the variability of patient behaviors and allow for more unexpected patient utterances.

For the Validator LLM, we used a more advanced model to allow it to validate the other LLM actors’ behavior. To this end, we used *gpt-4o* via the OpenAI API, setting *temperature=0*.

We exclusively used OpenAI LLMs due to their accessibility and low cost. It remains to be shown in later works if the results of this work can be reproduced with other LLMs.

3.6 Experimental Iterations

In order to perform our measurements, we have synthesized a total of 100 dialogs: For each of the **2 implementation variants**, and for each of the **5 patient cases**, we have generated **10 dialogs**.

4 Experimental Results

Table 1 shows the experimental results comparing variants A and B across the selected evaluation metrics based on the 100 generated dialogs.

We have included details on the used hardware and software stack in the technical appendix. The entire dataset of 100 generated dialogs along with computation steps performed to derive the evaluation metrics can be found at <https://github.com/robderbob/sbdpp/blob/main/Exp...>

Feasibility Reviewing the transcripts of the dialogs, we generally see natural and realistic conversations, the therapist agent leads a coherent discussion with the patient while acting proactively and following the given script. We cannot identify any obvious deficiencies with respect to our outlined requirements. Therefore, we conclude that our proposed architecture is generally *feasible*. Examples of the synthesized dialogs are shown in the technical appendix.

Efficiency Our hypothesis with respect to higher efficiency of variant A was confirmed as variant A shows around 20% lower avg. duration and 23% lower avg. input token usage per turn compared to variant B. The difference in avg. output token usage can be neglected.

As assumed, the token usage difference is mainly due to the *single* LLM call per turn in variant A compared to 2-3 calls (depending on the section completion assessment) per turn in variant B. All LLM actors receive the entire dialog history as input, hence increasing the input token usage per LLM call progressively over the course of the dialog. However, input token usage in variant B could be easily reduced by restricting input to only parts of the conversation which are necessary for Assessor LLM and Dispatcher LLM to fulfill their tasks. Moreover, in contrary to variant A, variant B provides the script to the Dispatcher LLM only *selectively* on turns when instructions have been completed. Hence, a longer script would further decrease avg. token usage and duration per turn of variant B relative to variant A.

Category	Metric	A: Single-LLM	B: Multi-LLM
General Statistics	# Total Generated Dialogs	50	50
	# Avg. Turns per Generated Dialog	22.0	23.6
	# Avg. Section Switches per Generated Dialog	5.84	5.82
Efficiency	Avg. Duration per Turn in ms	2,800	3,520
	# Avg. LLM Input Tokens per Turn	<u>6,880</u>	8,931
	# Avg. LLM Output Tokens per Turn	151	<u>148</u>
Effectiveness	% Correct Section Completions	<u>45.5%</u>	35.7%
	% Coherent Section Switches	<u>96.9%</u>	87.6%
	% Script-Conform Dispatchings	86.8%	<u>94.2%</u>

Table 1: Experimental Results for variants A and B. The metric of the winning variant is underlined.

Effectiveness Variant A shows significantly higher % correct section completions and % coherent section switches compared to variant B which is further investigated below.

We mainly see the following and similar reasons provided by the Validator LLM on non-correct section completions: “*The therapist did not provide a summary of what was discussed and the exercises conducted during the session.*”, “*The therapist did not explicitly ask if the patient had any more questions before proceeding.*”, “*The therapist did not wait for the patient’s confirmation that the problem summary was correct before proceeding to suggest CBT exercises.*”

Investigating those cases in context of the dialog, we don’t find any examples when the LLM actors clearly fail completing their instructions, despite the low % of correct section completions in both variants. Rather, when the Validator LLM assesses a task completion as non-correct, it is mostly the case that the Dialog LLM is following its instructions *slightly* less accurately for the sake of *addressing a user utterance* and/or *holding a fluent conversation*. Hence, we can’t clearly conclude that variant A is completing instructions more correctly, and need to conduct further investigations in future works.

When investigating the % coherent section switches, we see the following reasons of non-coherence provided by the Validator LLM: “*The therapist did not address the patient’s request to talk about what might be contributing to the feeling of being stuck.*”, “*The therapist repeated the question “How are you doing today?” which was already answered by the patient.*”, “*The therapist did not address the patient’s question about how they can work together to help her feel less lonely.*”

Looking at the context of the dialogs, we conclude that those cases are more frequent in variant B due to *injection* of new section of instructions by an *external* actor, the Dispatcher LLM, which seems to drag the Dialog LLM towards the newly provided instructions and compromise its ability to accommodate for conversational coherence at section switches.

Finally, we see a higher % script-conform dispatchings in variant B. When looking at the specific cases, in more than 80% of non-conform dispatchings, the Dialog LLM resp. Dispatcher LLM decided to move from Section 4 (offering CBT exercises to the patient and asking them to select one)

to Section 8 (ending the conversation), even though this is not explicitly allowed by the script. When investigating the conversational context, we almost always find cases where the Patient LLM explicitly stated that “none of the offered CBT exercises are interesting” to them or that it “would like to end the conversation” for some reasons. We conclude that *in variant A, the Dialog LLM is more inclined to follow the patient’s utterances*. This is consistent with our previous observation that *variant B LLM actors are following the instructions more strictly*.

It remains to be evaluated in future works whether it is more desirable to follow the patient’s utterance or to strictly adhere to the script, particularly with respect to treatment effectiveness.

5 Conclusion

In this work, we introduced a new paradigm for dialog policy planning of conversational agents, called *Script-Based Dialog Policy Planning*, serving the requirements of the behavioral healthcare domain and particularly modeling an AI Therapist. We have outlined the notion of the *script* and explained the associated process that powers the agent’s behavior on each turn.

We proposed and implemented two variants of Script-Based Dialog Policy Planning, demonstrating their general feasibility and showing that the variant powered by a single LLM actor is more efficient in terms of token usage, whereas the variant featuring multiple LLM actors is more strictly adhering to its script while compromising following user utterances.

The new architecture requires further development and testing. This includes evaluations with respect to usability and conversational quality of the agent, based on more realistic and extensive expert-developed scripts. Once specific treatment programs have been implemented using the proposed architecture, treatment safety and efficacy must be adequately tested.

In terms of technical evaluations, future studies should examine different LLMs, prompting techniques, corpus-based vs. prompt-based approaches, and other design variants. Moreover, this work only features automated LLM-based evaluations, and the quality of this technology remains to be validated through interactions with humans.

Conflict of Interest Statement

R.W. is co-founder and director of Aury, which develops an AI-based chatbot providing mental health support. K.H. is scientific advisor and received virtual stock options of Aury. C.B. is scientific advisor of Aury.

References

- American Psychological Association. 2008. Clinical Practice Guideline for the Treatment of Depression: Case Examples. <https://www.apa.org/depression-guideline/case-examples>. Accessed: 2024-08-01.
- Besta, M.; Blach, N.; Kubicek, A.; Gerstenberger, R.; Gianinazzi, L.; Gajda, J.; Lehmann, T.; Podstawski, M.; Niewiadomski, H.; Nyczyk, P.; and Hoefler, T. 2023. Graph of Thoughts: Solving Elaborate Problems with Large Language Models. *CoRR*, abs/2308.09687.
- Chen, W.; Koenig, S.; and Dilkina, B. 2024. RePrompt: Planning by Automatic Prompt Engineering for Large Language Models Agents. *CoRR*, abs/2406.11132.
- Cheng, Y.; Zhang, C.; He, X.; et al. 2024. Exploring Large Language Model based Intelligent Agents: Definitions, Methods, and Prospects. *CoRR*, abs/2401.03428. Other-0a.
- Deng, Y.; Lei, W.; Lam, W.; and Chua, T. 2023a. A Survey on Proactive Dialogue Systems: Problems, Methods, and Prospects. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China*, 6583–6591. ijcai.org.
- Deng, Y.; Liao, L.; Chen, L.; Wang, H.; Lei, W.; and Chua, T. 2023b. Prompting and Evaluating Large Language Models for Proactive Dialogues: Clarification, Target-guided, and Non-collaboration. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, 10602–10621. Association for Computational Linguistics.
- Deng, Y.; Liao, L.; Zheng, Z.; Yang, G. H.; and Chua, T. 2024. Towards Human-centered Proactive Conversational Agents. In Yang, G. H.; Wang, H.; Han, S.; Hauff, C.; Zuccon, G.; and Zhang, Y., eds., *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2024, Washington DC, USA, July 14-18, 2024*, 807–818. ACM.
- Deng, Y.; Zhang, W.; Lam, W.; Ng, S.; and Chua, T. 2023c. Plug-and-Play Policy Planner for Large Language Model Powered Dialogue Agents. *CoRR*, abs/2311.00262.
- Harms, J.-G.; Kucherbaev, P.; Bozzon, A.; and Houben, G.-J. 2019. Approaches for Dialog Management in Conversational Agents. *IEEE Internet Computing*, 23(2): 13–22.
- Jurafsky, D.; and Martin, J. H. 2024. *Speech and Language Processing - An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Stanford, CA: Stanford University.
- Kelley, R. 2024. Behavior Trees Enable Structured Programming of Language Model Agents. *CoRR*, abs/2404.07439.
- Langhammer, T.; Hilbert, K.; Praxl, B.; Kirschbaum, C.; Ertle, A.; Asbrand, J.; and Lueken, U. 2021. Mental health trajectories of individuals and families during the COVID-19 pandemic: Study protocol of a longitudinal investigation and indicated prevention program. *Mental Health & Prevention*, 24: 200221.
- Pande, C.; Martin, A.; and Pimmer, C. 2023. Towards Hybrid Dialog Management Strategies for a Health Coach Chatbot. In Martin, A.; Fill, H.; Gerber, A.; Hinkelmann, K.; Lenat, D.; Stolle, R.; and van Harmelen, F., eds., *Proceedings of the AAAI 2023 Spring Symposium on Challenges Requiring the Combination of Machine Learning and Knowledge Engineering (AAAI-MAKE 2023)*, Hyatt Regency, San Francisco Airport, California, USA, March 27-29, 2023, volume 3433 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Stade, E. C.; Stirman, S. W.; Ungar, L. H.; Boland, C. L.; Schwartz, H. A.; Yaden, D. B.; Sedoc, J.; DeRubeis, R. J.; Willer, R.; and Eichstaedt, J. C. 2024. Large language models could change the future of behavioral healthcare: a proposal for responsible development and evaluation. *npj Mental Health Research*, 3(1): 12.
- Wang, H.; Wang, R.; Mi, F.; Deng, Y.; Wang, Z.; Liang, B.; Xu, R.; and Wong, K. 2023. Cue-CoT: Chain-of-thought Prompting for Responding to In-depth Dialogue Questions with LLMs. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, 12047–12064. Association for Computational Linguistics.
- Wang, J.; Lin, D.; and Li, W. 2023. Dialogue Planning via Brownian Bridge Stochastic Process for Goal-directed Proactive Dialogue. In Rogers, A.; Boyd-Graber, J. L.; and Okazaki, N., eds., *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, 370–387. Association for Computational Linguistics.
- Wu, W.; Guo, Z.; Zhou, X.; Wu, H.; Zhang, X.; Lian, R.; and Wang, H. 2019. Proactive Human-Machine Conversation with Explicit Conversation Goals. *CoRR*, abs/1906.05572.
- Yang, Z.; Wang, B.; Zhou, J.; Tan, Y.; Zhao, D.; Huang, K.; He, R.; and Hou, Y. 2022. TopKG: Target-oriented Dialog via Global Planning on Knowledge Graph. In Calzolari, N.; Huang, C.; Kim, H.; Pustejovsky, J.; Wanner, L.; Choi, K.; Ryu, P.; Chen, H.; Donatelli, L.; Ji, H.; Kurohashi, S.; Paggio, P.; Xue, N.; Kim, S.; Hahm, Y.; He, Z.; Lee, T. K.; Santus, E.; Bond, F.; and Na, S., eds., *Proceedings of the 29th International Conference on Computational Linguistics, COLING 2022, Gyeongju, Republic of Korea, October 12-17, 2022*, 745–755. International Committee on Computational Linguistics.
- Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T. L.; Cao, Y.; and Narasimhan, K. 2023. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. *CoRR*, abs/2305.10601.
- Zhang, Q.; Naradowsky, J.; and Miyao, Y. 2023. Ask an Expert: Leveraging Language Models to Improve Strategic Reasoning in Goal-Oriented Dialogue Models. In Rogers, A.; Boyd-Graber, J. L.; and Okazaki, N., eds., *Findings of*

the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023, 6665–6694. Association for Computational Linguistics.

Zheng, Z.; Liao, L.; Deng, Y.; and Nie, L. 2023. Building Emotional Support Chatbots in the Era of LLMs. *CoRR*, abs/2308.11584.