

# Pathways of Thoughts: Multi-Directional Thinking for Long-form Personalized Question Answering

Alireza Salemi  
University of  
Massachusetts Amherst  
Amherst, MA, USA  
asalemi@cs.umass.edu

Cheng Li  
Google DeepMind  
Mountain View, CA, USA  
chgli@google.com

Mingyang Zhang  
Google DeepMind  
Mountain View, CA, USA  
mingyang@google.com

Qiaozhu Mei  
University of Michigan  
Ann Arbor, MI, USA  
qmei@umich.edu

Zhuowan Li  
Google DeepMind  
Mountain View, CA, USA  
zhuowan@google.com

Spurthi Amba  
Hombaiah  
Google DeepMind  
Mountain View, CA, USA  
spurthiah@google.com

Weize Kong  
Google DeepMind  
Mountain View, CA, USA  
weize@google.com

Tao Chen  
Google DeepMind  
Mountain View, CA, USA  
taochen@google.com

Hamed Zamani  
University of  
Massachusetts Amherst  
Amherst, MA, USA  
zamani@cs.umass.edu

Michael Bendersky  
Google DeepMind  
Mountain View, CA, USA  
bemike@google.com

## Abstract

Personalization is well studied in search and recommendation, but personalized question answering remains underexplored due to challenges in inferring preferences from long, noisy, implicit contexts and generating responses that are both accurate and aligned with user expectations. To address this, we propose *Pathways of Thoughts (PoT)*, an inference-stage method that applies to any large language model (LLM) without task-specific fine-tuning. PoT models the thinking as an iterative decision process, where the model dynamically selects among cognitive operations such as reasoning, revision, personalization, and clarification. This enables exploration of multiple reasoning trajectories, producing diverse candidate responses that capture different perspectives. PoT then aggregates and reweights these candidates according to inferred user preferences, yielding a final personalized response that benefits from the complementary strengths of diverse reasoning paths. Experiments on the LaMP-QA benchmark show that PoT consistently outperforms competitive baselines, achieving up to a 10.8% relative improvement. Human evaluation further validates these improvements, with annotators preferring PoT in 66% of cases compared to the best-performing baseline and reporting ties in 15% of cases.

## CCS Concepts

• **Information systems** → **Question answering; Personalization; Computing methodologies** → **Markov decision processes; Natural language generation.**

## Keywords

Personalization, Test-time compute scaling, Thinking LLMs

### ACM Reference Format:

Alireza Salemi, Cheng Li, Mingyang Zhang, Qiaozhu Mei, Zhuowan Li, Spurthi Amba Hombaiah, Weize Kong, Tao Chen, Hamed Zamani, and Michael Bendersky. 2026. Pathways of Thoughts: Multi-Directional Thinking for Long-form Personalized Question Answering. In *Proceedings of the ACM Web Conference 2026 (WWW '26)*, April 13–17, 2026, Dubai, United Arab Emirates. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3774904.3792145>

## 1 Introduction

Personalizing large language models (LLMs) has become increasingly important due to applications in content generation, writing assistance, education, and recommendation [16, 18, 22, 23, 29]. By conditioning outputs on user preferences, context, and needs, personalization improves engagement, satisfaction, and overall effectiveness [28]. It is particularly valuable in information-seeking systems such as search engines and question answering (QA) [31, 32], where appropriate responses depend on user-specific factors like background knowledge and learning style. Despite its clear relevance to web search and emerging systems (e.g., AI-powered search modes), personalization in QA remains relatively underexplored.

Augmenting LLM inputs with the user’s personalized context have proven effective for personalized generation [16, 18, 27, 29]. This can be further enhanced by training the LLM to perform reasoning over the personalized context to extract the user’s preferences before response generation [28]. However, in most real-world cases, fine-tuning the LLM is often infeasible for two key reasons. On one hand, many LLMs are black boxes with frozen parameters, offering no control over their internal mechanisms. On the other hand, even with full access to parameters, the fine-tuning process requires substantial amount of personal data and computational resources for individual users. Consequently, a practical way to influence the



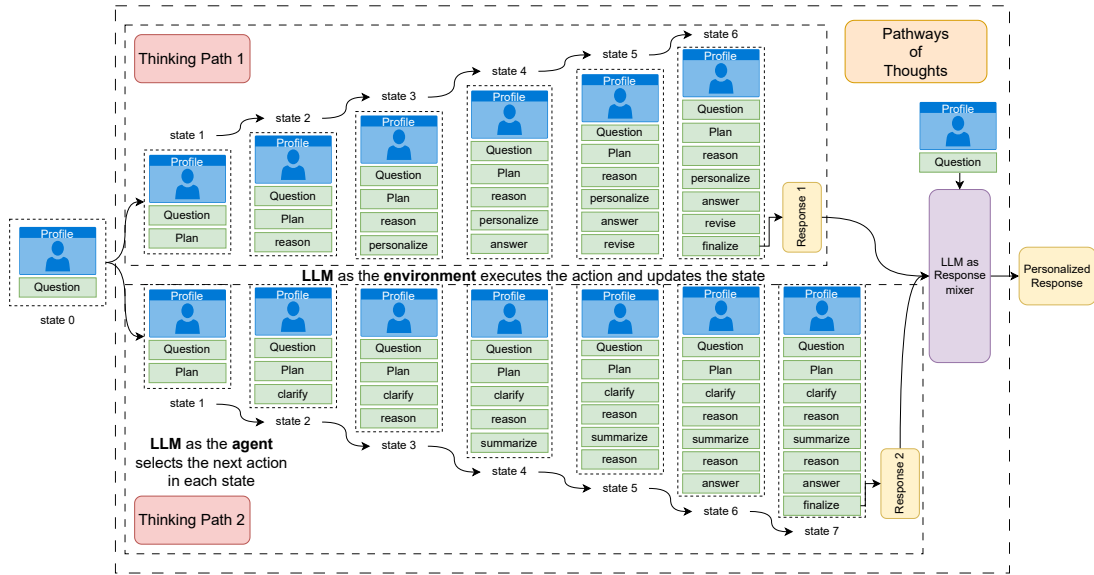
This work is licensed under a Creative Commons Attribution International 4.0 License.

WWW '26, Dubai, United Arab Emirates

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2307-0/2026/04

<https://doi.org/10.1145/3774904.3792145>



**Figure 1: Pathways of Thoughts overview.** PoT explores multiple directions of thinking about answering the question, generates a response for each, and mix them into a final response.

model’s output is by applying advanced strategies during inference. Recognizing this, our paper focuses on this scenario and proposes an **inference-stage strategy** to achieve better personalization.

Developing effective personalization models presents its own unique challenges. First, the quality of a response in personalized QA is determined by the user who asks the question, which is influenced by their preferences, context, and background. Therefore, even when two answers are factually accurate, one may be preferred by the user. Second, user preferences are often not explicitly stated and must be inferred from personalized context, such as profiles or prior interactions. This context is frequently lengthy and noisy, making it difficult for a model to extract relevant signals. These issues highlight that simply augmenting an LLM’s input with user context is often insufficient. Effective personalization requires the model to perform reasoning over noisy, implicit user data to deduce user preferences before generating a response.

To address the aforementioned challenges of personalized QA, this paper introduces *Pathways of Thoughts (PoT)*, demonstrated in Figure 1. This approach, which is generic and can be applied to any LLM without additional training, enables the model to explore multiple thinking pathways when generating responses, each leading to a distinct response. We formalize each thinking pathway as a Markov Decision Process (MDP), where the LLM serves as both the agent and the environment. At each step, the LLM, acting as the agent, selects the most suitable action from a predefined set of cognitive operations, such as planning, reasoning, personalizing, and revising. The LLM then assumes the role of the environment, executing the chosen action and updating the state with newly acquired knowledge. This iterative process allows the model to refine its understanding of the user’s intent, extract key preferences from the user profile, and explore diverse strategies for answering the question. The process continues until the model determines that no further actions are needed, producing an initial set of responses. To generate the final personalized response, the LLM mixes outputs

from different thinking pathways into a single response, aligning them with the user’s preferences extracted from their profile. Pathways of Thoughts addresses the challenges of LLMs in personalized question answering by allowing the model to iteratively think about the next step in responding to the question and how the lengthy user context can be effectively utilized. Additionally, it explores different diverse directions and paths, considering a range of possible solutions for answering the question. Finally, generating a response by aggregating all pathways’ responses ensures that the output captures the strengths of all explored pathways, resulting in a more personalized and aligned response.

We conduct our experiments on the *LaMP-QA* benchmark [31]—a recent large-scale benchmark designed for evaluating personalized QA in three diverse domains: (1) Art & Entertainment, (2) Lifestyle & Personal Development, and (3) Society & Culture. Our results demonstrate that PoT outperforms competitive baselines across all categories, with up to 10.8% statistically significant improvement on average compared to the best-performing baselines. Additionally, we investigate the impact of pathway length, the number of pathways, and various methods for diversifying pathways and aggregating responses on performance. We show that PoT generalizes to different long-context backbone LLMs without requiring additional training. Finally, human evaluation of comparing outputs side by side shows that *PoT* better satisfies users’ information needs in 66% of cases, compared to 19% for the best-performing baseline.

## 2 Related Work

**LLM Personalization.** Personalization plays a central role in search, recommendation, and many text generation applications [6, 23, 29, 39]. To enable personalization in LLMs, Salemi et al. [29] introduced a retrieval-augmented generation framework alongside the LaMP benchmark for short-form text, which was later extended to long-form content generation with LongLaMP [16] and more recently to personalized QA with LaMP-QA [31]. Personalized assistants have

also been studied in several recent works [18, 21, 22, 43]. Proposed approaches span a variety of techniques, including training retrieval models with user feedback [27], optimizing LLMs with personalized supervision [13], and generating personalized prompts [17]. Parameter-efficient fine-tuning has emerged as another direction [35], including its integration with retrieval-augmented methods [30]. Beyond model adaptation, reasoning and self-training strategies have been shown to improve personalization in text generation [28]. More recently, learning from natural language feedback has proven effective for adapting QA systems to individual users [32]. Nevertheless, strategies for improving personalized QA at test time—without incurring the cost of additional training—remain underexplored, which is the focus of this work.

**Reasoners & Agents for Problem Solving.** Reasoning enables models to solve problems with step-by-step processing, known as chain-of-thought (CoT), enhancing performance in complex tasks like mathematical problem-solving, logical reasoning, and commonsense understanding [19, 37, 42]. Training LLMs to perform intermediate reasoning steps before generating responses has been effective, as seen in GPT-o1 [25] and DeepSeek-R1 [5]. Additionally, agentic and multi-turn or multi-agent systems that interact with external tools [14], code [8, 20, 38], simulated environments [10, 26, 36, 41], or with other agents [9] have been explored for complex problem-solving. Markov Decision Processes (MDP) have shown promise in reasoning for games and mathematical problems [10, 36]. However, these methods have not been applied to free-form text generation, particularly personalization which requires user-specific responses rather than objectively correct answers. This study extends these methods to personalized text generation by integrating them with graph-based search in place of traditional tree search in MDPs, enabling the model to produce responses that reflect the aggregated quality of all explored candidate solutions rather than relying on a single path.

**Scaling Inference Compute.** Recent advances in LLM reasoning for logic and mathematics show that increasing compute resources during inference can significantly enhance performance [4, 33]. This allows LLMs to explore answer spaces more thoroughly, improving accuracy in logical reasoning, math problem-solving, and code generation [2, 3, 44]. While previous research has focused on math, coding, and logic, its potential for free-form text generation, especially personalized question answering, is underexplored. Our work addresses this by using increased inference compute to better navigate response spaces, enhancing personalization.

### 3 Problem Formulation

We assume a question  $x_u$  originates from a user  $u$ , who has a set of  $n_u$  personalized information elements, denoted as  $P_u = \{p_i\}_{i=1}^{n_u}$ . Following previous work that utilizes long-term user history as the user profile [15, 29, 31],  $P_u$  consists of the user’s previously asked questions along with their question narrative. The objective is to use this personalized information, along with the question  $x_u$ , to generate a personalized response using an LLM  $\pi$ , formulated as  $\hat{y}_u = \pi(x_u, P_u)$ . To evaluate the personalized response generated for the user  $u$ ’s question, following Salemi and Zamani [31], we assume access to a set of aspects  $E_{x_u} = \{e_i\}_{i=1}^{|E_{x_u}|}$  that user  $u$  expects

to be addressed in the response, extracted from a question narrative  $r_{x_u}$  written by the user. These aspects are used solely for evaluation and are not accessible to the model during response generation. Finally, a metric  $\mu(x_u, \hat{y}_u, E_{x_u}, r_u)$  is used to score the generated response based on the extent to which these aspects are covered.

## 4 Pathways of Thoughts

Personalization is inherently user-specific, meaning that a response generation strategy effective for one user may not perform well for another. Consequently, relying on a single solution to generate a response for a user may not always yield the most preferred outcome for each individual user. To address this limitation, as well as those discussed in Section 1, we propose *Pathways of Thoughts (PoT)*, illustrated in Figure 1. In this framework, the LLM explores multiple thinking pathways, each corresponding to a distinct line of thought that leads to a potential response. We model each thinking pathway as an MDP, where the LLM serves as both the agent and the environment. At each step, the LLM (acting as the agent) selects an action from a set of fundamental cognitive operations such as planning, reasoning, clarification, or revising. The same LLM (now acting as the environment) then executes the chosen action, updating the current state with newly acquired information about the question and user. This iterative process continues until convergence, producing a final response for each pathway. To generate a single personalized response, the LLM aggregates the responses obtained from different pathways, integrating their complementary strengths while aligning with the user’s preferences inferred from their profile. The complete procedure is detailed in Algorithm 1, and the next section elaborates on this method.

### 4.1 Thinking as a Markov Decision Process

The thinking process in the human brain involves using and integrating a set of fundamental cognitive skills—e.g., reasoning, planning, and revision—that collectively form the foundation of structured thought [7, 34]. Inspired by theories about sequential decision making in the human brain, we model the thinking process of LLMs as an MDP. Given the current thoughts and steps taken to generate a response, the model selects the next fundamental action from a predefined set to enhance its representation of the question, refine its reasoning, and improve the generated response.

An MDP is defined as a quadruple  $(S, A_{s_t}, P_{a_t}, R_{a_t})$ , where  $S$  is the set of states,  $A_{s_t}$  represents the set of available actions in state  $s_t$ ,  $P_{a_t}(s_{t+1} | s_t)$  defines the probability of transitioning from state  $s_t$  to  $s_{t+1}$  given action  $a_t$ , and  $R_{a_t}(s_{t+1}, s_t)$  specifies the reward function, which assigns a reward for taking action  $a_t$  in state  $s_t$  and transitioning to  $s_{t+1}$ . In this setting, the agent  $\pi$  aims to assign a probability  $p_\pi(a_t | s_t)$  to each action based on the current state  $s_t$ .

To generate a personalized response for the question  $x_u$  from the user  $u$ , we assume the agent  $\pi$  starts from an initial state  $s_0$ , which consists of the question  $q_u$  and the user profile  $P_u$ . The agent iteratively selects the most probable action at each step using the prompt shown in Figure 7 in Appendix A, continuing this process until it either chooses a halting action (as determined by the LLM) or reaches the maximum number of allowed actions, denoted as  $T$ . Since our task involves text generation and lacks an external environment enforcing termination, the agent itself determines when

```


Initial State & Action Definition Prompt


Your task is to answer a user's question from a post in a personalized way by considering the user's past post questions and detailed descriptions of these questions. This is a multi turn task and you don't need to answer the question in one turn. In each turn, you can perform one action in that turn.
# Your context information:
- The user's current question from a post.
- The user's past post questions and detailed descriptions of these questions.
# Actions:
## answer:
### Definition: This action asks you to generate a personalized response to the question, if the question can benefit from personalization. Therefore, you should choose this action when you are ready to answer the question. This action should be executed when there is no ambiguity and all the information needed for answering the question is available.
### Output format:
You should generate a valid JSON, enclosed in ```json ``` block, with the following fields:
- personalizedAnswer: a string representing your personalized answer to the user's current question considering the personalized context and previous taken steps.
## check_personalization:
### Definition: This action ask you to see if the question asked by the user can benefit from personalization or not. This action can be useful when you want to decide if you need to personalize the response or not.
### Output format:
You should generate a valid JSON, enclosed in ```json ``` block, with the following fields:
- needs_personalization: a boolean that is true if the question can benefit from personalization and false if personalization cannot help.
- reason: a string representing the reason why the question can or cannot benefit from personalization.
## personalize
### Definition: This action asks you to find the relevant information about the question from the context and summarize how it can be helpful for answering the question. This action should be executed when you think personalized information is needed to answer the question.
### Output format:
You should generate a valid JSON, enclosed in ```json ``` block, with the following fields:
- personalizedSummary: a string representing your summary of user context and how it is helpful to answer the question.
## plan
### Definition: This action asks you to generate a plan of steps that we can take to answer all the aspects of the question. This action should be executed when you think the the question has multiple aspects and planning future steps that you need to take to answer the question is necessary.
### Output format:
You should generate a valid JSON, enclosed in ```json ``` block, with the following fields:
- plan: a string representing your plan of steps that we can take to answer the question.
## reason
### Definition: This action asks you to think step by step about one aspect of the question. This action should be executed when you know different aspects of the question and you want to think step by step about one aspect of the question and how it can be answered.
### Output format:
You should generate a valid JSON, enclosed in ```json ``` block, with the following fields:
- aspect: a string representing the aspect of the question.
- reason: a string representing your reasoning of the question aspect and how it can be answered.
## clarify
### Definition: This action asks you to check if the question is ambiguous and how it can be clarified. This action should be executed when you think the current information needs more clarifying.
### Output format:
You should generate a valid JSON, enclosed in ```json ``` block, with the following fields:
- ambiguous: a boolean indicating whether the question is ambiguous or not.
- clarification: a string representing your clarification questions of the current information and steps.
- clarification_response: a string representing your response to the clarification question. You should generate a reasonable response to clarification questions. This field should be empty if the question is not ambiguous.
## summarize
### Definition: This action summarizes all available information and findings about the question. This action should be executed when you think the current information needs more summarization.
### Output format:
You should generate a valid JSON, enclosed in ```json ``` block, with the following fields:
- summarization: a string representing your summarization of the current information and steps.
## revise
### Definition: This action asks you to revise the current response to the question. This action should be executed when you provided an answer to the question and you want to revise the current response.
### Output format:
You should generate a valid JSON, enclosed in ```json ``` block, with the following fields:
- revised: a string representing your revised response to the question.
## finalize
### Definition: This action will finalize the response to the question. This action should be executed when you think the response is finalized and no more action is required to answer or revise the answer to the question.
### Output format:
You should generate a valid JSON, enclosed in ```json ``` block, with the following fields:
- personalizedAnswer: a string representing your finalized response to the question.
# Your task:
This is a multi turn task and you don't need to answer the question in one turn. In each turn, you can perform multiple actions from an action list. Then, based on the detention of the action, you can do that action in that turn and provide the response. Therefore, your task in each turn consists of two steps:
## Step 1: Selecting one of the actions that the user provides for you as the possible next actions. In this case, you should choose one of the actions considering that the main goal is generating a personalized response to the question. The action you choose should be the next best step toward the final goal, which is generating a personalized response to the question.
## Step 2: In this step, based on the action definition, you can do that action in the same turn.
### Output format:
You should generate a valid JSON, enclosed in ```json ``` block, with the following fields:
- action: a string representing your selected action. This text should be one of the valid actions that user provides.
- reason: a string representing the reason for why the action is selected and how it is helpful to reach the final goal, which is generating a personalized response to the question.
- actionOutput: a json representing the output of the action based on the defined output format for the selected action.
# Your input:
## Past post questions and detailed descriptions of these questions:
(profile)
## Current post question:
(query)

```

Figure 2: Initial state and actions definition prompt used in Pathways of Thoughts.

to stop or continues until the predefined action limit is reached. See = Algorithm 1 (lines 3-17) for details. The following subsections outlines the definitions of each MDP component in our framework:

**Set of States ( $S$ ):** The state is defined as the concatenation of the inputs, outputs, and decisions of the model into a sequence of tokens using a predefined prompt format. This representation captures the full context of the system like a conversation. The initial state of the system consists of the user's question  $x_u$ , the user profile  $P_u$ , and a set of instructions specifying the actions and the model's expected behavior for each action, as shown in Figure 2. Given this setup,

the set of all possible states corresponds to the set of all textual sequences that this initial prompt is their prefix. As a result, the state space forms an uncountable discrete set.

**Set of Possible Actions for State  $s_t$  ( $A_{s_t}$ ):** We consider the following fundamental actions that the agent can choose from in thinking process: *answering*, *planning*, *personalization*, *personalizing*, *reasoning*, *clarifying*, *summarizing*, *revising*, and *finalizing*.<sup>1</sup> The detailed definition of each action and output format of them is provided in Appendix A.1. In the initial state, the agent is restricted to

<sup>1</sup>This action, *finalizing*, also referred to as halting, terminates the process.

**Algorithm 1** Pathways of Thoughts Implementation using MDP and Global Search.

---

**Input:** prompt  $x_u$ , profile  $P_u$ , max path length  $T$ , number of pathways  $N$ , backbone LLM  $\pi$

**Output:** personalized output  $\hat{y}_{x_u}$

```

1:  $R_{PoT} = \{\}$   $\triangleright$  Set of all responses from different pathways
2: for  $i = 0$  to  $N$  do  $\triangleright$  Generate  $N$  pathways of thoughts
3:    $t = 0$ 
4:    $s_t = \text{init\_prompt}(x_u, P_u)$   $\triangleright$  Prompt in Figure 2
5:   while  $t < T$  do  $\triangleright$  Thinking for maximum  $T$  steps
6:      $A_{s_t} = \text{possible\_actions}(s_t, t)$   $\triangleright$  Find the possible
       actions for state  $s_t$ .
7:      $\bar{a}_p = \text{action\_selection\_prompt}(s_t, A_{s_t})$   $\triangleright$  Prompt in
       Figure 7
8:      $a_t = \pi(s_t; \bar{a}_p)$   $\triangleright$  Selecting the action (LLM as agent)
9:      $\hat{y}_t = \pi(s_t; \bar{a}_p; a_t)$   $\triangleright$  Performing the action (LLM as
       environment)
10:     $s_{t+1} = s_t; \bar{a}_p; a_t; \hat{y}_t$   $\triangleright$  new state by concatenating
       input/outputs in this time step
11:    if  $a_t = \text{finalize}$  then
12:       $R_{PoT} = R_{PoT} \cup \{\hat{y}_t\}$   $\triangleright$  Capturing the final response
       of this pathway
13:      break
14:    end if
15:     $t = t + 1$   $\triangleright$  updating the step variable
16:  end while
17: end for
18:  $\bar{I}_p = \text{important\_aspects\_prompt}(P_u)$   $\triangleright$  Prompt in Figure 8
19:  $I_u = \pi(\bar{I}_p)$   $\triangleright$  Extracting important aspects from user profile
20:  $\bar{y}_p = \text{mixture\_prompt}(I_u, x_u, R_{PoT})$   $\triangleright$  Prompt in Figure 10
21:  $\hat{y}_{x_u} = \pi(\bar{y}_p)$   $\triangleright$  Combining responses based on important
       aspects
22: return  $\hat{y}_{x_u}$   $\triangleright$  Returning final response

```

---

selecting the *planning* action, as it must first plan the steps required to address the user’s question. In subsequent states, the agent can choose from all the actions; however, the model can only select the *revising* action if it has previously selected the *answering* action, as revision typically follows the generation of a response.

**Transitioning from State  $s_t$  to  $s_{t+1}$  with Action  $a_t$  ( $P_{a_t}(s_{t+1} | s_t)$ ):** We use the LLM as both the agent and the environment simultaneously: when it selects an action, it acts as the agent, and when it executes the action, it acts as the environment. This dual role allows the LLM to control the flow of thinking and performing it simultaneously. Accordingly, with the agent decision to perform action  $a_t$  at time step  $t$ , the LLM receives the state  $s_t$  and action  $a_t$  as input and generates the output  $\hat{y}_t$ . The concatenation of  $s_t$ ,  $a_t$ , and  $\hat{y}_t$  forms the next state  $s_{t+1}$ . Thus, the transition probability  $P_{a_t}(s_{t+1} | s_t)$  is defined by the LLM itself.

**Reward Function ( $R_a(s_{t+1}, s_t)$ ):** Since our agent is used without training, an explicit reward function is not required; we assume the agent makes reasonable choices. However, in general, an outcome reward function can be employed to assign rewards to actions based on the evaluation of the LLM’s final response to the question.

## 4.2 Multi-Directional Thinking by Diverse Planning and Mixture-of-Responses

We observed that when the agent selects *planning* as the first action, the subsequent steps consistently follow the outline of the generated plan. This suggests that encouraging multi-directional thinking about responding to the question can be achieved by generating a variety of plans, each exploring a different strategy for addressing the question. We study two methods for diverse planning:

- **Initial State Alteration:** If the initial state is modified before executing the planning action, the LLM can produce different plans, as the changes in personalized context influences the planning process. This can be modeled by varying the amount of past user interactions accessible to the agent, effectively addressing the challenge of long personal context. To implement this, we randomly select  $N$  subsets of the user profile  $P_u$ , each of size  $\tau \times |P_u|$ , where  $0 \leq \tau \leq 1$ . Starting from these  $N$  diverse initial states lead to different planning trajectories and thinking pathways.
- **Planning Action Variation:** Another way to generate different plans from the same initial state is to use a higher sampling temperature  $0 \leq \tau \leq 1$  during plan generation, which increases the randomness in the generated plans, leading to more diverse outputs. Thus, we generate  $N$  different plans from the same initial state (created using the whole profile  $P_u$ ) by setting a high temperature  $\tau$ , leading to different thinking pathways.

Each thinking pathway results in a response, and together they form a set of response proposals  $R_{PoT} = \{r_i\}_{i=1}^N$ . To produce the final response  $\hat{y}_{x_u}$  for the user, we need to ensure that it is the most suitable one for the user. To achieve this, we first use the LLM  $\pi$  to extract the important aspects  $I_u$  for user  $u$  from their user profile  $P_u$  using the prompt shown in Figure 8 in Appendix A.2. This serves as an intermediate step to learn about the user’s individual preferences. Based on this, there are two methods to generate the final response:

- **Best-of-N:** Given the extracted important aspects  $I_u$  for the user  $u$  and the input  $x_u$ , the LLM  $\pi$  selects the best response from the set of responses resulted from multi-directional thinking  $R_{PoT}$ , using the prompt shown in Figure 9 in Appendix A.2.
- **Mixture-of-N:** Given the extracted aspects  $I_u$  for the user  $u$  and the input  $x_u$ , the LLM  $\pi$  combines different aspects of the generated responses from the set of multi-directional thinking outputs  $R_{PoT}$ . The final response is denoted as  $\hat{y}_{x_u} = \pi(x_u, R_{PoT}, I_u)$ , using the prompt shown in Figure 10 in Appendix A.2. This mimics how humans solve problems by exploring multiple solutions and synthesizing a new solution that incorporates the strengths of them. This is beneficial for open-ended text generation, where there is no definitive right or wrong response. By aggregating the strengths of all responses, it generates a response that aligns with user preferences, ensuring a more personalized output. PoT uses this approach as the primary response aggregation method.

## 5 Experiments

### 5.1 Experimental Setup

**Datasets.** We conduct experiments on the LaMP-QA benchmark [31], the only publicly available dataset for long-form personalized QA. LaMP-QA covers three diverse domains: (1) Art & Entertainment (767 questions), (2) Lifestyle & Personal Development (989

**Table 1: Performance on LaMP-QA. † and ‡ show statistically significant improvement over the single and multiple inference best baseline, respectively, using t-test ( $p < 0.05$ ). Gemini 1.5 Pro has been used as the backbone LLM for all the methods.**

Method		Arts & Entertainment	Lifestyle & Personal Development	Society & Culture	Average (macro)
Single Inference					
(1)	No Personalization	0.1741	0.2863	0.2996	0.2533
(2)	In-Context Personalization	0.1722	0.3022	0.3482	0.2742
(3)	In-Context Personalization w/ CoT	0.2119	0.3862	0.4230	0.3403
(4)	<b>Pathways of Thoughts (<math>N = 1, T = 8</math>)</b>	<b>0.2499<sup>†</sup></b>	<b>0.3966</b>	<b>0.4490<sup>†</sup></b>	<b>0.3651<sup>†</sup></b>
Multiple Inference					
(5)	In-Context Personalization w/ Best-of-32	0.2248	0.3753	0.4303	0.3434
(6)	In-Context Personalization w/ CoT & Best-of-32	0.2607	0.4380	0.4789	0.3925
(7)	In-Context Personalization w/ Mix-of-32	0.2310	0.3781	0.4411	0.3500
(8)	In-Context Personalization w/ CoT & Mix-of-32	0.2689	0.4443	0.4891	0.4007
(9)	Tree of Thoughts ( $K = 32, B = 2$ )	0.2054	0.4178	0.4546	0.3592
(10)	<b>Pathways of Thoughts (<math>N = 16, T = 8</math>)</b>	<b>0.2999<sup>†‡</sup></b>	<b>0.4960<sup>†‡</sup></b>	<b>0.5362<sup>†‡</sup></b>	<b>0.4440<sup>†‡</sup></b>

questions), and (3) Society & Culture (1074 questions). Each LaMP-QA example contains a user question, the user’s question history (serving as the user profile), a user-written narrative that reflects the user’s information need and intent, and a set of personalized rubrics specifying the aspects an ideal response should address. For evaluation, we follow the evaluation recipe of Salemi and Zamani [31] with Gemini 1.5 Pro as the evaluator language model. For each question, the evaluator checks whether the generated response addresses each personalized aspect, assigning scores in the range  $[0, 2]$ . These are normalized to  $[0, 1]$ , and the final response score is computed as the mean normalized score across all aspects. For details, we refer the reader to Salemi and Zamani [31]. In one experiment, we corroborate our findings through a side-by-side human evaluation comparing PoT to the strongest baseline.

**PoT’s Configurations.** We use Gemini 1.5 Pro as the backbone LLM. To show the generalizability of PoT, we also use GPT-4o-mini<sup>2</sup> [24] using our most representative experiment configurations. For single pathways, we use Nucleus sampling [11] with a temperature of  $\tau = 0.1$ . For multi-directional pathways, we use the same setting, except for the *Planning Action Variation*, where a temperature of  $\tau = 0.9$  is used exclusively for the *planning* action. For generating diverse plans, we use *Planning Action Variation* by default and *Mixture-of-N* for response aggregation, unless otherwise specified. We set the context length to  $32k$ , the maximum number of actions per pathway to  $T = 8$ , and the number of pathways to  $N = 16$ . In ablations, we reduce the number of pathways to  $N = 4$  to save costs, unless otherwise specified. For personalized context, we use the 10 recent questions of the user with the same order as they appear in the LaMP-QA [31] dataset.

**Baselines.** We use six baselines: (1) LLM without personalized context that directly answers the question, (2) LLM with access to the same personalized context as PoT, and (3) LLM with access to personalized context that first performs reasoning over the context to extract important aspects relevant to the question before generating the response. The two latter use the prompts shown in Figures 11 and 12 in Appendix B. These baselines all use Nucleus sampling

with a temperature of  $\tau = 0.1$ . We extend the latter two baselines by doing multiple inferences and Best-of-N [3, 12] and Mixture-of-N to evaluate their performance when test-time compute scales, to form baselines (5), (6), (7), and (8). Additionally, we include the (9) Tree of Thoughts (ToT) [40] as a baseline, leveraging its test-time compute scaling for comparison with our method. We follow their setup, particularly in the creative writing task, as it is the closest task they studied to personalized question answering. This setup consists of an intermediate planning step, where a set of plans is generated and the best plan is selected, followed by a final response generation step, where a set of responses is generated and the best response is chosen. For planning, we use the prompts shown in Figure 13, and for response generation, we refer to Figure 14, both in Appendix B. We observed that the baselines generate, on average, half the number of tokens compared to *PoT*. To ensure a fair comparison, we set the Best-of-N and Mixture-of-N baselines to generate double the number of responses, with  $N = 32$  and a sampling temperature of  $\tau = 0.9$  to ensure diverse responses. For the Best-of-N and Mixture-of-N approach, we use the method introduced in Section 4.2. For ToT, following Yao et al. [40] and their setup for the creative writing task that is the most similar to ours, we set the tree depth to  $B = 2$ . For fair comparison with PoT, we generate  $N = 32$  plans and responses at each node with a sampling temperature of  $\tau = 0.9$ . The prompts for selecting the best plan is shown in Figure 15 in Appendix B. For selecting the best response, we utilize the method introduced in Section 4.2.

## 5.2 Empirical Findings

**Comparison with baselines.** The results of this experiment are reported in Table 1. The findings show that PoT outperforms all the baselines across all categories, achieving a **10.8%** relative improvement on average over the best baseline (row 8). The improvements are statistically significant with respect to t-test with 95% confidence. This demonstrates the superiority of PoT compared to the baselines. To compare our thinking-as-an-MDP method with baselines, we also report the results using only a single inference (rows 1-4). The results from this experiment show that our method (row

<sup>2</sup>Available at: <https://platform.openai.com/docs/models/gpt-4o-mini>

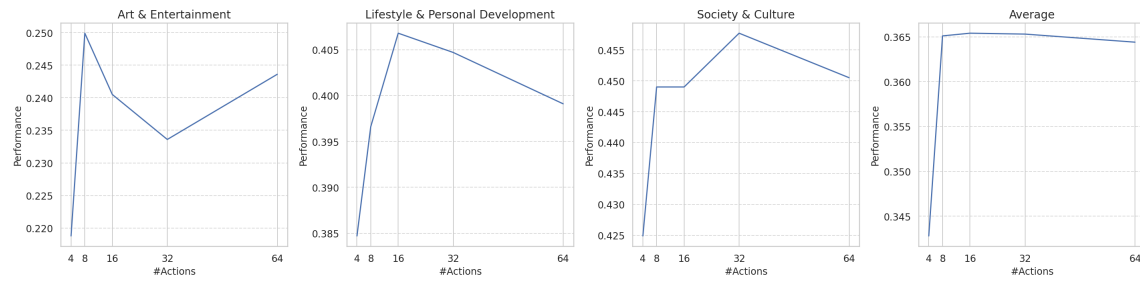


Figure 3: Effect of the maximum number of actions per pathway on the performance.

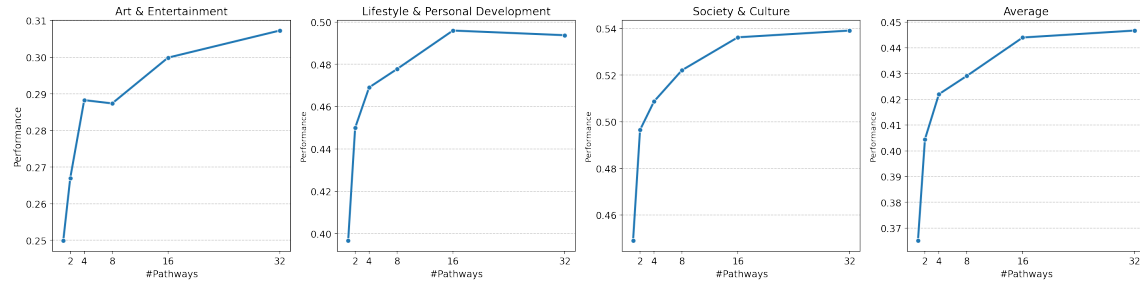


Figure 4: Effect of the number of pathways on the performance.

4) outperforms all single inference baselines (rows 1-3), with statistically significant improvements in two out of three categories and overall performance. Specifically, our method achieves a 7.2% relative improvement in average performance, suggesting that our thinking approach is significantly more effective than CoT.

**Effect of the maximum number of actions in PoT.** For this, we focus on a single inference ( $N = 1$ ) using PoT and vary the maximum number of actions it can perform. The performance across all categories is shown in Figure 3. The findings show that increasing the number of actions improves performance up to a certain point (8 actions per pathway), after which performance declines. This occurs because, with too few actions, the LLM cannot construct effective pathways to solve the problem. Conversely, when given too many actions, the LLM may overthink the problem, performing unnecessary actions that degrade performance [1].

**Effect of the number of pathways in PoT.** We set the maximum number of actions per pathway to  $T = 8$  and vary the number of pathways. The performance across all categories is shown in Figure 4. The findings indicate that increasing the number of pathways consistently improves performance; however, as the number grows, the improvements become thinner. This occurs because there is a limited set of different ways for answering the question, and after a certain point, the solutions start to repeat, leading to diminishing returns in performance with additional pathways.

**Effect of different plan diversification and response aggregation in PoT.** We evaluate all possible combinations of the methods introduced in Section 4.2. Specifically, we consider four configurations: 1) *Planning Action Variation w/ Best-of-N*, 2) *Planning Action Variation w/ Mixture-of-N*, 3) *Initial State Alteration w/ Best-of-N*, and 4) *Initial State Alteration w/ Mixture-of-N*. Additionally, we explore different values of the hyperparameter  $\tau$  for both *Planning Action Variation* and *Initial State Alteration*. The results of this experiment are shown in Figure 5. Our analysis reveals several key

observations. Most notably, *Mixture-of-N* consistently outperforms *Best-of-N* regardless of the plan diversification method. This can be attributed to the fact that *Best-of-N* selects only one response from the generated candidates, whereas *Mixture-of-N* enables the combination of strengths from multiple responses, resulting in a more comprehensive and higher-quality output. We observe that increasing the hyperparameter  $\tau$  improves performance. In *Planning Action Variation*, a higher  $\tau$  leads to a more diverse set of generated plans due to higher sampling temperature, resulting in more varied thinking pathways that may produce more tailored responses. In *Initial State Alteration*, increasing  $\tau$  provides each pathway with more user-specific information, enabling more personalized responses. Comparing the two, *Planning Action Variation* generally outperforms *Initial State Alteration* in terms of average performance. However, as  $\tau$  increases, the performance gap narrows. Overall, *Planning Action Variation w/ Mixture-of-N* achieves the best performance for personalized question answering.

**PoT's generalizability and independence of LLM.** In the same setting as Gemini 1.5 pro, we apply PoT and the best performing baselines to another backbone LLM, GPT-4o-mini, to demonstrate its effectiveness across different LLMs with long context windows. The average performance across all categories is presented in Figure 6. The results indicate that, similar to Gemini 1.5 Pro, GPT-4o-mini with PoT outperforms all baselines on all categories. It achieves a statistically significant 9.4% relative improvement on average across all datasets compared to the best baseline. These findings suggest that PoT enhances the performance of LLMs in personalized QA regardless of the underlying backbone LLM.

**Human evaluation.** We randomly sample 100 examples from LaMP-QA and generate outputs using PoT and In-Context Personalization w/ CoT & Best-of-32. Human annotators evaluate the outputs based on their alignment with the user's question narrative (i.e., a detailed description of user's information need and

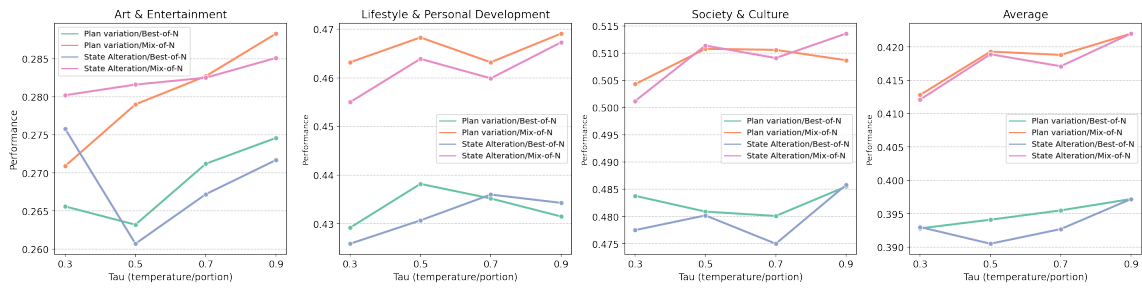


Figure 5: Effect of planning diversification and response aggregation methods on the performance.

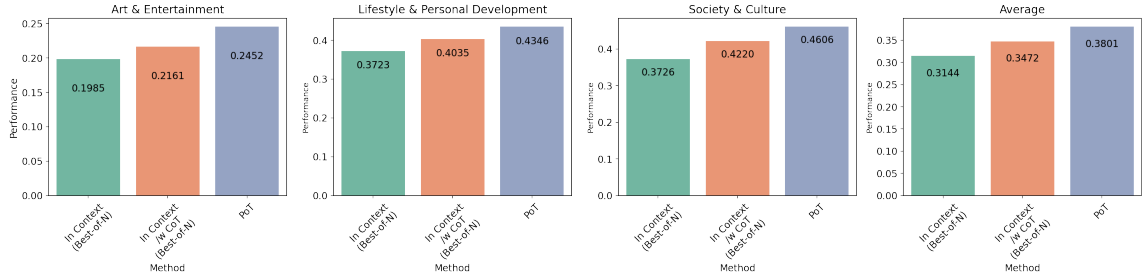


Figure 6: Performance of Pathways of Thoughts and baselines with GPT-4o-mini as the backbone LLM.

perspective). Each example is annotated twice, yielding an inter-annotator agreement Cohen’s kappa of 0.613. The results indicate that annotators preferred PoT in 66% of the cases, the baseline in 19%, and reported tie in 15% of cases. These results show that PoT achieves considerably higher alignment with human preferences.

### 5.3 Case Study

We observe that PoT ( $N = 16, T = 8$ ) generates 1943 unique pathways, defined by distinct action sequences chosen. Specifically, *Art & Entertainment* yields 1109, *Lifestyle & Personal Development* produces 814, and *Society & Culture* results in 1182 unique pathways. On average, each example generates 11.1 unique pathways with on average 7.19 actions. As a case study demonstrating the effectiveness of PoT, we present an example in Figure 16 in Appendix:

- Thinking Pathways & Generated Responses:** To conserve space, we present in Figure 16 only the sequence of steps taken in each pathway, rather than the full generated text, and the response in each pathway. Each pathway corresponds to a distinct reasoning trajectory generated by PoT. For example, in the illustrated case, PoT selects different pathways depending on whether personalization is required. Some pathways generate a direct response without personalization (e.g., Pathway 7), while others involve multiple rounds of reasoning that integrate various aspects of the query with the user profile before producing an answer (Pathways 1–6, 9–11, 13–15). Certain pathways also involve revision (Pathways 12, 15, and 16) or clarification (Pathway 4). Notably, we observed that for some trajectories PoT omit personalization entirely (Paths 7, 8, and 11). This shows that PoT deliberately explores a diverse range of reasoning strategies, producing responses from multiple perspectives that are subsequently combined into a final answer encompassing the advantages of each. Although responses across pathways share

common elements, each retains distinctive characteristics, reflecting the variation in underlying reasoning strategies.

- Final Response:** Figure 16 shows the final response to the user’s question, produced by aggregating and integrating components from multiple pathway outputs into a single coherent answer. The contributions of different pathways are indicated in brackets ([ ]). This shows how PoT leverages the complementary strengths of diverse reasonings, combining them into a unified response that aligns with the user’s preferences and key information needs.

## 6 Conclusions

This paper addresses the critical yet under-explored problem of personalized question answering. we propose *Pathways of Thoughts (PoT)*, a novel approach that models the LLM’s thinking process as an MDP during test time. By enabling the exploration of multiple thinking pathways and aggregating diverse responses, *PoT* effectively captures user preferences to generate personalized answers. Our results on *LaMP-QA* demonstrate that *PoT* achieves substantial improvements over baseline methods, with performance gains of up to 10.8%. Furthermore, human evaluation confirms the effectiveness of our approach, showing a strong user preference over existing baselines. This paper focuses solely on personalized question answering; however Pathways of Thoughts is a general framework and future work can explore the potential of Pathways of Thoughts for long-form text generation beyond personalization.

## Acknowledgement

This work was supported in part by the Center for Intelligent Information Retrieval, in part by NSF grant #2143434, in part by the Office of Naval Research contract #N000142412612, and with support from by Google.org. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

## References

- [1] Pranjal Aggarwal, Seungone Kim, Jack Lanchantin, Sean Welleck, Jason Weston, Ilya Kulikov, and Swarnadeep Saha. 2025. OptimalThinkingBench: Evaluating Over and Underthinking in LLMs. arXiv:2508.13141 [cs.CL] <https://arxiv.org/abs/2508.13141>
- [2] Zhenni Bi, Kai Han, Chuanjian Liu, Yehui Tang, and Yunhe Wang. 2024. Forest-of-Thought: Scaling Test-Time Compute for Enhancing LLM Reasoning. arXiv:2412.09078 [cs.CL]
- [3] Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V. Le, Christopher Ré, and Azalia Mirhoseini. 2024. Large Language Monkeys: Scaling Inference Compute with Repeated Sampling. arXiv:2407.21787 [cs.LG]
- [4] Yanxi Chen, Xuchen Pan, Yaliang Li, Bolin Ding, and Jingren Zhou. 2024. A Simple and Provable Scaling Law for the Test-Time Compute of Large Language Models. arXiv:2411.19477 [cs.CL]
- [5] DeepSeek-AI. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. arXiv:2501.12948 [cs.CL]
- [6] Andrew Fowler, Kurt Partridge, Ciprian Chelba, Xiaojun Bi, Tom Ouyang, and Shumin Zhai. 2015. Effects of Language Modeling and Its Personalization on Touchscreen Typing Performance. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 649–658. doi:10.1145/2702123.2702503
- [7] Peter Frensch and Joachim Funke. 2002. Thinking and problem solving. *Encyclopedia of Life Support Systems (EOLSS), Developed under the Auspices of the UNESCO, Eolss Publishers, Oxford, UK* (01 2002).
- [8] Yao Fu, Dong-Ki Kim, Jaekyeom Kim, Sungryull Sohn, Lajanugen Logeswaran, Kyunghoon Bae, and Honglak Lee. 2024. AutoGuide: Automated Generation and Selection of Context-Aware Guidelines for Large Language Model Agents. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=mRIQz8Zd6O>
- [9] Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. 2024. Large Language Model Based Multi-agents: A Survey of Progress and Challenges. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, Kate Larson (Ed.). International Joint Conferences on Artificial Intelligence Organization, 8048–8057. doi:10.24963/ijcai.2024/890 Survey Track.
- [10] Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu. 2023. Reasoning with Language Model is Planning with World Model. In *The 2023 Conference on Empirical Methods in Natural Language Processing*. <https://openreview.net/forum?id=V7WWvYf1R>
- [11] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. The Curious Case of Neural Text Degeneration. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=rygQYrFvH>
- [12] Yuki Ichihara, Yuu Jinnai, Tetsuro Morimura, Kenshi Abe, Kaito Ariu, Mitsuki Sakamoto, and Eiji Uchibe. 2025. Evaluation of Best-of-N Sampling Strategies for Language Model Alignment. *Transactions on Machine Learning Research* (2025). <https://openreview.net/forum?id=H4S4ETc8c9>
- [13] Joel Jang, Seungone Kim, Bill Yuchen Lin, Yizhong Wang, Jack Hessel, Luke Zettlemoyer, Hannaneh Hajishirzi, Yejin Choi, and Prithviraj Ammanabrolu. 2023. Personalized Soups: Personalized Large Language Model Alignment via Post-hoc Parameter Merging. arXiv:2310.11564
- [14] Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan O Arik, Dong Wang, Hamed Zamani, and Jiawei Han. 2025. Search-R1: Training LLMs to Reason and Leverage Search Engines with Reinforcement Learning. In *Second Conference on Language Modeling*. <https://openreview.net/forum?id=Rwhi91Ideu>
- [15] Pranav Kasela, Marco Braga, Gabriella Pasi, and Raffaele Perego. 2024. SE-PQA: Personalized Community Question Answering. In *Companion Proceedings of the ACM Web Conference 2024* (Singapore, Singapore) (WWW '24). Association for Computing Machinery, New York, NY, USA, 1095–1098. doi:10.1145/3589335.3651445
- [16] Ishita Kumar, Snigdha Viswanathan, Sushrita Yerra, Alireza Salemi, Ryan A. Rossi, Franck Dernoncourt, Hanieh Deilamsalehy, Xiang Chen, Ruiyi Zhang, Shubham Agarwal, Nedim Lipka, and Hamed Zamani. 2024. LongLaMP: A Benchmark for Personalized Long-form Text Generation. arXiv:2407.11016 [cs.CL]
- [17] Cheng Li, Mingyang Zhang, Qiaozhu Mei, Weize Kong, and Michael Bendersky. 2024. Learning to Rewrite Prompts for Personalized Text Generation. In *Proceedings of the ACM on Web Conference 2024* (WWW '24). ACM. doi:10.1145/3589334.3645408
- [18] Cheng Li, Mingyang Zhang, Qiaozhu Mei, Yaqing Wang, Spurthi Amba Hombaiah, Yi Liang, and Michael Bendersky. 2023. Teach LLMs to Personalize – An Approach inspired by Writing Education. arXiv:2308.07968
- [19] Hammeng Liu, Zhiyang Teng, Leyang Cui, Chaoli Zhang, Qiji Zhou, and Yue Zhang. 2023. LogiCoT: Logical Chain-of-Thought Instruction Tuning. In *The 2023 Conference on Empirical Methods in Natural Language Processing*. <https://openreview.net/forum?id=qlCtkvgQJH>
- [20] Yanming Liu, Xinyue Peng, Jiannan Cao, Shi Bo, Yuwei Zhang, Xuhong Zhang, Sheng Cheng, Xun Wang, Jianwei Yin, and Tianyu Du. 2025. Tool-Planner: Task Planning with Clusters across Multiple Tools. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=dRz3cizftU>
- [21] Zhuoran Lu, Sheshera Mysore, Tara Safavi, Jennifer Neville, Longqi Yang, and Mengting Wan. 2024. Corporate Communication Companion (CCC): An LLM-empowered Writing Assistant for Workplace Social Media. arXiv:2405.04656
- [22] Sheshera Mysore, Zhuoran Lu, Mengting Wan, Longqi Yang, Steve Menezes, Tina Baghaee, Emmanuel Barajas Gonzalez, Jennifer Neville, and Tara Safavi. 2023. PEARL: Personalizing Large Language Model Writing Assistants with Generation-Calibrated Retrievers. arXiv:2311.09180
- [23] Maxim Naumov, Dheevatsa Mudigere, Hao-Jun Michael Shi, Jianyu Huang, Narayanan Sundaraman, Jongsoo Park, Xiaodong Wang, Udit Gupta, Carole-Jean Wu, Alisson G. Azzolini, Dmytro Dzhulgakov, Andrey Malleevich, Ilya Cherniavskii, Yinghai Lu, Raghuraman Krishnamoorthi, Ansha Yu, Volodymyr Kondratenko, Stephanie Pereira, Xianjie Chen, Wenlin Chen, Vijay Rao, Bill Jia, Liang Xiong, and Misha Smelyanskiy. 2019. Deep Learning Recommendation Model for Personalization and Recommendation Systems. arXiv:1906.00091 [cs.IR]
- [24] OpenAI. 2024. GPT-4o System Card. arXiv:2410.21276 [cs.CL]
- [25] OpenAI. 2024. OpenAI o1 System Card.
- [26] Sungjin Park, Xiao Liu, Yeyun Gong, and Edward Choi. 2024. Ensembling Large Language Models with Process Reward-Guided Tree Search for Better Complex Reasoning. arXiv:2412.15797 [cs.CL]
- [27] Alireza Salemi, Surya Kallumadi, and Hamed Zamani. 2024. Optimization Methods for Personalizing Large Language Models through Retrieval Augmentation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Washington DC, USA) (SIGIR '24). Association for Computing Machinery, New York, NY, USA, 752–762. doi:10.1145/3626772.3657783
- [28] Alireza Salemi, Cheng Li, Mingyang Zhang, Qiaozhu Mei, Weize Kong, Tao Chen, Zhuowan Li, Michael Bendersky, and Hamed Zamani. 2025. Reasoning-Enhanced Self-Training for Long-Form Personalized Text Generation. arXiv:2501.04167 [cs.CL] <https://arxiv.org/abs/2501.04167>
- [29] Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. 2024. LaMP: When Large Language Models Meet Personalization. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 7370–7392. <https://aclanthology.org/2024.acl-long.399>
- [30] Alireza Salemi and Hamed Zamani. 2024. Comparing Retrieval-Augmentation and Parameter-Efficient Fine-Tuning for Privacy-Preserving Personalization of Large Language Models. arXiv:2409.09510 [cs.CL]
- [31] Alireza Salemi and Hamed Zamani. 2025. LaMP-QA: A Benchmark for Personalized Long-form Question Answering. arXiv:2506.00137 [cs.CL]
- [32] Alireza Salemi and Hamed Zamani. 2025. Learning from Natural Language Feedback for Personalized Question Answering. arXiv:2508.10695 [cs.CL] <https://arxiv.org/abs/2508.10695>
- [33] Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2024. Scaling LLM Test-Time Compute Optimally can be More Effective than Scaling Model Parameters. arXiv:2408.03314 [cs.LG]
- [34] Robert J. Sternberg and Talia Ben-Zeev. 2001. Complex Cognition: The Psychology of Human Thought.
- [35] Zhaoxuan Tan, Zheyuan Liu, and Meng Jiang. 2024. Personalized Pieces: Efficient Personalized Large Language Models through Collaborative Efforts. arXiv:2406.10471
- [36] Chaojie Wang, Yanchen Deng, Zhiyi Lyu, Liang Zeng, Jujie He, Shuicheng Yan, and Bo An. 2024. Q\*: Improving Multi-step Reasoning for LLMs with Deliberative Planning. arXiv:2406.14283 [cs.AI]
- [37] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA) (NIPS '22). Curran Associates Inc., Red Hook, NY, USA, Article 1800, 14 pages.
- [38] Chengxing Xie and Difan Zou. 2024. A Human-Like Reasoning Framework for Multi-Phases Planning Task with Large Language Models. In *ICML 2024 Workshop on LLMs and Cognition*. <https://openreview.net/forum?id=BX6wC2452j>
- [39] Gui-Rong Xue, Jie Han, Yong Yu, and Qiang Yang. 2009. User Language Model for Collaborative Personalized Search. *ACM Trans. Inf. Syst.* 27, 2, Article 11 (mar 2009), 28 pages. doi:10.1145/1462198.1462203
- [40] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik R Narasimhan. 2023. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. In *Thirty-seventh Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=5Xc1ecx0Ih>
- [41] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. ReAct: Synergizing Reasoning and Acting in Language Models. *arXiv preprint arXiv:2210.03629* (2022).
- [42] Han Yin, Jianxing Yu, Miaopei Lin, and Shiqi Wang. 2024. Answering Spatial Commonsense Questions Based on Chain-of-Thought Reasoning with Adaptive Complexity. In *Web and Big Data*, Wenjie Zhang, Anthony Tung, Zhonglong Zheng, Zhengyi Yang, Xiaoyang Wang, and Hongjie Guo (Eds.). Springer Nature

```

Action Selection Prompt

# Your Task
Your task is to select one of the actions between the given actions. The actions should be chosen considering that the main goal is generating a personalized response to the question. The action you choose should be the next best step toward the final goal, which is generating a personalized response to the question. The valid actions that you can choose from are:

(list of actions)
(action 1)
(action 2)
...
(action n)

You cannot choose any action that is not in the aforementioned valid actions list.

# Your output
Your output for selecting an action should be a valid JSON, enclosed in ""json"" block, with the following fields:
- action: a string representing your selected action.
- reason: a string representing the reason for why the action is selected and how it is helpful to reach the final goal, which is generating a personalized response to the question.
- actionOutput: a json representing the output of the action based on the defined output format for the selected action.
    
```

Figure 7: Action selection prompt used in PoT.

```

Important Aspects Extraction from Profile

You are a personalized search assistant with great rating capabilities. Your task is to look at the previous questions and post details of a user, and infer what are the aspects that is important for the user and the user expect to see them in a response to their query.

# your input:
Your input is the previous questions and details of them that asked by the user.
# your output:
You should generate a valid JSON list of objects, enclosed in ""json"" block, each object with the following fields:
- aspect: The aspect that is important for the user.
- description: The description of the aspect.

previous questions and details:
(question & details 1)
...
(question & details |P|)

output: ""json
    
```

Figure 8: Prompt for extracting important aspects from the user profile used in PoT.

```

Best-of-N

You are a personalized search assistant with great rating capabilities. Your task is to select the best answer to the question of the user, considering the aspects that are important for the user.

# your input:
- aspects: a json list of aspects that are important for the user. Aspects are in the following format:
  - aspect: The aspect that is important for the user.
  - description: The description of the aspect.
- question: The question asked by the user.
- responses: a list of responses to the question of the user. Responses are in the following format:
  - response: The generated response to the question of the user.
  - index: The index of the response.

# Your task: Your task is to select the answer to the question of the user that covers the most of aspects that are important for the user.
## your output:
You should generate a valid JSON object, enclosed in ""json"" block that contains the best answer to the question of the user, in the following format:
- index: The index of the best answer to the question of the user that covers the most number of aspects that are important for the user.
- reason: a string that shows the reason for this selection.

# aspects:
(aspect: aspect 1, description: description aspect 1)
...
(aspect: aspect 1, description: description aspect K)
# question: {question}
# responses:
(index: 1, response: response 1)
...
(index: N, response: response N)

output: ""json
    
```

Figure 9: Prompt for selecting the best response out of N generated responses used in PoT.

Singapore, Singapore, 186–200.

[43] Kai Zhang, Yangyang Kang, Fubang Zhao, and Xiaozhong Liu. 2024. LLM-based Medical Assistant Personalization with Short- and Long-Term Memory Coordination. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Kevin Duh, Helena Gomez, and Steven Bethard (Eds.). Association for Computational Linguistics, Mexico City, Mexico, 2386–2398. <https://aclanthology.org/2024.naacl-long.132>

[44] Kexun Zhang, Shang Zhou, Danqing Wang, William Yang Wang, and Lei Li. 2024. Scaling LLM Inference with Optimized Sample Compute Allocation. arXiv:2410.22480 [cs.CL]

## A PoT’s Implementation Details

### A.1 Actions Definition

We define an action as a means by which the agent interacts with the text generation environment, specifically the LLM, by instructing it to perform a particular operation. Each action is associated with

```

Mixture-of-N

You are a personalized search assistant with great rating capabilities. Your task is to look at a set of personalized responses to the user’s question and combine them into a single personalized response, considering the aspects that are important for the user.

# your input:
- aspects: a json list of aspects that are important for the user. Aspects are in the following format:
  - aspect: The aspect that is important for the user.
  - description: The description of the aspect.
- question: The question asked by the user.
- responses: a list of responses to the question of the user. Responses are in the following format:
  - response: The generated response to the question of the user.
  - index: The index of the response.

# Your task: Your task is to look at a set of personalized responses to the user’s question and combine them into a single personalized response, considering the aspects that are important for the user.
## your output:
You should generate a valid JSON object, enclosed in ""json"" block that contains the best answer to the question of the user, in the following format:
- personalizedAnswer: The final combined personalized response to the user’s question considering all the generated responses and the important aspects for the user.

# aspects:
(aspect: aspect 1, description: description aspect 1)
...
(aspect: aspect 1, description: description aspect K)
# question: {question}
# responses:
(index: 1, response: response 1)
...
(index: N, response: response N)

output: ""json
    
```

Figure 10: Prompt for mixing the N generated responses into a single response used in PoT.

a specific output format that the LLM generates in response to the instruction. In this paper, we consider the following actions as the operations the model can take to respond to a question:

- *answering*: This action instructs the model to generate a personalized response to the question, provided that personalization can enhance the response. It should be selected when the model is prepared to answer the question, ensuring that all necessary information is available and there is no ambiguity.
- *planning*: This action instructs the model to generate a plan of steps required to address all aspects of the question. It should be executed when the question has multiple components, and the model recognizes that planning the next steps is necessary to adequately respond to the user’s query.
- *personalization*: This action asks the model to evaluate whether the user’s question can benefit from personalization. It is useful when determining whether personalizing the response is necessary or if a generic response would suffice.
- *personalizing*: This action asks the model to identify relevant information from the user profile and summarize how it can be used to answer the question. It is executed when personalized information is necessary to enhance the quality of the response.
- *reasoning*: This action prompts the model to break down one aspect of the question and think through it step by step. It is executed when questions involve multiple aspects, and the model needs to focus on and consider how to address a specific aspect.
- *clarifying*: This action prompts the model to assess whether the question is ambiguous and determine how it can be clarified. It is executed when the model identifies that the current information is insufficient or unclear, requiring further elaboration.
- *summarizing*: This action instructs the model to summarize all available information and findings. It should be executed when the model determines that the current information needs to be consolidated or synthesized to ensure a clearer understanding before continuing with the response generation.
- *revising*: This action instructs the model to revise the current response to the question. It should be executed when the model has already generated an initial response but believes the response can be improved or adjusted to better address the user’s query or to incorporate additional insights or information.

```

    In-context Personalization

    Your task is to answer a user's question from a post in a personalized way by considering this user's past post questions and detailed descriptions of these questions.

    # Your input:
    - The user's current question from a post.
    - The user's past post questions and detailed descriptions of these questions.

    # Your task:
    Answer the user's current question in a personalized way by considering this user's past post questions and detailed descriptions of these questions.

    # Your output:
    You should generate a valid JSON, enclosed in ```json``` block, with the following fields:
    - personalizedAnswer: a string representing your personalized answer to the user's current question by considering this user's past post questions and detailed descriptions of these questions.

    # input:
    ## Past post questions and detailed descriptions of these questions:
    (profile)
    ## Current post question:
    (query)
    output: ```json
  
```

Figure 11: Prompt for the In-Context Personalization.

```

    In-context Personalization /w CoT

    Your task is to answer a user's question from a post in a personalized way by considering this user's past post questions and detailed descriptions of these questions.

    # Your input:
    - The user's current question from a post.
    - The user's past post questions and detailed descriptions of these questions.

    # Your task:
    Answer the user's current question in a personalized way by considering this user's past post questions and detailed descriptions of these questions. You should do this using the following steps:
    - step 1: From the previous questions and their detailed descriptions, you should first identify the aspects that are important for the user to include in the response to their question.
    - step 2: Based on the extracted important aspects for the user, you should answer their current question in a personalized way.

    # Your output:
    You should generate a valid JSON, enclosed in ```json``` block, with the following fields:
    - importantAspects: a string representing the important aspects for the user extracted from their previous questions and their details explanation.
    - personalizedAnswer: a string representing your personalized answer to the user's current question by considering this user's past post questions and detailed descriptions of these questions.

    # input:
    ## Past post questions and detailed descriptions of these questions:
    (profile)
    ## Current post question:
    (query)
    output: ```json
  
```

Figure 12: Prompt for In-Context Personalization with CoT.

- *finalizing*: This action finalizes the response. It should be executed when the model believes the response is complete, accurate, and addresses the user's query. No further actions or revisions are necessary, and the response is ready to be presented.

## A.2 Details of Response Aggregation Methods

An essential component of Pathways of Thoughts involves aggregating the candidate responses generated across different thinking pathways into a single output. Each pathway produces a candidate response, collectively forming a set of response proposals, denoted as  $R_{PoT}$ . To generate the final response  $\hat{y}_{x_u}$  personalized for the user, it is crucial to select the most suitable candidate. To achieve this, we first employ the LLM  $\pi$  to extract the key aspects  $I_u$  relevant to user  $u$  from their user profile  $P_u$ . This extraction process is guided by the prompt illustrated in Figure 8, acting as an intermediate step to capture the user's individual preferences. Based on this, there are two methods to generate the final response:

- *Best-of-N*: Given the extracted important aspects for the user and the input, the LLM selects the optimal response from the set of candidate responses generated through multi-directional thinking. This is guided by the prompt shown in Figure 9.
- *Mixture-of-N*: Given the extracted important aspects for the user and the input, the LLM integrates different aspects of the candidate responses from the set of multi-directional thinking outputs, guided by the prompt in Figure 10. This method emulates human

```

    Planning in Tree-of-Thoughts

    Your task is to answer a user's question from a post in a personalized way by considering this user's past post questions and detailed descriptions of these questions.

    # Your input:
    - The user's current question from a post.
    - The user's past post questions and detailed descriptions of these questions.

    # Your task:
    Your task is to generate a plan for responding to the user's question in a personalized way by considering this user's past post questions and detailed descriptions of these questions.

    # Your output:
    You should generate a valid JSON, enclosed in ```json``` block, with the following fields:
    - plan: a string representing a plan of steps that we can take to answer all the aspects of the question.

    # Your input:
    ## Past post questions and detailed descriptions of these questions:
    (profile)
    ## Current post question:
    (query)
    output: ```json
  
```

Figure 13: Prompt for the planning step in ToT.

```

    Response Generation in Tree-of-Thoughts

    Your task is to answer a user's question from a post in a personalized way by considering this user's past post questions and detailed descriptions of these questions.

    # Your input:
    - The user's current question from a post.
    - The user's past post questions and detailed descriptions of these questions.
    - The plan that guides you to answer the question.

    # Your task:
    Your task is to generate a response to the user's question in a personalized way by considering this user's past post questions and detailed descriptions of them and the given plan on how to answer the question.

    # Your output:
    You should generate a valid JSON, enclosed in ```json``` block, with the following fields:
    - personalizedAnswer: a string representing your personalized answer to the user's current question by considering this user's past post questions and detailed descriptions of these questions and the given plan on how to answer the question.

    # Your input:
    ## Past post questions and detailed descriptions of these questions:
    (profile)
    ## Current post question:
    (query)
    ## Plan:
    (plan)
    output: ```json
  
```

Figure 14: Prompt for the response generation step in ToT.

```

    Plan Selection in Tree-of-Thoughts

    You are a personalized search assistant with great rating capabilities. Your task is to select the best plan to answer the question of the user, considering the aspects that are important for the user.

    # your input:
    - aspects: a json list of aspects that are important for the user. Aspects are in the following format:
      - aspect: The aspect that is important for the user.
      - description: The description of the aspect.
    - question: The question asked by the user.
    - plans: a list of plans to answer the question of the user. Plans are in the following format:
      - plan: The generated response to the question of the user.
      - index: The index of the response.

    # Your task:
    Your task is to select the plan to answer the question of the user that covers the most of aspects that are important for the user.

    ## your output:
    You should generate a valid JSON object, enclosed in ```json``` block that contains the best answer to the question of the user, in the following format:
    - index: The index of the best plan to answer the question of the user that covers the most number of aspects that are important for the user.
    - reason: a string that shows the reason for this selection.

    # aspects:
    (aspects)
    # question:
    (question)
    # plans:
    (plans)
    output: ```json
  
```

Figure 15: Prompt for the plan selection step in Tree of Thoughts [40]. The important aspects extracted for the user are generated using the prompt shown in Figure 8.

problem-solving by exploring multiple solutions and synthesizing a new response that leverages the strengths of all candidates.

## B Baselines Prompts

Figures 11 and 12 present the prompts used for the personalized baselines without and with chain-of-thought reasoning, respectively. For the Tree of Thoughts baseline, Figure 13 shows the planning prompt, Figure 14 illustrates the response generation prompt, and Figure 15 provides the plan selection prompt.

<p><b>Pathway 1: [plan, personalize, reason, reason, reason, reason, reason, summarize, answer, finalize]</b></p> <p>Bringing up a political concern at work, especially in a minority-led business, requires careful consideration and sensitivity. Based on your previous posts, it seems you value clear and effective communication, which is essential in this situation. First, it's important to clarify your own thoughts. What is the specific political topic you're concerned about, and what outcome are you hoping to achieve by bringing it up? Are you aware of any company policies regarding political discussions in the workplace? It's understandable that you're mindful of the context of working in a minority-led business. These environments can sometimes be more sensitive to political discussions, especially if the topic relates to issues that directly impact the community the business represents. When approaching the conversation, focus on shared values and goals that resonate with the company's mission and the concerns of its leadership. Choose the right time and place for a private conversation, avoiding public forums or emotionally charged moments. Be respectful, listen actively, and approach the conversation with a willingness to understand different perspectives. Instead of dwelling on problems, propose concrete actions or solutions that align with the company's values and address your concerns. Finally, it's important to be aware of relevant laws, such as freedom of speech and workplace discrimination. While you have the right to express your political views, employees are often required to remain neutral and avoid making generalizations or assumptions. It's best to familiarize yourself with your company's policies and relevant laws in your jurisdiction. Clear and concise communication is key to navigating this sensitive situation effectively.</p>	<p><b>Pathway 2: [plan, personalize, reason, reason, reason, reason, reason, summarize, answer, finalize]</b></p> <p>It's understandable that you're being mindful of the business being minority-led. It's important to approach sensitive topics like politics with respect and consideration for everyone's perspectives. First, consider the specific political concern you want to raise and how it relates to the business. Is it directly relevant to the business's operations, values, or industry? How might this concern affect employees, customers, or the business's reputation? Remember when you were asking about writing your petition clearly and concisely? Those same principles apply here. When discussing a sensitive political topic, it's crucial to be clear about your concerns, provide specific examples if possible, and avoid overly emotional or accusatory language. You've also shown a strong ability to communicate passionately and effectively in your grant writing, so tap into those skills to express your concerns constructively. Here are a few ways you could bring up the topic: "Open dialogue with management." If you feel comfortable, consider having an open and honest conversation with your manager or a trusted supervisor. Explain your concerns clearly and respectfully, focusing on how the issue might impact the business or work environment. "Create a safe space for discussion." Perhaps suggest creating a safe space for employees to discuss their concerns and perspectives on political matters. This could be a facilitated discussion group or an open forum where everyone feels comfortable sharing their thoughts. "Written communication." If you prefer a more formal approach, consider writing a well-structured letter or email outlining your concerns and suggestions. This allows you to carefully articulate your thoughts and ensure clarity. Remember how you emphasized clear and passionate communication in your grant writing? Those same principles apply here, but with an added layer of sensitivity. "Use respectful language." Acknowledge that others may have different viewpoints and try to understand their reasoning. "Practice active listening." Listen attentively to what others have to say, even if you disagree. "Find common ground." Focus on shared values and goals, and try to find areas of agreement. "Use respectful language." Avoid inflammatory language or personal attacks. It's important to be prepared for the possibility of disagreement or conflict. Not everyone may share your views, and that's okay. If you encounter resistance or differing points, try to: "Stay calm and respectful." Avoid making personal attacks. "Focus on the facts." Stick to the issue at hand and avoid making generalizations or assumptions. "Seek common ground." Even if you disagree on the specifics, try to find areas of shared values or goals. "Know when to disengage." If the conversation becomes unproductive or heated, it's okay to politely disengage and revisit the topic at another time. By approaching the conversation with respect and a willingness to listen, you can create a more productive and positive dialogue.</p>	<p><b>Pathway 3: [plan, check, personalization, personalize, reason, reason, reason, summarize, answer, finalize]</b></p> <p>It's understandable to feel apprehensive about bringing up a political concern, especially within a minority-led business. Your past experiences navigating travel as a blind person, managing medical needs abroad, and advocating for yourself as a grant recipient demonstrate your resilience and ability to navigate potentially sensitive situations. Remember those successes as you approach this conversation. Here are some tips to help you navigate this conversation effectively: "Focus on shared values." Frame your concern in a way that highlights shared values with the business leadership. Emphasize how it relates to the business's mission, impacts employees or the community, or aligns with principles of fairness and respect. "Choose the right time and setting." Request a private meeting to discuss your concerns, rather than raising them in a public or informal setting. "Use respectful language." Communicate clearly and calmly, avoiding accusatory or inflammatory language. Focus on expressing your perspective and understanding theirs. "Be an active listener." Listen attentively to the perspectives of the business leadership and be open to dialogue. This demonstrates respect and a willingness to find common ground. "Have a clear goal." Define what you hope to achieve through the conversation. Are you seeking change, raising awareness, or simply expressing your concerns? "Prepare for different outcomes." The business leadership might be receptive to your concerns and open to dialogue, or they might be dismissive or defensive. Consider how you will respond in each scenario, maintaining your composure and advocating for your needs respectfully. Remember, you have the right to express your concerns, and doing so with respect and sensitivity can lead to a productive conversation. Your past experiences have equipped you to handle this situation with grace and determination.</p>	<p><b>Pathway 4: [plan, check, personalization, personalize, reason, reason, reason, summarize, answer, finalize]</b></p> <p>It's commendable that you want to advocate for immigrant rights, especially given the current political climate. However, discussing political issues in a minority-led business can be challenging. Here's a breakdown of key points: "Sensitivity." Recognize that political discussions can be divisive. Approach the topic with respect and empathy, understanding that others may have different viewpoints. "Business Values." Research if your company has any stated positions on immigration. Aligning your advocacy with existing company values can make your approach more effective. "Power Dynamics." Be mindful of your position within the company and the potential impact of your words. Consider the perspectives of the minority leadership, who may have unique experiences and sensitivities related to immigration. "Communication Strategies." Start by having informal conversations with colleagues to gauge their perspectives and build consensus. "If you decide to approach leadership, do so respectfully and professionally." Frame your concerns in a way that highlights the potential benefits to the business, such as attracting diverse talent or strengthening community ties. "Be prepared to listen to opposing viewpoints and engage in constructive dialogue." "Legal Considerations." Familiarize yourself with your company's policies on political discussions in the workplace. "Understand and respect professional boundaries." While you advocate for your cause, your goal is to advocate for a cause you believe in while maintaining a positive and professional relationship. "Be open to dialogue." Remember, your goal is to have a conversation thoughtfully and strategically, you can increase your chances of making a meaningful impact.</p>
<p><b>Pathway 5: [plan, check, personalization, personalize, reason, reason, reason, summarize, answer, finalize]</b></p> <p>It's commendable that you're being mindful of the fact that your workplace is minority-led. This shows a sensitivity to the potential complexities of discussing political topics in this context. Your experience with grant writing has likely taught you the importance of presenting information persuasively and respectfully. These skills will be valuable when discussing political concerns. As you'll need to articulate your points clearly and consider the perspectives of others. Additionally, your emphasis on conciseness will help ensure your message is impactful and easy to understand. It's understandable that you're feeling apprehensive about bringing up a political topic at work, especially since it's a minority-led business. These situations require a delicate balance of advocating for your beliefs while respecting the company's values and the potential vulnerabilities of a minority-led organization. Before taking any action, it's important to clarify your own goals. Do you want to advocate for a specific policy change, express dissatisfaction, or simply engage in open dialogue? Understanding your motivation will help you tailor your approach. Given your experience with petitions and grant applications, you clearly value clear and concise communication. This will be a valuable asset in navigating this conversation. Here's how you can approach this: "Research." Understand the business and industry context. How do your views align with the company's values? "Choose the right setting." Initiate a private conversation, away from customers or colleagues. "Prepare your points." Organize your thoughts and present them in a clear, concise, and respectful manner. "Focus on shared values." Frame the conversation around common ground and potential benefits to the business. "Be a good listener." Engage in active listening and be open to different perspectives. "Offer solutions." Instead of just raising concerns, propose actionable steps the business could take. Remember, your goal is to navigate sensitive situations while traveling with a disability can be helpful here. Assess the potential risks before proceeding. Political discussions at work can strain relationships, impact the company's reputation, or even jeopardize your job security. Weigh the potential benefits against these risks and make a responsible decision that aligns with your values and goals.</p>	<p><b>Pathway 6: [plan, check, personalization, personalize, reason, reason, reason, summarize, answer, finalize]</b></p> <p>It's understandable that you're feeling apprehensive about bringing up a political topic at work, especially since it's a minority-led business. These situations require a delicate balance of advocating for your beliefs while respecting the company's values and the potential vulnerabilities of a minority-led organization. Before taking any action, it's important to clarify your own goals. Do you want to advocate for a specific policy change, express dissatisfaction, or simply engage in open dialogue? Understanding your motivation will help you tailor your approach. Given your experience with petitions and grant applications, you clearly value clear and concise communication. This will be a valuable asset in navigating this conversation. Here's how you can approach this: "Research." Understand the business and industry context. How do your views align with the company's values? "Choose the right setting." Initiate a private conversation, away from customers or colleagues. "Prepare your points." Organize your thoughts and present them in a clear, concise, and respectful manner. "Focus on shared values." Frame the conversation around common ground and potential benefits to the business. "Be a good listener." Engage in active listening and be open to different perspectives. "Offer solutions." Instead of just raising concerns, propose actionable steps the business could take. Remember, your goal is to navigate sensitive situations while traveling with a disability can be helpful here. Assess the potential risks before proceeding. Political discussions at work can strain relationships, impact the company's reputation, or even jeopardize your job security. Weigh the potential benefits against these risks and make a responsible decision that aligns with your values and goals.</p>	<p><b>Pathway 7: [plan, check, personalization, answer, finalize]</b></p> <p>It's commendable that you're being mindful of the dynamics within your workplace. Here's how to approach a political topic with sensitivity, especially in a minority-led business: "Choose the right time and place." Avoid bringing up politics during work hours or in common areas. If you feel a discussion is necessary, consider a more private setting and time. "Focus on shared values." Frame your concerns in a way that aligns with the company's mission and values. This helps create common ground and fosters a more collaborative atmosphere. "Be prepared to listen to perspectives that may differ from your own." Show empathy and understanding, even if you disagree. "Avoid inflammatory language." Use neutral and respectful language. "Acknowledge generalizations or assumptions about others' political beliefs." "Be mindful of power dynamics." Be aware of the potential power imbalances that may exist in a minority-led business. Approach the conversation with humility and a willingness to learn from others' experiences. "Focus on solutions." Instead of dwelling on problems, try to focus on potential solutions or ways to address the issue constructively. Remember, open and respectful communication is key. By approaching the conversation thoughtfully and with a willingness to understand different perspectives, you can create a more inclusive and productive work environment.</p>	<p><b>Pathway 8: [plan, check, personalization, reason, answer, finalize]</b></p> <p>Bringing up political topics at work can be tricky, especially in a minority-led business where there might be diverse perspectives. Here's how to approach it: "Prioritize Respect." Start by acknowledging that everyone has different viewpoints. Listen actively, be empathetic, and avoid making assumptions about others' beliefs. "Be Mindful of Context." While it's important not to stereotype, recognize that a minority-led business might have unique sensitivities related to certain political issues. Be sensitive to this context. "Find Common Ground." Frame the conversation around shared goals and values. Look for areas of agreement and emphasize the benefits to the business. "Choose the Right Channel." Consider whether a private conversation with a supervisor or HR is more appropriate than a public announcement. "Stay Professional." Avoid personal attacks, inflammatory language, and emotional outbursts. Stick to the facts and present your concerns calmly and respectfully. Remember, open and honest communication is important, but it's crucial to approach sensitive topics with tact and understanding.</p>
<p><b>Pathway 9: [plan, check, personalization, personalize, reason, reason, reason, summarize, answer, finalize]</b></p> <p>It's great that you're being mindful of the business's background as minority-led. It shows your sensitivity to the potential complexities of discussing political topics in this context. Raising political concerns can be tricky in any workplace, as it can sometimes feel like there's a power imbalance. It's important to approach the conversation with sensitivity and respect for everyone's perspectives. Based on your past questions, it seems you have a good understanding of navigating potentially sensitive situations, especially when it comes to cultural contexts and legal considerations. This experience can be valuable when approaching a political discussion in your workplace. Just as you carefully considered accessibility and legal matters while traveling, it's important to approach this conversation with sensitivity and awareness. Here are some tips for having a constructive conversation: "Find common ground." Instead of focusing on differences, try to identify shared values and goals. "Frame concerns constructively." Phrase your concerns in a way that emphasizes collaboration and finding solutions. You could say something like, "I'm concerned about [issue] and I'm wondering if we can discuss ways to improve this together." "Listen actively and understand respectfully." Be open to hearing different perspectives and try to understand where others are coming from. Finally, it's always a good idea to familiarize yourself with your company's policies regarding political discussions. This will help you understand the appropriate channels for raising concerns and ensure you're approaching the situation in a way that aligns with company guidelines.</p>	<p><b>Pathway 10: [plan, check, personalization, personalize, reason, reason, reason, summarize, answer, finalize]</b></p> <p>It's understandable that you're cautious about bringing up a political topic at work, especially in a minority-led business. It's important to approach this conversation with respect and sensitivity, drawing on your past experiences advocating for your needs. Here's how you can approach this: "Clarify your goals." Before taking any action, it would be helpful to understand what you hope to achieve by bringing up this topic. Are you looking to raise awareness, advocate for change, or simply express your views? Also, consider how this topic aligns with the business's values and mission. This will help you frame the conversation in a way that resonates with the leadership. "Choose the right time and place." Find an appropriate time and place, perhaps during a one-on-one meeting or a team discussion where open dialogue is encouraged. "Frame your concerns constructively." Focus on shared values and potential solutions rather than criticism. Be prepared to listen actively to the perspectives of others, even if they differ from your own. "Keep in mind that you may encounter differing viewpoints or sensitivities. This is natural, especially when discussing political topics. If this happens, try to remain calm and respectful, and focus on understanding the other person's perspective. If the conversation becomes heated, it might be best to politely disengage and revisit the topic at another time.</p>	<p><b>Pathway 11: [plan, check, personalization, reason, answer, summarize, answer, finalize]</b></p> <p>Bringing up political concerns at work can be tricky, especially in a minority-led business where sensitivities might be heightened. Here's how to approach it: "1. Consider the context." "Company culture." Is your workplace generally open to discussing political issues? "Impact on the business." Does the political issue directly affect your work or the company's mission? "2. Choose the right time and place." "Avoid bringing up politics during busy work hours or in public settings where others might feel uncomfortable. If you need to discuss it, consider scheduling a private meeting with your supervisor or HR, ensuring a respectful and inclusive." "Use respectful language and avoid generalizations or stereotypes about any group of people." "Acknowledge that others may have different viewpoints or sensitivities." "Focus on shared values and potential solutions." "3. Prepare and Practice." Before initiating, research the political topic thoroughly and consider how it might intersect with the business's mission and values. Practice framing your concerns in a clear, concise, and respectful manner. This preparation will help you communicate effectively and confidently. "4. Choose the Right Time and Place." Find a private and appropriate setting to have this conversation, ensuring that you and your employer feel comfortable and have ample time to discuss the matter. "5. Be a Good Listener." Active listening is crucial. Pay attention to your employer's perspective, ask clarifying questions, and demonstrate that you value their input. This approach will help build trust and understanding. Remember, your goal is to express your concerns while maintaining a positive and respectful relationship. By approaching the conversation thoughtfully and leveraging your communication skills, you can navigate this sensitive situation effectively.</p>	<p><b>Pathway 12: [plan, check, personalization, personalize, reason, answer, summarize, answer, finalize]</b></p> <p>It's understandable that you're cautious about bringing up political topics at work, especially in a minority-led business. Your past questions show you're aware of navigating sensitive situations and understanding the importance of context. Here's how you can approach this: "Consider the context." Is the topic relevant to your work or the company's mission? Is there a specific issue affecting the workplace or the community that you feel strongly about? "Choose your moment wisely." Avoid bringing up politics during busy times or when emotions are running high. Opt for a calm and respectful setting. "Focus on shared values." Frame your concerns in a way that aligns with the company's values and mission. Emphasize common ground and shared goals. "Be mindful of power dynamics." Remember your position within the company and be respectful of those in leadership roles. "Listen actively." Engage in a two-way conversation, actively listening to others' perspectives, even if you disagree. "Be clear and concise." Just like your successful grant writing, where you had to communicate your goals persuasively, use those same skills to articulate your points clearly and effectively, avoiding jargon or overly emotional language. "Share your company's policies." Familiarize yourself with any workplace policies regarding political discussions. Remember, your past experiences navigating sensitive situations, like ensuring accessibility for those with special needs or understanding regulations around medication, have equipped you with the skills to approach this conversation thoughtfully and respectfully. By focusing on shared values, clear communication, and respectful dialogue, you can contribute to a positive and productive workplace environment.</p>
<p><b>Pathway 13: [plan, personalize, reason, reason, reason, reason, reason, summarize, answer, finalize]</b></p> <p>It's commendable that you're being mindful of the sensitivities involved in discussing politics, especially within a minority-led business. It shows your thoughtful and considerate of the potential impact of your words. Drawing from your experience in writing your petition, remember that focusing on shared values and common goals can be a good starting point. Just as you sought to persuade your audience in the petition, you can aim to find common ground and build understanding when discussing political topics. It's wise to be particularly aware of the power dynamics at play, especially since the business is minority-led. Consider the perspectives and experiences of those in leadership positions, and how your views might be received. Active listening and showing respect for their experiences can go a long way in fostering a constructive dialogue. Remember your own emphasis on clear and concise communication. Apply that same principle to political discussions, focusing on shared values, understanding different viewpoints, and expressing your concerns respectfully. This approach can help foster a constructive dialogue rather than a heated debate. Here are some practical tips: 1) Choose an appropriate time and place for discussion, avoiding sensitive moments or public settings. 2) Frame your concerns in a way that focuses on shared values and common goals. 3) Be prepared to listen actively and understand different perspectives, even if you disagree. 4) Avoid accusatory language or personal attacks. 5) Be willing to compromise and find common ground. Just as you carefully researched accessibility before traveling, take time to observe the workplace culture and dynamics. Pay attention to how colleagues interact, what topics are openly discussed, and the general political climate. This will give you valuable insights into the best approach for bringing up your concerns.</p>	<p><b>Pathway 14: [plan, check, personalization, personalize, reason, reason, reason, summarize, answer, finalize]</b></p> <p>It's understandable that you want to address this concern, especially given your history of advocating for important issues. However, approaching this topic with a minority business owner requires a delicate touch. Here's how you can do this: "Acknowledge the Sensitivity." Begin by recognizing the potential sensitivities of the discussion and your position within a minority-led business. Your employer may have unique experiences and perspectives related to their identity, so approach with empathy and respect. "Focus on Open Dialogue." Frame your concerns constructively, emphasizing shared values instead of making accusations or engaging in divisive language. Focus on understanding your employer's viewpoint and common ground. Remember, your past successes in grant writing demonstrate your ability to articulate your thoughts persuasively. Use those skills to foster a productive conversation. "3. Prepare and Practice." Before initiating, research the political topic thoroughly and consider how it might intersect with the business's mission and values. Practice framing your concerns in a clear, concise, and respectful manner. This preparation will help you communicate effectively and confidently. "4. Choose the Right Time and Place." Find a private and appropriate setting to have this conversation, ensuring that you and your employer feel comfortable and have ample time to discuss the matter. "5. Be a Good Listener." Active listening is crucial. Pay attention to your employer's perspective, ask clarifying questions, and demonstrate that you value their input. This approach will help build trust and understanding. Remember, your goal is to express your concerns while maintaining a positive and respectful relationship. By approaching the conversation thoughtfully and leveraging your communication skills, you can navigate this sensitive situation effectively.</p>	<p><b>Pathway 15: [plan, check, personalization, personalize, reason, reason, summarize, answer, finalize]</b></p> <p>Raising political concerns at work can be tricky, especially in a minority-led business where sensitivities might be heightened. Here's how to approach this: "1. Consider the context." "Company culture." Is your workplace generally open to discussing political issues? "Impact on the business." Does the political issue directly affect your work or the company's mission? "2. Choose the right time and place." "Avoid bringing up politics during busy work hours or in public settings where others might feel uncomfortable. If you need to discuss it, consider scheduling a private meeting with your supervisor or HR, ensuring a respectful and inclusive." "Use respectful language and avoid generalizations or stereotypes about any group of people." "Acknowledge that others may have different viewpoints or sensitivities." "Focus on shared values and potential solutions." "3. Prepare and Practice." Before initiating, research the political topic thoroughly and consider how it might intersect with the business's mission and values. Practice framing your concerns in a clear, concise, and respectful manner. This preparation will help you communicate effectively and confidently. "4. Choose the Right Time and Place." Find a private and appropriate setting to have this conversation, ensuring that you and your employer feel comfortable and have ample time to discuss the matter. "5. Be a Good Listener." Active listening is crucial. Pay attention to your employer's perspective, ask clarifying questions, and demonstrate that you value their input. This approach will help build trust and understanding. Remember, your goal is to express your concerns while maintaining a positive and respectful relationship. By approaching the conversation thoughtfully and leveraging your communication skills, you can navigate this sensitive situation effectively.</p>	
<p><b>Pathway 16: [plan, check, personalization, personalize, reason, answer, summarize, answer, finalize]</b></p>			
<p><b>Combined Response with Mixture-of-16</b></p> <p>It's commendable that you're being mindful of the dynamics within your workplace, especially given that it's minority-led. Bringing up political concerns in any workplace can be tricky, but your past experiences with grant writing, advocating for accessibility, and navigating complex situations while traveling have equipped you with valuable skills for this conversation. Here's a combined approach, drawing from the best advice and considering your emphasis on clear, concise, and respectful communication: "Preparation and Understanding." "Clarify your goals." What do you hope to achieve by bringing up this topic? Are you seeking change, raising awareness, or simply expressing your concerns? [Response 5] "Research." Understand the business's stance on the issue and consider how your views align with their values. [Response 5] "Contextualize." Is the topic directly relevant to the business's operations, values, or industry? How might it affect employees, customers, or the business's reputation? [Response 1] "Company Culture." Is your workplace generally open to discussing political matters? [Response 10] "Framing the Conversation." "Focus on shared values." Frame your concerns in a way that aligns with the company's mission and values, emphasizing common ground and potential benefits for the business. [Responses 2, 4, 5, 6, 14] "Use respectful language." Avoid inflammatory rhetoric, personal attacks, and emotional outbursts. Stick to the facts and present your concerns calmly and respectfully. [Responses 1, 7, 11, 12] "Be a good listener." Engage in active listening and be open to understanding different perspectives, even if you disagree. [Responses 1, 2, 4, 5, 6, 8, 9, 11, 13] "Offer solutions." Instead of just raising concerns, propose actionable steps the business could take. [Response 5, 10, 14] "Choosing the Right Approach." "Open dialogue with management." If you feel comfortable, consider having a private conversation with your manager or a trusted supervisor. [Response 1] "Create a safe space for discussion." Suggest a facilitated discussion group or an open forum where employees feel comfortable sharing their thoughts. [Response 1] "Written communication." Consider a well-structured letter or email outlining your concerns and suggestions if you prefer a more formal approach. [Response 1] "Choose the right time and setting." Avoid bringing up politics during busy work hours or in inappropriate settings. Opt for a private meeting or a time when open dialogue is encouraged. [Responses 2, 6, 9, 10, 11, 13, 14] "Remember." "Power dynamics." Be mindful of your position within the company and the potential impact of your words, especially in a minority-led business. [Responses 3, 6, 8, 11, 12] "Be prepared for different outcomes." The business leadership might be receptive or dismissive. Consider how you will respond in each scenario, maintaining your composure and advocating for your needs respectfully. [Response 2] "Know when to disengage." If the conversation becomes unproductive or heated, it's okay to politely disengage and revisit the topic at another time. [Response 1, 9] "Company policies." Familiarize yourself with any workplace policies regarding political discussions. [Responses 3, 11, 8] By approaching the conversation with sensitivity, respect, and a willingness to listen, you can create a more productive and positive dialogue, even on potentially sensitive topics. Your experience with clear and concise communication will be a valuable asset in navigating this conversation effectively.</p>			

Figure 16: A case study of Pathways of Thoughts with  $N = 16$  and  $T = 8$  for question: "What's the best way for me to bring up a topic of political concern if the business I work for is minority-led?"