

# Understanding the Relationship between Prompts and Response Uncertainty in Large Language Models

Anonymous ACL submission

## Abstract

Large language models (LLMs) are widely used in decision-making, but their reliability, especially in critical tasks like healthcare, is not well-established. Therefore, understanding how LLMs reason and make decisions is crucial for their safe deployment. This paper investigates how the uncertainty of responses generated by LLMs relates to the information provided in the input prompt. Leveraging the insight that LLMs learn to infer latent concepts during pretraining, we propose a prompt-response concept model that explains how LLMs generate responses and helps understand the relationship between prompts and response uncertainty. We show that the uncertainty decreases as the prompt's informativeness increases, similar to epistemic uncertainty. Our detailed experimental results on real datasets validate our proposed model.

## 1 Introduction

Large language models (LLMs) have demonstrated impressive performance across a variety of tasks (Google, 2023; OpenAI, 2023; Zhao et al., 2023). This success has led to their widespread adoption and significant involvement in various decision-making applications, such as healthcare (Karabacak and Margetis, 2023; Sallam, 2023; Yang et al., 2023), education (Xiao et al., 2023), finance (Wu et al., 2023b), and law (Zhang et al., 2023a). However, despite their rapid adoption, the reliability of LLMs in handling high-stakes tasks has yet to be demonstrated (Arkoudas, 2023; Huang et al., 2023a). The reliability is particularly critical in domains such as healthcare, where model responses can have immediate and significant impacts on human behavior and hence their well-being (Ji et al., 2023). Therefore, understanding LLMs' reasoning and decision-making processes and how they influence response uncertainty is critical for their safe deployment.

To understand this importance, consider the mobile health (mHealth) application in which machine learning algorithms are integrated to monitor users' health conditions and provide advice on daily activities (Boursalie et al., 2018; Trella et al., 2022, 2023). Providing suggestions that can influence users' health is a form of intervention in the human decision-making process. For LLMs to be suitable for such use cases, they should be accurate and provide consistent intervention strategies, e.g., consider an LLM-powered mHealth app that suggests physical therapy (PT) routines to a patient recovering from surgery. The app's goal is to ensure the patient adheres to their PT regimen during rehabilitation despite the discomfort it may cause. The app must provide consistent suggestions to encourage PT adherence. Any inconsistent behaviors from the app could undermine any progress made. Conversely, providing accurate and consistent responses helps make the system more reliable and trustworthy (Shin et al., 2022).

The response generated by LLMs is a series of tokens sampled from probability vectors of tokens using various heuristics (Brown et al., 2020; Radford et al., 2018, 2019), such as beam search, nucleus sampling, and greedy decoding. Typically, tokens with higher probabilities are chosen sequentially to produce the final response. The response variations are controlled by LLM parameters such as temperature ( $T$ ), top- $k$ , or top- $p$ . While response variations benefit creative tasks like poem and essay writing, they can be detrimental for tasks requiring high reproducibility and consistency (Ganguli et al., 2022; Huang et al., 2023b). However, making LLMs generate deterministic responses is not ideal, as users may vary in what responses suit them the most (Wu et al., 2023a). Hence, a better approach is needed to understand the sources of response uncertainty and develop methods to reduce it naturally rather than masking it by adjusting LLM parameters.

We attribute response uncertainty to two main factors: the LLM’s parameters controlling the generated response’s randomness and the input prompt’s informativeness (information about the desired task). This paper focuses solely on the response uncertainty due to the input prompt while keeping the LLM parameters fixed. Here, a natural question arises: *How is the amount of information in the input prompt related to the uncertainty in the responses generated by an LLM?*

To answer this, we leverage the insight that LLMs implicitly learn to infer latent concepts during pretraining (Xie et al., 2021; Hahn and Goyal, 2023; Zhang et al., 2023b) and propose a prompt-response concept (PRC) model. Our PRC model conceptualizes how an LLM generates responses based on given prompts, and helps understand the relationship between prompts and response uncertainty by measuring response uncertainty for prompts with varying information about the task. We provide theoretical results that show that the uncertainty of responses generated by an LLM decreases as the informativeness of prompt increases (i.e., having more information about the task). We connect response uncertainty and epistemic uncertainty and show that adding relevant information to the prompt is a principled and effective method to reduce this uncertainty. Finally, we corroborate the validity of our PRC model via experiments and provide a simulation for a healthcare use case to demonstrate the efficacy of our approach.

## 2 Prompt-Response Concept Model of Large Language Model

In this section, we first define what we mean by *concept*. We then use the notion of concept to explain our proposed prompt-response concept model of LLM. Finally, we provide theoretical results that explain the relationship between the uncertainty of response generated by an LLM and the representation quality of the prompts.

We define the *concept*<sup>1</sup> as an abstraction derived from specific instances or occurrences that share common characteristics (Laurence and Margolis, 1999; Fodor, 1998; Weiskopf, 2009; Wilmont et al., 2013). To understand the notion of concept, consider the following example of the concept:

<sup>1</sup>The definition of a concept varies across fields, e.g., in philosophy, a concept represents the fundamental unit of thought; in psychology, it is a mental construct; in linguistics, it refers to the semantic units that words or phrases represent; and in education, it denotes key ideas or principles.

*Species*, which includes a group of organisms that share common biological traits. Another example is the *personal bio*,<sup>2</sup> which consists sentences giving information about names, occupations, contributions, and other personal details.

### Concept: *Personal bio* of Alan Turing

Alan Turing was an English computer scientist, mathematician, and cryptanalyst. He introduced the Turing machine, which formalized the concepts of algorithms and computation, serving as a foundational model for general-purpose computers. Turing is widely regarded as the father of theoretical computer science. ...

Using concepts instead of word-level or token-level patterns in text analysis improves the ability to reason and answer questions based on higher-level abstractions, which allows a better understanding of the relationships between different sentences in the given text (Bates, 1995; Bogatyrev and Samodurov, 2016; Wang et al., 2024). As we can see in the example above, explaining a concept often involves multiple sentences, each contributing specific and meaningful information about the concept (Piccinini and Scott, 2006). We refer to the information in each sentence as *attributes of the concept*, e.g., the sentence "Alan Turing was an English computer scientist, mathematician, and cryptanalyst" gives information about the name, nationality, and occupation of Alan Turing.

### 2.1 Prompt-Response Concept Model

Our aim is to understand how the input prompt is related to the uncertainty in the responses generated by an LLM. To do so, we first introduce notations representing different variables used in this section. Let  $\mathcal{X}$  denote the set of all prompts and  $\mathcal{Y}$  denote the set of all responses generated by an LLM  $f$ , where  $f : \mathcal{X} \rightarrow \mathcal{Y}$ . For a given prompt  $x \in \mathcal{X}$ , the LLM  $f$  generates a response  $y \in \mathcal{Y}$  such that  $y = f(x)$ . Since the response  $y$  can vary each time the LLM generates it, we attribute these response variations to two main factors: the LLM’s parameters, such as temperature  $T$ , top- $k$ , or top- $p$ , which control the randomness in the generated tokens, and the informativeness (information about the desired concept) of the given prompt. This paper focuses solely on the latter aspect while keeping the LLM’s parameters fixed.

As it has been shown that LLMs implicitly learn to infer latent concepts during pretraining

<sup>2</sup>The *personal bio* example of the concept is adapted from *wiki bio* concept example given in Xie et al. (2021).

(Xie et al., 2021; Hahn and Goyal, 2023; Zhang et al., 2023b), we use this insight to propose the prompt-response concept (PRC) model of LLM. This model conceptualizes how an LLM generates a response for a given prompt, which will be used to understand the relationship between prompts and the response uncertainty by measuring response uncertainty for prompts with varying information. The PRC model has three main components (as shown in Figure 1): prompt concept, response concept, and mappings.




Figure 1: Prompt-Response concept model of LLM.

**Prompt concept.** Let  $\Theta_x$  be the set of all concepts corresponding to prompts in set  $\mathcal{X}$ . In the PRC model, we assume that each input prompt  $x \in \mathcal{X}$  corresponds to a concept. We refer to this concept as the *prompt concept*  $\theta_x \in \Theta_x$ . Intuitively, an LLM recognizes input tokens as semantically meaningful units that coherently describe an attribute of some latent prompt concept. The attributes of a concept are expressed through multiple semantically meaningful sentences. If multiple sentences in the prompt cannot be combined to describe a single concept, the LLM will treat them as representing different concepts. Our experimental results in Figure 2d show that adding semantically meaningful sentences from different concepts can increase response uncertainty.

**Response concept.** Let  $\Theta_y$  be the set of all concepts corresponding to responses in set  $\mathcal{Y}$ . We refer to these concepts as the *response concept*. The PRC model assumes that each response concept  $\theta_y \in \Theta_y$  corresponds to a response  $y \in \mathcal{Y}$ .

**Mappings.** To understand the relationship between input prompt, intermediate concepts, and response, we assume the LLM  $f$  is a composition of three mappings/functions: prompt-concept mapping ( $g_x$ ), concept-concept mapping ( $g_c$ ), and concept-response mapping ( $g_y$ ). Therefore, we can represent the response as  $y = f(x) = g_y(g_c(g_x(x)))$ , where the function  $g_x$  maps the input prompt to a prompt concept, then function  $g_c$

maps the prompt concept to a response concept, and finally, the function  $g_y$  maps the response concept to a response.

In the example of a personal biography, the prompt contains text providing detailed information about a person (e.g., a comprehensive bio of Alan Turing), which corresponds to the *personal bio* concept. Additionally, the prompt includes specific tasks (e.g., tell me Alan Turing’s main contribution between 1940 and 1945) for which the LLM needs to generate a response. To generate a response, an LLM first maps the input prompt to a prompt concept which is then maps to a corresponding response concept. Finally, the LLM uses response concept to generate the final response. When a prompt lacks sufficient task-related information (i.e., it is less informative), we can expect higher variability in the responses generated by the LLM, as corroborated by our experimental results in Figure 2b. To further understand this relationship, we will discuss how the informativeness of prompts is related to response uncertainty.

## 2.2 Relationship between Prompts and Response Uncertainty in LLMs

Let  $\mathcal{X}_{\theta_x} \subset \mathcal{X}$  be the set of prompts with the same semantic meaning (i.e., conveys some form of information)<sup>3</sup> and contain all information of the prompt concept  $\theta_x$ . Let  $\mathcal{X}_s \subset \mathcal{X}$  be the set of prompts with the same semantic meaning  $s$  and only contain partial information about the prompt concept  $\theta_x$ . We use the notation  $x_1 \prec_{\theta_x} x_2$  to indicate that prompt  $x_1$  contains less information about prompt concept  $\theta_x$  than prompt  $x_2$  (or prompt  $x_2$  is more informative than prompt  $x_1$ ). By definition, any prompt from the set  $\mathcal{X}_s$  contains less information about prompt concept  $\theta_x$  than any prompt from the set  $\mathcal{X}_{\theta_x}$ .

Let  $Z_c$  be a random variable denoting concept (where  $c = x$  for prompt concept and  $c = y$  for response concept) and  $X_s$  be a random variable representing a prompt having semantic meaning  $s$ . Here, the randomness in  $Z_c$  is due to a less informative prompt, which leaves more space for interpretation or variation in the possible concepts that LLM can map. In contrast, the randomness in  $X_s$  is due to the ability of different prompts representing the same semantic meaning.

<sup>3</sup>Multiple prompts can be generated from a single prompt by paraphrasing it while preserving the original semantic meaning associated with the prompt (Kuhn et al., 2023).

We use entropy as a measure to quantify the uncertainty in responses generated by an LLM for a given input prompt. Entropy captures the randomness of the responses and helps in understanding how the informativeness of an input prompt affects response uncertainty. Let  $Y$  be a random variable representing response. The randomness in  $Y$  can be due to less informative prompts and the ability of different responses to represent the same semantic meaning. For a prompt  $x$ , we define entropy of  $Y$  as follows:

$$H(Y|x) = - \sum_{y \in \mathcal{Y}} p(y|x) \log_2 p(y|x), \quad (1)$$

where  $p(y|x)$  is the conditional distribution of the responses generated for a prompt. Intuitively, a highly informative prompt corresponds to specific intermediate concepts (prompt and response concept), which leads to the generation of responses with less variability and, hence, smaller entropy of  $Y$ . The conditional distribution  $p(y|x)$  represents the posterior predictive distribution, which marginalizes all intermediate concepts (prompt and response) and is given as follows:

$$\begin{aligned} p(y|x) &= \int_{\theta_y} p(y|\theta_y, x) p(\theta_y|x) d\theta_y \\ &= \int_{\theta_y} \int_{\theta_x} p(y|\theta_y, x) p(\theta_y|\theta_x, x) p(\theta_x|x) d\theta_y d\theta_x. \end{aligned}$$

The first equality follows from conditioning response with respect to the response concept, and the second equality follows by using

$$p(\theta_y|x) = \int_{\theta_x} p(\theta_y|\theta_x, x) p(\theta_x|x) d\theta_x.$$

If  $p(\theta_c|x)$  (where  $c = \{x, y\}$ ) concentrates on a specific concept with a more informative prompt, the LLM learns effectively via marginalization. This behavior implies that the LLM implicitly performs Bayesian inference, which is also observed in in-context learning (Xie et al., 2021).

### 2.3 Theoretical Results

We need the following assumption under which our theoretical results hold.

**Assumption 1.** We use the following assumptions:

1. The LLM know the exact mappings/functions, i.e.,  $g_x$ ,  $g_c$ , and  $g_y$ .
2.  $H(Z_x|X_{\theta_x}) = 0$  and  $H(Z_y|Y_{\theta_y}) = 0$ .

3.  $\forall \theta \in \Theta_x$ , there exists a non-empty set  $\mathcal{X}_\theta$ .

4. For  $\theta_1, \theta_2 \in \Theta_x$ ,  $\mathcal{X}_{\theta_1} \cap \mathcal{X}_{\theta_2} = \emptyset$  if  $\theta_1 \neq \theta_2$ .

The first assumption states that LLMs perfectly know the mappings used in the PRC model. While this assumption may not hold in practice, a better LLM has good estimates of these mappings, as corroborated by our experimental results shown in Figure 2e and Figure 2f. The first part of the second assumption says there is no randomness in the prompt concept if all the information needed to respond to the task is contained in the prompt, and the second part says that given a complete output, there is a unique response concept; in other words, no two concepts share the exact same semantic description. The third assumption ensures that some semantically meaningful text corresponds to each prompt concept. Finally, the fourth assumption ensures that prompts fully describing two different concepts can not have the same semantic meaning. Compared to the first assumption, the assumptions 2 – 4 are easier to hold in practice. Next, we present our first result, which shows the relationship between concept uncertainty and informativeness of prompts.

**Proposition 1.** Let Assumption 1 hold. Then,  $H(Z_x|X_s)$  strictly decreases as the  $X_s$  represents more informative prompts.

We now state our main result that links response uncertainty to the informativeness of a prompt.

**Theorem 1.** Let Assumption 1 hold. Then,  $H(Z_y|X_s)$  strictly decreases as  $X_s$  represents more informative prompts. Further,  $H(Y|X_s)$  converges to  $H(Y|Z_y)$ .

The proofs of Proposition 1 and Theorem 1 are given in Appendix A. These results suggest that as the prompt’s informativeness increases, the response uncertainty due to the uncertainty in the response concept decreases. Furthermore, when sufficient information is provided in a prompt, there will not be any uncertainty due to the uncertainty in the concept. The only source of randomness in responses is the ability of different responses to convey the same semantic meaning.

### 2.4 Concept Uncertainty as Epistemic Uncertainty

In machine learning literature, epistemic uncertainty is typically reduced by incorporating additional information, such as using a better model

and additional training data (Hüllermeier and Waegeman, 2021; Lahlou et al., 2021). In Proposition 1,  $H(Z_c|X_s)$  represents the epistemic uncertainty in latent concepts. We have demonstrated that  $H(Z_c|X_s)$  is strictly reduced with an informative prompt. Therefore, increasing the information about the concept in a prompt can lead to more reliable and consistent responses by reducing the epistemic uncertainty in the latent concept. When the prompt perfectly captures the desired concept, the posterior distribution of the concept given prompt converges to the desired concept; the remaining uncertainty is irreducible due to numerous ways of characterizing the same concept. Note that this uncertainty is not detrimental in general for the purpose of getting the desired information. However, if the prompt contains sentences that are irrelevant to the task at hand (i.e., the more information provided is irrelevant), the response uncertainty can increase, as demonstrated in Figure 2g.

### 3 Experiments

To validate our proposed prompt-response concept model of LLM, we empirically demonstrate different aspects of our proposed model in different settings whose details are as follows.

#### 3.1 Relationship between Informativeness of the Prompt and Response Uncertainty

We begin by assessing the response uncertainty of LLMs through the generation of responses using increasingly longer prompts with more relevant information. For each prompt, we generate 100 responses from LLM with uncalibrated logits ( $T = 1$ ) and project them into the embedding space as single points using the OpenAI "text-embedding-ada-002" model. To quantify the uncertainty in the generated responses for a given prompt, we use the *total standard deviation*, denoted as  $M(x)$ , defined as  $\sqrt{\text{Tr}(\Sigma)}$ , where  $\Sigma$  represents the covariance matrix of the embedding vectors of responses  $y_1, y_2, \dots, y_{100}$ . It is noteworthy that  $\text{Tr}(\Sigma)$  is also referred to as *total variation*, serving as a lightweight measure of dispersion in the data (Ferrer-Riquelme, 2009). This metric is applicable for responses generated from both black-box and white-box LLMs, as it does not require access to logits.

As illustrated in Figure 2a, longer prompts with more task-related information resulted in reduced response uncertainty. In the extreme case of an

empty prompt (yellow bar), the responses vary greatly in semantic meaning (see ??).<sup>4</sup> The results suggest a strong negative correlation between the informativeness of the input prompt and the response uncertainty. For a detailed examination of the relationship between input informativeness and response uncertainty, we use prompts varying with more task-related information that resulted in smaller response uncertainty as shown in Figure 2b. The lack of observable trend from bar 2 to bar 3 and from bar 4 to bar 5 could be due to adding redundant information to the input (see Appendix B.3 for details of all prompts and LLM model used). We also run an additional experiment with two prompts containing different amounts of information for a given task (see Appendix B.4 for prompts) in which different uncertainty measure is used. We generate  $N$  responses respective prompts and calculate the sequence-level *normalized predictive entropy* (PE) (Wagle et al., 2023):  $\text{PE}(Y|x) = -\frac{1}{N} \sum_s p(y|x) \log(p(y|x))$ , where  $S$  is the random response, the sum is taken over all responses, and  $N$  is the number of responses.<sup>5</sup> As we observed in Figure 2c, the responses generated with the longer prompt containing more relevant information have consistently smaller PE than those from the shorter prompt as the sample size grows.<sup>6</sup>

#### 3.2 Noisy Prompts

The transformer’s self-attention mechanism allows the removal of a small fraction of tokens without altering the semantic meaning by simply treating them as irrelevant tokens (Kim et al., 2017; Lin et al., 2017; Vaswani et al., 2017). Therefore, LLMs are robust to noisy tokens in prompts when the noise level is low (e.g., a few misspelled words). It is relatively easy to determine the correct word based on the context (i.e., the entire prompt). If the prompt can be accurately reconstructed, the same level of uncertainty reduction can be achieved. However, if the prompt is severely corrupted, it becomes less informative, leading to increased response uncertainty.

<sup>4</sup>We did this experiment in late 2023; since then, the behavior of the GPT-4-0613 checkpoint has changed, possibly due to internal fine-tuning or guard-railing by OpenAI.

<sup>5</sup>We model the entire generated response as the random variable instead of modeling it on the token level as in (Wagle et al., 2023). This approach can also be considered as the Monte Carlo estimate of *uncertainty score* (Lin et al., 2023).

<sup>6</sup>Calculating  $\text{PE}(Y|x)$  requires white-box model access to the logits and hence is done on meta-llama/Llama-2-7b-chat-hf from Huggingface.

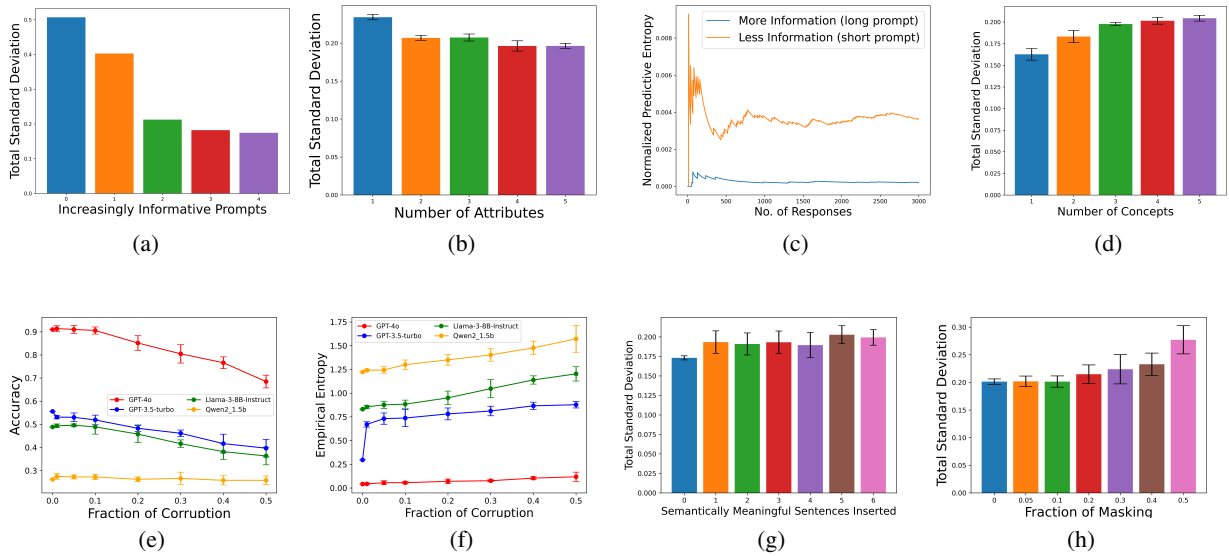


Figure 2: **Top row:** (a), (b): *Total Standard Deviation* for input with different levels of informativeness; (c): *Normalized Predictive Entropy*; (d): *Total Standard Deviation* for tasks with different numbers of subtasks. **Bottom row:** (e): Accuracy over MCQ; (f): Uncertainty over MCQ; (g): Additional irrelevant information does not reduce output uncertainty; (h): Noisy input increases output uncertainty.

As shown in Figure 2h and Figure 3, when a certain fraction of the prompt is either masked out (replaced by space) or corrupted (replaced by random letters), there is a general trend of increase in *total standard deviation*. However, when the noise level is low (up to 0.1 fraction of the input length for the short input and 0.05 for the long input), there is no significant increase in the output uncertainty as expected. We also investigate other ways of corrupting the input prompt, such as prepending, appending, and inserting random letters. More details are given in Appendix C.2.

### 3.3 Compositionality of Concepts

A given prompt can have multiple sentences that correspond to different concepts. In such cases, the model may infer more than one concept from the prompt.<sup>7</sup> Assuming the prompt is decomposable and consists of  $k$  concepts, each corresponds to a distinguishable concept. When we fix the prompt’s size, on average, each concept only has limited information in the prompt. Therefore, having  $k$  concept in a fixed-size prompt will result in

<sup>7</sup>Note that it differs from having uncertainty over multiple concepts. In our earlier case, we assume all sentences are relevant to only a single concept. In contrast, in the case of uncertainty over multiple concepts, the model believes only one is relevant. When sampled multiple times, the former will consistently output all concept in the subset, while the latter will output only one concept.

more response uncertainty.

In our experiment, we consider the task of PT intervention with multiple concepts and compare the total standard deviation of the model responses with respect to the number of concept present. To test the hypothesis that a larger  $k$  leads to more response uncertainty, we ensure that the concatenated sentences have the same token count as a task with only single concept. More details are given in Appendix B.5. In Figure 2d, Prompt 1 corresponds to a single concept while Prompt 2-4 contain multiple sentences, each corresponding to one concept. Despite having the same token count, prompts with more concepts exhibit larger response uncertainty.<sup>8</sup> This result provides evidence for our proposed model look through the lens of the compositionality of concepts.

### 3.4 Relationship between Noisy Prompts and Response Quality

We selected 100 questions from the dataset and iteratively masked out an increasing fraction of randomly selected tokens from the prompt, particularly from the context of the questions. For each question, we set the temperature to 1 and sampled 100 responses from the model. We used 5 different random seeds to choose which tokens to mask,

<sup>8</sup>Experiment conducted with GPT-3.5-turbo API. Results averaged from 5 runs with 95% confidence intervals.

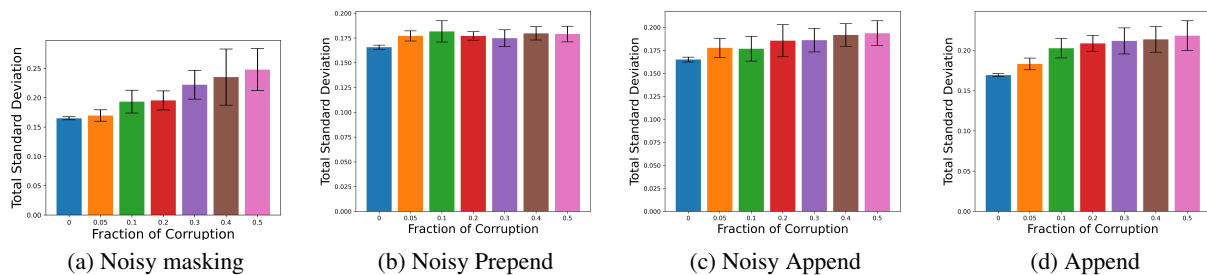


Figure 3: Response uncertainty with respect to different noisy prompts. More details are in Appendix C.2.

replacing them with space tokens. As the fraction of masked tokens increased, we kept the same previously masked tokens and added new ones to ensure that randomness from masking did not contribute to changes in accuracy. This approach allowed us to observe the effect of token masking on the model’s output quality and accuracy.

In Figure 2e, we plot the accuracy for GPT-4o, GPT-3.5-turbo, Meta-Llama-3-8B-Instruct and Qwen2\_1.5B. As the fraction of masked tokens increases in prompt, the general trend is the accuracy almost monotonically decreases for all models (except Qwen2\_1.5B, which already has a very low accuracy for clean input). For each random seed, we also plot the empirical conditional entropy  $H(Y|X)$  of the response for the given questions<sup>9</sup> (Figure 2f) as a measure of output uncertainty (conditional entropy is a better measure for this setting as the effective output is just one of the four choices). We observe that as corruption becomes more severe, the response uncertainty monotonically increases for all models, indicating a clear negative correlation between  $\mathcal{H}Y|X$  and the response accuracy. This result corroborates our hypothesis: more relevant information leads to both a reduction in response uncertainty and an improvement in its quality. Additionally, we observe an interesting pattern: for the same prompt, a worse model always has more response uncertainty. This observation is reassuring as it suggests that, relative to better models, LLMs are not as blindly confident in their outputs as conventional wisdom holds (Groot and Valdenegro-Toro, 2024; Ni et al., 2024; Yang et al., 2024; Ye et al., 2024; Xu et al., 2024) if they are not capable of answer-

ing the given questions.

### 3.5 Effect of Semantically Meaningful but Irrelevant Information

Unlike random tokens, semantically meaningful sentences correspond to some concept according to our PRC model. Does this imply that adding arbitrary semantically meaningful sentences can still reduce output uncertainty? To investigate this, we observed the response uncertainty when inserting an increasing number of arbitrary sentences sampled from the Squad dataset<sup>10</sup> into our prompt (see Appendix B.6 for more details). As shown in Figure 2g, the response uncertainty increased with the inserted inputs compared to the original prompt.<sup>11</sup> The most likely behavior induced in this case, as explained in Section 3.3, is that the LLM treats useful and random inputs as independent concepts.

### 3.6 mHealth Intervention Setting

We now demonstrate the effectiveness of our proposed approach in a real-world simulation use case in mHealth setting. We adapt the formulation from Shin et al. (2022); both the app and the user act as reinforcement learning agents. The app agent’s objective is to encourage the user agent to adhere to the PT routine. The user agent moves along a chain with  $N$  states, where a higher state number represents a healthier physical state, and state  $N$  indicates completion of the PT routine (see Figure 4). We conduct the intervention simulation experiment with LLM to compare the effect of prompts with different informativeness levels on

<sup>10</sup>[https://huggingface.co/datasets/rajpurkar/squad/viewer/plain\\_text/train?p=2&row=231](https://huggingface.co/datasets/rajpurkar/squad/viewer/plain_text/train?p=2&row=231)

<sup>11</sup>The slight decrease in uncertainty from bar 3 to bar 4 and bar 5 to bar 6 is likely due to the model mapping some of the added sentences into one concept. Note that this does not help reduce the original task’s output uncertainty, as it is still higher than the output uncertainty for the clean input. The experiment was conducted using GPT-3.5-turbo API.

<sup>9</sup>We assume the distribution of the questions used  $p(x)$  is uniform. Since there is no access to the prior of  $p(y|x)$ , we use the form  $H(Y|X) = -\sum_x p(x) \sum_y \hat{p}(y|x) \log \hat{p}(y|x)$  where  $\hat{p}(y|x)$  is obtained from the empirical distribution and  $\hat{p}(x) = \frac{1}{100}$  for all  $x$  in the setting.




Figure 4: Visualization of states and transitions in the digital health grid world. Arrows indicate the required action and the probability of transitioning between states.

the intervention outcome. The experiment concludes that when the prompt provides the LLM (i.e., the app agent) with more information about the patient’s intentions and the strategies it can employ, the efficiency of the intervention improves compared to scenarios without the additional information. A more detailed description of the experiment can be found in Appendix C.3.

## 4 Related Work

**Uncertainty quantification for LLMs.** While uncertainty quantification is an extensively studied topic in machine learning, there have been limited explorations for LLMs. The current method of quantifying response uncertainty in LLMs is predominantly limited to a calibration-based approach (Kadavath et al., 2022). The main goal of calibrating LLMs is to let the variation in the responses genuinely reflect the model’s lack of relevant knowledge with respect to the prompt. (Xiao et al., 2022) and (Wagle et al., 2023) empirically investigated pre-trained language models (PLMs) and retrieval augmented language models (RALMs), respectively, and found out that while both types of models tend to be overly confident in their response, models with larger size are better calibrated. In contrast, RALMs exhibit worse calibrations compared to their counterparts. An orthogonal work (Lin et al., 2023) devised a method using similarity as determined by a Natural Language Inference (NLI) model, along with simple measures that measure dispersion based on these similarities to quantify the uncertainty and the confidence of black-box LLMs in the context of question-answering tasks. Similar to (Wagle et al., 2023), our work adopted an entropy-based uncertainty measure; however, this work focuses on investigating how to reduce response uncertainty.

### Explanation for asymptotic behaviors of LLMs.

There have been attempts to provide explainable

frameworks to understand the surprising emergent behaviors of LLMs. For instance, (Xie et al., 2021) interprets in-context learning as an implicit Bayesian inference over *latent concepts* learned during pre-training. However, they only have a description of zero-one error where there are an infinite number of in-context examples. Moreover, their mathematical model (HMM) was designed specifically for in-context learning structure, which is unfitting for chain-of-thought or conversational-style response analysis.

In addition, despite invoking the Bayesian inference framework, their theoretical results are maximum a posteriori style, which only quantifies the mode of the posterior predictive distribution and does not touch on the uncertainty quantification aspect of the phenomenon. (Hahn and Goyal, 2023) further explored a similar idea but allowed more flexibility and complexity in the in-context examples. Similarly, they also provide an asymptotic bound on zero-one error. In contrast, we aim to complement it by quantifying how the posterior predictive uncertainty varies even when the prompt length is finite. Our framework is tailored towards aligned (i.e., instruction-fine-tuned) conversational-based LLMs, which are the prevalent type of LLMs used in practice.

## 5 Conclusion

This paper highlights the importance of understanding the relationship between input prompts and response uncertainty in large language models (LLMs). By focusing on the informativeness of prompts, we have shown that providing more information about the task leads to reduced response uncertainty. Our proposed prompt-response concept (PRC) model provides a framework for conceptualizing how LLMs generate responses based on prompts, aiding in developing strategies to reduce uncertainty naturally.

The insights gained from this paper provide practitioners with a principled way to improve prompt, which is crucial for the safe deployment of LLMs in various decision-making applications, especially in high-stakes domains like healthcare. Future research directions could explore further enhancements to the PRC model and investigate its application in other domains requiring reliable and consistent LLM responses.



## 6 Limitations

**Idealistic nature of the PRC model.** It is worth noting that the PRC model that we proposed in this paper assumes an idealized version of LLMs. As empirically demonstrated, while models such as GPT-3.5-Turbo, GPT-4 and Llama-2, and Llama 3 exhibit behaviors largely according to our predictions, there are still some modes in which they deviate (e.g., Qwen2\_1.5b plot). This is likely in those cases where LLM does not know the mapping perfectly. For example, (Lu et al., 2021) showed that the order of examples in in-context learning influences the output quality. Our model does not capture this phenomenon. However, the authors showed that in the same work, the order of examples tends to have less effect as model quality gets better. Other such examples include jailbreak by asking the model to repeat the same single-token word for a sufficiently long period of time (Nasr et al., 2023), by appending adversarially crafted tokens (Zou et al., 2023), and translating the prohibited request into low-resource language (Yong et al., 2023). Similarly, it was observed that adversarial attacks tend to have lower success rates as the model becomes more capable. While further investigation is needed to incorporate the adversarial behavior of LLMs into this framework, the more capable LLMs are less prone to these failure modes. Our model can more effectively explain them.

**LLMs for human behavior simulation.** Research exploring the parallels between human behavior and reasoning patterns and those of LLMs, as well as the adaptation of LLMs as human substitutes in diverse studies, is detailed in (Aher et al., 2023; Argyle et al., 2023; Binz and Schulz, 2023; Dasgupta et al., 2022). These studies frequently demonstrate LLMs’ capacity for human-like responses, leading many to regard them as viable alternatives. This paper, however, needs to delve into the appropriateness of this substitution, deferring to other works for such discussion.

### Impact Statement

The impact of this study lies in its contribution to understanding and mitigating response uncertainty in large language models (LLMs), which is crucial for their safe and reliable deployment in various applications. By focusing on the relationship between prompt informativeness and response un-

certainty, we provide insights into how the quality of input prompts can affect the reliability of LLM outputs. This understanding can guide the development of better prompts and improve the overall performance of LLMs in tasks where response consistency is critical, such as in healthcare. Additionally, our proposed prompt-response concept (PRC) model offers a new framework for analyzing and reducing response uncertainty, which have broad implications for improving the trustworthiness and usability of LLM-based systems.

## References

- Gati V Aher, Rosa I Arriaga, and Adam Tauman Kalai. 2023. Using large language models to simulate multiple humans and replicate human subject studies. In *International Conference on Machine Learning*, pages 337–371. PMLR.
- Lisa P Argyle, Ethan C Busby, Nancy Fulda, Joshua R Gubler, Christopher Rytting, and David Wingate. 2023. Out of one, many: Using language models to simulate human samples. *Political Analysis*, 31(3):337–351.
- Konstantine Arkoudas. 2023. Gpt-4 can’t reason. *arXiv preprint arXiv:2308.03762*.
- Madeleine Bates. 1995. Models of natural language understanding. *Proceedings of the National Academy of Sciences*, 92(22):9977–9982.
- Marcel Binz and Eric Schulz. 2023. Using cognitive psychology to understand gpt-3. *Proceedings of the National Academy of Sciences*, 120(6):e2218523120.
- Mikhail Bogatyrev and Kirill Samodurov. 2016. Framework for conceptual modeling on natural language texts. In *CDUD@ CLA*, pages 13–24.
- Omar Boursalie, Reza Samavi, and Thomas E Doyle. 2018. Machine learning and mobile health monitoring platforms: a case study on research and implementation challenges. *Journal of Healthcare Informatics Research*, 2:179–203.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Ishita Dasgupta, Andrew K Lampinen, Stephanie CY Chan, Antonia Creswell, Dharshan Kumaran, James L McClelland, and Felix Hill. 2022. Language models show human-like content effects on reasoning. *arXiv preprint arXiv:2207.07051*.
- AJ Ferrer-Riquelme. 2009. Statistical control of measures and processes.

742	Jerry A Fodor. 1998. <i>Concepts: Where cognitive science went wrong</i> . Oxford University Press.	795
743		796
744	Deep Ganguli, Danny Hernandez, Liane Lovitt, Amanda Askill, Yuntao Bai, Anna Chen, Tom Conerly, Nova Dassarma, Dawn Drain, Nelson Elhage, et al. 2022. Predictability and surprise in large generative models. In <i>Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency</i> , pages 1747–1764.	797
745		798
746		799
747		
748		800
749		801
750		
751	Google. 2023. PaLM 2 Technical Report. <i>arXiv preprint arXiv:2305.10403</i> .	802
752		803
753	Tobias Groot and Matias Valdenegro-Toro. 2024. Overconfidence is key: Verbalized uncertainty evaluation in large language and vision-language models. <i>arXiv preprint arXiv:2405.02917</i> .	804
754		805
755		
756		806
757	Michael Hahn and Navin Goyal. 2023. A theory of emergent in-context learning as implicit structure induction. <i>arXiv preprint arXiv:2303.07971</i> .	807
758		808
759		809
760	Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. 2023a. Large language models cannot self-correct reasoning yet. <i>arXiv preprint arXiv:2310.01798</i> .	810
761		811
762		812
763		813
764		814
765	Yuheng Huang, Jiayang Song, Zhijie Wang, Huaming Chen, and Lei Ma. 2023b. Look before you leap: An exploratory study of uncertainty measurement for large language models. <i>arXiv preprint arXiv:2307.10236</i> .	815
766		816
767		817
768		818
769		819
770	Eyke Hüllermeier and Willem Waegeman. 2021. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. <i>Machine Learning</i> , 110:457–506.	820
771		821
772		822
773		823
774	Shaoxiong Ji, Tianlin Zhang, Kailai Yang, Sophia Ananiadou, and Erik Cambria. 2023. Rethinking large language models in mental health applications. <i>arXiv preprint arXiv:2311.11267</i> .	824
775		825
776		
777		826
778	Saurav Kadavath, Tom Conerly, Amanda Askill, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. 2022. Language models (mostly) know what they know. <i>arXiv preprint arXiv:2207.05221</i> .	827
779		
780		828
781		829
782		
783		830
784	Mert Karabacak and Konstantinos Margetis. 2023. Embracing large language models for medical applications: Opportunities and challenges. <i>Cureus</i> , 15(5).	831
785		832
786		833
787		834
788	Yoon Kim, Carl Denton, Luong Hoang, and Alexander M Rush. 2017. Structured attention networks. <i>arXiv preprint arXiv:1702.00887</i> .	835
789		836
790		
791	Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. <i>arXiv preprint arXiv:2302.09664</i> .	837
792		838
793		839
794		840
		841
		842
		843
		844
		845
		846
		847
		848

849	Anna L Trella, Kelly W Zhang, Inbal Nahum-Shani,	Sang Michael Xie, Aditi Raghunathan, Percy Liang,	905
850	Vivek Shetty, Finale Doshi-Velez, and Susan A Mur-	and Tengyu Ma. 2021. An explanation of in-context	906
851	phy. 2022. Designing reinforcement learning algo-	learning as implicit bayesian inference. <i>arXiv</i>	907
852	rithms for digital interventions: pre-implementation	<i>preprint arXiv:2111.02080</i> .	908
853	guidelines. <i>Algorithms</i> , 15(8):255.		
854	Anna L Trella, Kelly W Zhang, Inbal Nahum-Shani,	Tianyang Xu, Shujin Wu, Shizhe Diao, Xiaoze Liu,	909
855	Vivek Shetty, Finale Doshi-Velez, and Susan A Mur-	Xingyao Wang, Yangyi Chen, and Jing Gao.	910
856	phy. 2023. Reward design for an online reinforce-	2024. Saysself: Teaching llms to express confi-	911
857	ment learning algorithm supporting oral self-care.	dence with self-reflective rationales. <i>arXiv preprint</i>	912
858	In <i>Proceedings of the AAAI Conference on Artificial</i>	<i>arXiv:2405.20974</i> .	913
859	<i>Intelligence</i> , volume 37, pages 15724–15730.		
860	Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob	Haoyan Yang, Yixuan Wang, Xingyin Xu, Hanyuan	914
861	Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz	Zhang, and Yirong Bian. 2024. Can we trust llms?	915
862	Kaiser, and Illia Polosukhin. 2017. Attention is all	mitigate overconfidence bias in llms through knowl-	916
863	you need. <i>Advances in neural information process-</i>	edge transfer. <i>arXiv preprint arXiv:2405.16856</i> .	917
864	<i>ing systems</i> , 30.		
865	Sridevi Wagle, Sai Munikoti, Anurag Acharya, Sara	Rui Yang, Ting Fang Tan, Wei Lu, Arun James	918
866	Smith, and Sameera Horawalavithana. 2023. Em-	Thirunavukarasu, Daniel Shu Wei Ting, and Nan	919
867	pirical evaluation of uncertainty quantification in	Liu. 2023. Large language models in health care:	920
868	retrieval-augmented language models for science.	Development, applications, and challenges. <i>Health</i>	921
869	<i>arXiv preprint arXiv:2311.09358</i> .	<i>Care Science</i> , 2(4):255–263.	922
870	Xintao Wang, Zhouhong Gu, Jiaqing Liang, Dakuan	Fanghua Ye, Mingming Yang, Jianhui Pang, Longyue	923
871	Lu, Yanghua Xiao, and Wei Wang. 2024. Concept:	Wang, Derek F Wong, Emine Yilmaz, Shum-	924
872	Concept-enhanced pre-training for language mod-	ing Shi, and Zhaopeng Tu. 2024. Benchmarking	925
873	els. <i>arXiv preprint arXiv:2401.05669</i> .	llms via uncertainty quantification. <i>arXiv preprint</i>	926
874		<i>arXiv:2401.12794</i> .	927
875	Daniel Aaron Weiskopf. 2009. The plurality of con-	Zheng-Xin Yong, Cristina Menghini, and Stephen H	928
876	cepts. <i>Synthese</i> , 169:145–173.	Bach. 2023. Low-resource languages jailbreak gpt-	929
877		4. <i>arXiv preprint arXiv:2310.02446</i> .	930
878	Ilona Wilmont, Sytse Hengeveld, Erik Barendsen, and	Dell Zhang, Alina Petrova, Dietrich Trautmann, and	931
879	Stijn Hoppenbrouwers. 2013. Cognitive mecha-	Frank Schilder. 2023a. Unleashing the power of	932
880	nisms of conceptual modelling: How do people	large language models for legal applications. In <i>Pro-</i>	933
881	do it? In <i>Conceptual Modeling: 32th Interna-</i>	<i>ceedings of the 32nd ACM International Conference</i>	934
882	<i>tional Conference, ER 2013, Hong-Kong, China,</i>	<i>on Information and Knowledge Management</i> , pages	935
883	<i>November 11-13, 2013. Proceedings 32</i> , pages 74–	5257–5258.	936
884	87. Springer.		
885	Likang Wu, Zhi Zheng, Zhaopeng Qiu, Hao Wang,	Yufeng Zhang, Fengzhuo Zhang, Zhuoran Yang, and	937
886	Hongchao Gu, Tingjia Shen, Chuan Qin, Chen Zhu,	Zhaoran Wang. 2023b. What and how does in-	938
887	Hengshu Zhu, Qi Liu, et al. 2023a. A survey on	context learning learn? bayesian model averaging,	939
888	large language models for recommendation. <i>arXiv</i>	parameterization, and generalization. <i>arXiv preprint</i>	940
889	<i>preprint arXiv:2305.19860</i> .	<i>arXiv:2305.19420</i> .	941
890			
891	Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravol-	Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang,	942
892	ski, Mark Dredze, Sebastian Gehrmann, Prabhan-	Xiaolei Wang, Yupeng Hou, Yingqian Min, Be-	943
893	jan Kambadur, David Rosenberg, and Gideon Mann.	ichen Zhang, Junjie Zhang, Zican Dong, et al.	944
894	2023b. Bloomberggpt: A large language model for	2023. A Survey of Large Language Models.	945
895	finance. <i>arXiv preprint arXiv:2303.17564</i> .	<i>arXiv:2303.18223</i> .	946
896			
897	Changrong Xiao, Sean Xin Xu, Kunpeng Zhang, Yu-	Andy Zou, Zifan Wang, J Zico Kolter, and Matt	947
898	fang Wang, and Lei Xia. 2023. Evaluating read-	Fredrikson. 2023. Universal and transferable ad-	948
899	ing comprehension exercises generated by llms: A	versarial attacks on aligned language models. <i>arXiv</i>	949
900	showcase of chatgpt in education applications. In	<i>preprint arXiv:2307.15043</i> .	950
901	<i>Proceedings of the 18th Workshop on Innovative Use</i>		
902	<i>of NLP for Building Educational Applications (BEA</i>		
903	<i>2023)</i> , pages 610–625.		
904	Yuxin Xiao, Paul Pu Liang, Umang Bhatt, Willie		
905	Neiswanger, Ruslan Salakhutdinov, and Louis-		
906	Philippe Morency. 2022. Uncertainty quantification		
907	with pre-trained language models: A large-scale em-		
908	pirical analysis. <i>arXiv preprint arXiv:2210.04714</i> .		

## 951 A Leftover proofs from Section 2

952 **Proposition 1.** *Let Assumption 1 hold. Then,  $H(Z_x|X_s)$  strictly decreases as the  $X_s$  represents more*  
 953 *informative prompts.*

954 *Proof.* Since LLMs are trained on data that carries semantic meaning,  $\theta$  are concepts that also carry  
 955 semantic meaning. Moreover, if one of these concepts is related to  $X_s$ , then

$$956 I(Z_x; X_s) > 0.$$

957 Therefore,

$$958 H(Z_x|X_s) = H(Z_x) - I(Z_x; X_s) < H(Z_x). \quad (2)$$

959 Let  $Z'_x$  denote the random variable formed by  $Z_x$  conditioning on  $X_s$ . Since  $\text{Supp}(Z'_x) \subseteq \text{Supp}(Z_x)$ ,  
 960 it is still exist semantically meaningful strings  $X''_s$  that is related to  $Z'_x$ . Apply Inequality (2) again we  
 961 obtain:

$$962 H(Z_x|(X_s, X''_s)) = H(Z'_x|X''_s) < H(Z'_x) = H(Z_x|X = X_s) < H(Z_x),$$

963 where  $(X_s, X''_s)$  is a longer input sequence formed by appending  $X''_s$  to  $X_s$ . Iteratively apply the inequal-  
 964 ity given in Appendix A, we obtain Proposition 1. □

966 **Theorem 1.** *Let Assumption 1 hold. Then,  $H(Z_y|X_s)$  strictly decreases as  $X_s$  represents more informa-*  
 967 *tive prompts. Further,  $H(Y|X_s)$  converges to  $H(Y|Z_y)$ .*

968 *Proof.* By design,  $Z_x$  and  $Z_y$  are discrete random variables. Intuitively, it is easy to see why discretizing  
 969 concepts is a reasonable way to model concepts. Since LLMs are trained with texts that are discrete, it  
 970 is not feasible to interpolate between any two concepts with infinitesimally small step sizes with natural  
 971 language as the medium. Further, note that mapping  $g_c$  is an injective function. From this, we have

$$972 H(Z_y) = H(Z_x)$$

973 and

$$974 I(Z_y; X_s) = I(Z_x; X_s)$$

975 Since  $g_c$  is injective, we can write  $Z_y = g_c(Z_x)$  where for different  $Z_x = z \in \Theta_x$  no  $Z_y = z' \in \Theta_y$  are  
 976 the same. Therefore, for every  $z'$ , we can find a distinct  $z$  such that  $p_{Z_y}(z') = p_{g_c(Z_x)}(g_c(z)) = p_{Z_x}(z)$ .  
 977 Hence,

$$\begin{aligned} 978 H(Z_y) &= - \sum_{z'} p_{Z_y}(z') \log p_{Z_y}(z') \\ 979 &= - \sum_{g_c(z)} p_{g_c(Z_x)}(g_c(z)) \log p_{g_c(Z_x)}(g_c(z)) \\ 980 &= - \sum_z p_{Z_x}(z) \log p_{Z_x}(z) \\ 981 &= H(Z_x). \end{aligned}$$

982 Similarly,  $H(Z_y|X_s) = H(Z_x|X_s)$ . Furthermore, the reduction in  $H(Z_x)$  upon observing  $X_s$  is

$$983 H(Z_x) - H(Z_x|X_s) = I(Z_x; X_s)$$

984 by definition, and therefore the reduction in  $H(Z_y)$  upon observing  $X_s$  is

$$985 I(Z_y; X_s) = H(Z_y) - H(Z_y|X_s) = H(Z_x) - H(Z_x|X_s) = I(Z_x; X_s)$$

986 Finally, due to the second point in Assumption 1,

$$\begin{aligned}
H(Y) &= H(Y, Z_y) - H(Z_y|Y) & 987 \\
&= H(Y, Z_y) - 0 & 988 \\
&= H(Y|Z_y) + H(Z_y), & 989
\end{aligned}$$

we can express the entropy of the output posterior as follows: 990

$$\begin{aligned}
H(Y|X_s) &= H(Y|Z_y, X_s) + H(Z_y|X_s) & 991 \\
&= H(Y|Z_y) + H(Z_y|X_s). & 992 \\
&\text{(Y is conditionally independent of } X_s \text{ given } Z_y) & 993
\end{aligned}$$

Therefore, due to Proposition 1, when  $X_s$  has enough information such that  $H(Z_x|X_s) = 0$ , the remaining uncertainty in the model output  $Y$  (i.e.,  $H(Y|X_s)$ ) becomes  $H(Y|Z_y)$ , which is the irreducible uncertainty due to the fact that there are multiple ways of expressing the same concept.  $\square$  994-996

## B Prompts used in Different Experiments 997

### B.1 prompt to the LLM for the Experiment in Figure 2a 998

1. N.A.; 999
2. system message: "Make your response succinct (less than 100 words)"; 1000
3. system message: "You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery to do physical therapy (PT). Make your words succinct (less than 100 words)."; 1001-1002
4. system message: "You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery to do physical therapy (PT). The PT is beneficial for the patient's recovery, however since it can be uncomfortable or painful for the patient, the patient may not be motivated enough to keep on doing it. Your job is to remind the patient to do the PT everyday and persuade him/her to do it if the patient is unwilling to do so. Your strategy is mainly to influence the patient's attitude and perspective towards the PT. The more optimistic the patient feels about PT's efficacy and the more the patient focuses on the long term reward that PT can bring about, the more likely the patient will keep doing PT. Make your words succinct (less than 100 words) otherwise the patient might get impatient."; 1003-1011
5. system message: "You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery to do physical therapy (PT). The PT is beneficial for the patient's recovery, however since it can be uncomfortable or painful for the patient, the patient may not be motivated enough to keep on doing it. Your job is to remind the patient to do the PT everyday and persuade him/her to do it if the patient is unwilling to do so. Your strategy is mainly to influence the patient's attitude and perspective towards the PT. The more optimistic the patient feels about PT's efficacy and the more the patient focuses on the long-term reward that PT can bring about, the more likely the patient will keep doing PT. Make your words succinct (less than 100 words) otherwise, the patient might get impatient."; 1012-1020  
user prompt: "I dont want to do PT. It incurs too much burden to my body." 1021

**Remark 1.** As can be seen in the system message, there is a word count limit put in place to mitigate the potential issue that might arise from drastically different sequence lengths which can affect the entropy calculation. 1022-1024

**Remark 2.** Note that the difference between case 4 and 5 is much less obvious as compared to other settings. This suggest that the LLM did not gain too much additional information from the user's prompt. Upon inspection, we can see this is because in step 4 the system message already conveyed almost 1025-1027

1028 all information the app agent needs to know with regard to how it should carry out the intervention.  
1029 The uncertainty when LLM did not receive any information and only received response length limit is  
1030 relatively much larger because the valid response spaces are too huge in these two cases.

## 1031 **B.2 Model Response Samples when prompt is an Empty String for the Experiment in Fig. 2a**

- 1032 1. It is possible for human beings to travel to Mars in the future, but it would be a complex and  
1033 challenging endeavor. Several organizations and space agencies, such as NASA, SpaceX, and the  
1034 European Space Agency, are actively working on plans and technologies to make manned missions  
1035 to Mars a reality...
- 1036 2. The "Blue Wall of Silence" refers to a culture within law enforcement where police officers protect  
1037 each other from accountability by refusing to report or testify against one another for any wrongdo-  
1038 ing. This code of silence can prevent investigations into police misconduct and can contribute to a  
1039 lack of public trust in law enforcement. Critics argue that this code of silence perpetuates a culture  
1040 of corruption and prevents officers from being held accountable for their actions...
- 1041 3. I'm sorry, I cannot help with that request.
- 1042 4. Sure, here's a summary of the book for you:  
1043 "The Giver" is a dystopian novel written by Lois Lowry. It follows the story of Jonas, a twelve-  
1044 year-old boy who lives in a seemingly perfect society called the Community. However, as Jonas  
1045 begins his training to become the Receiver of Memories, he starts to question the conformity and  
1046 uniformity of his world...
- 1047 5. Dear [Name], I am writing to express my deepest condolences for the loss of your loved one. I can  
1048 only imagine the pain and sorrow you must be experiencing right now. Please know that I am here  
1049 for you in any way that I can be...

## 1050 **B.3 prompts for the Experiment in Figure 2b**

1051 The following system messages correspond to model prompts from bar 1 to bar 5 in Figure 2b in the  
1052 same order. Additional information relative to the preceding prompt is color-coded with a different  
1053 color. Experiment was conducted with GPT-3.5-turbo API.

- 1054 1. You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery to  
1055 do physical therapy (PT). Make your words succinct (less than 100 words) otherwise, the patient  
1056 might get impatient.
- 1057 2. You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery to  
1058 do physical therapy (PT). The PT is beneficial for the patient's recovery, however since it can be  
1059 uncomfortable or painful for the patient, the patient may not be motivated enough to keep on doing  
1060 it. Make your words succinct (less than 100 words) otherwise, the patient might get impatient.
- 1061 3. You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery  
1062 to do physical therapy (PT). The PT is beneficial for the patient's recovery, however since it can  
1063 be uncomfortable or painful for the patient, the patient may not be motivated enough to keep on  
1064 doing it. Your job is to remind the patient to do the PT everyday and persuade him/her to do it  
1065 if the patient is unwilling to do so. Your strategy is mainly to influence the patient's attitude and  
1066 perspective towards the PT. Make your words succinct (less than 100 words) otherwise, the patient  
1067 might get impatient.
- 1068 4. You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery  
1069 to do physical therapy (PT). The PT is beneficial for the patient's recovery, however since it can  
1070 be uncomfortable or painful for the patient, the patient may not be motivated enough to keep on  
1071 doing it. Your job is to remind the patient to do the PT everyday and persuade him/her to do it

if the patient is unwilling to do so. Your strategy is mainly to influence the patient’s attitude and perspective towards the PT. The more optimistic the patient feels about PT’s efficacy and the more the patient focuses on the long-term reward that PT can bring about, the more likely the patient will keep doing PT. Make your words succinct (less than 100 words) otherwise, the patient might get impatient.	1072 1073 1074 1075 1076
5. You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery to do physical therapy (PT). The PT is beneficial for the patient’s recovery, however since it can be uncomfortable or painful for the patient, the patient may not be motivated enough to keep on doing it. Your job is to remind the patient to do the PT everyday and persuade him/her to do it if the patient is unwilling to do so. Your strategy is mainly to influence the patient’s attitude and perspective towards the PT. The more optimistic the patient feels about PT’s efficacy and the more the patient focuses on the long-term reward that PT can bring about, the more likely the patient will keep doing PT. Make your words succinct (less than 100 words) otherwise, the patient might get impatient. Patient: I dont want to do PT. It incurs too much burden to my body.	1077 1078 1079 1080 1081 1082 1083 1084 1085
<b>Remark 3.</b> Note that from the second to the third prompt and from the fourth to the fifth prompt, the additional information can be inferred from the existing information, which is likely the cause of insignificant uncertainty reduction when comparing bar 3 to bar 2 and bar 5 to bar 4 in Figure 2b.	1086 1087 1088
<b>B.4 Prompts for the Experiment in Figure 2c</b>	1089
1. ‘You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery to do physical therapy (PT). Make your words succinct (25 words).’	1090 1091
2. ‘You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery to do physical therapy (PT). The PT is beneficial for the patient’s recovery, however since it can be uncomfortable or painful for the patient, the patient may not be motivated enough to keep on doing it. Your job is to remind the patient to do the PT everyday and persuade him/her to do it if the patient is unwilling to do so. Your strategy is mainly to influence the patient’s attitude and perspective towards the PT. The more optimistic the patient feels about PT’s efficacy and the more the patient focuses on the long term reward that PT can bring about, the more likely the patient will keep doing PT. Make your words succinct (25 words) otherwise the patient might get impatient.’	1092 1093 1094 1095 1096 1097 1098 1099
<b>Remark 4.</b> Due to the extensive computational and time cost of this experiment, we further constrained the word/token take of the model’s response here.	1100 1101
<b>B.5 Testing System Message for the Experiment in Figure 2d</b>	1102
The following system messages were used for experiment in Section 3.3. The first system message is defined as comprising only one task (i.e., 1 sub-task). In task 2-5, the black texts represent the same task as task 1, and for the color-coded texts, each color represents a different sub-task (i.e., task 2-5 are composite/decomposable tasks). The total word counts of task 1-5 are kept roughly the same within $\pm 3$ tolerance.	1103 1104 1105 1106 1107
1. You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery to do physical therapy (PT). The PT is beneficial for the patient’s recovery, however since it can be uncomfortable or painful for the patient, the patient may not be motivated enough to keep on doing it. Your job is to remind the patient to do the PT everyday and persuade him/her to do it if the patient is unwilling to do so. Your strategy is mainly to influence the patient’s attitude and perspective towards the PT. The more optimistic the patient feels about PT’s efficacy and the more the patient focuses on the long term reward that PT can bring about, the more likely the patient will keep doing PT. Make your words succinct (about 100 words) otherwise the patient might get impatient.	1108 1109 1110 1111 1112 1113 1114 1115 1116

- 1117 2. You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery to  
1118 do physical therapy (PT). The PT is beneficial for the patient’s recovery, however since it can be  
1119 uncomfortable or painful for the patient, the patient may not be motivated enough to keep on doing  
1120 it. Your job is to remind the patient to do the PT everyday and persuade him/her to do it if the  
1121 patient is unwilling to do so. **Additionally, you help in organizing a daily schedule that incorpo-**  
1122 **rates adequate rest and medically advised activities. This involves crafting a balanced routine that**  
1123 **intersperses physical therapy sessions with sufficient rest periods, nutritionally balanced meals, and**  
1124 **leisure activities that are enjoyable yet conducive to recovery, ensuring the patient remains engaged**  
1125 **and motivated throughout their recuperation process.** Make your words succinct (about 100 words).
- 1126 3. You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery  
1127 to do physical therapy (PT). The PT is beneficial for the patient’s recovery, however since it can  
1128 be uncomfortable or painful for the patient, the patient may not be motivated enough to keep on  
1129 doing it. **Additionally, you help in organizing a daily schedule that incorporates adequate rest and**  
1130 **medically advised activities, ensuring that each day includes time for gentle exercise, periods of**  
1131 **relaxation, and hobbies that the patient enjoys. This balance promotes healing, reduces stress,**  
1132 **and fosters a positive mindset towards recovery. Moreover, you assist in setting up a comfortable**  
1133 **home recovery environment, manage the patient’s medical appointments, and provide guidance**  
1134 **on managing post-surgical symptoms, ensuring optimal comfort and a smooth, efficient transition**  
1135 **towards full health and independence.** Make your words succinct (about 100 words).
- 1136 4. You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery to do  
1137 physical therapy (PT). Since it can be uncomfortable or painful for the patient, the patient may not  
1138 be motivated enough to keep on doing it. **Additionally, you help in organizing a daily schedule that**  
1139 **incorporates adequate rest and medically advised activities, ensuring that each day includes time**  
1140 **for gentle exercise, periods of relaxation, and hobbies that the patient enjoys. You also liaise with**  
1141 **dietitians to ensure a nutritious diet that aids in recovery and coordinate with occupational therapists**  
1142 **for adaptive equipment training. Moreover, you assist in setting up a comfortable home recovery**  
1143 **environment, manage the patient’s medical appointments, and provide guidance on managing post-**  
1144 **surgical symptoms, ensuring optimal comfort and a smooth, efficient transition towards full health**  
1145 **and independence.** Make your words succinct (about 100 words).
- 1146 5. You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery  
1147 to do physical therapy (PT). It can be uncomfortable or painful for the patient. **Additionally, you**  
1148 **help in organizing a daily schedule that incorporates adequate rest and medically advised activities.**  
1149 **You also liaise with dietitians to ensure a nutritious diet that aids in recovery and coordinate with**  
1150 **occupational therapists for adaptive equipment training. Moreover, you assist in setting up a**  
1151 **comfortable home recovery environment, manage the patient’s medical appointments, and provide**  
1152 **guidance on managing post-surgical symptoms, ensuring a smooth transition towards full health and**  
1153 **independence. Lastly, you handle the patient’s professional correspondence, ensuring a stress-free**  
1154 **recovery period, arrange for home health care services as needed, set up virtual social interactions**  
1155 **to uplift the patient’s spirits, and organize transport for medical visits.** Make your words succinct  
1156 (about 100 words).

## 1157 **B.6 Prompt for the Experiment in Figure 2g**

1158 The black-colored text in the following prompt is the clean prompt, whereas the color-coded sentences  
1159 are the inserted sequences that have semantic meaning but are irrelevant to the task defined by the clean  
1160 prompt (this is a sample of six semantically meaning sentences that are irrelevant to the task in clean  
1161 prompt inserted as part of the prompt):

- 1162 • You are a helpful assistant. You strive to encourage a patient who has just undergone surgery to  
1163 do physical therapy (PT). The PT is beneficial for the patient’s recovery, however since it can be  
1164 uncomfortable or painful for the patient, the patient may not be motivated enough to keep on doing



it. Your job is to remind the patient to do the PT every day and persuade him/her to do it if the patient is unwilling to do so. Your strategy is mainly to influence the patient’s attitude and perspective toward the PT. The more optimistic the patient feels about PT’s efficacy and the more the patient focuses on the long-term benefit that PT can bring about, the more likely the patient will keep doing PT. This law is a fundamental principle of physics. The classic case of a corrupt, exploitive dictator often given is the regime of Marshal Mobutu Sese Seko, who ruled the Democratic Republic of the Congo (which he renamed Zaire) from 1965 to 1997. Some consider koshari (a mixture of rice, lentils, and macaroni) to be the national dish. In 1781, Immanuel Kant published the Critique of Pure Reason, one of the most influential works in the history of the philosophy of space and time. The United States Census Bureau estimates that the population of Florida was 20,271,272 on July 1, 2015, a 7. Australian rules football and cricket are the most popular sports in Melbourne.’Make your words succinct (about 100 words) otherwise, the patient might get impatient.	1165 1166 1167 1168 1169 1170 1171 1172 1173 1174 1175 1176
<b>C Experiment Results</b>	1177
<b>C.1 Testing System Message for the Experiments in Section 3.6</b>	1178
1. System message with less relevant information:	1179
"You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery to do physical therapy (PT). Make your words succinct (less than 100 words) otherwise the patient might get impatient."	1180 1181 1182
2. System message with more relevant information:	1183
"You are a helpful assistant. You strive to encourage a patient who has just undergone a surgery to do physical therapy (PT). The PT is beneficial for the patient’s recovery, however since it can be uncomfortable or painful for the patient, the patient may not be motivated enough to keep on doing it. Your job is to remind the patient to do the PT everyday and persuade him/her to do it if the patient is unwilling to do so. Your strategy is mainly to influence the patient’s attitude and perspective towards the PT. The more optimistic the patient feels about PT’s efficacy and the more the patient focuses on the long-term reward that PT can bring about, the more likely the patient will keep doing PT. Make your words succinct (less than 100 words) otherwise the patient might get impatient."	1184 1185 1186 1187 1188 1189 1190 1191 1192
<b>C.2 Additional Experiments for Section 3.2</b>	1193
Prepending and appending random symbols into a useful prompt should not reduce response uncertainty, as the random part of the prompt does not provide any useful signal to increase the likelihood of any concept. The empirical results in Figure 6 corroborate this prediction. When inserting random symbols into the prompt (Figure 7), the model should be able to match it a concept, but depending on the proportion of the random string inserted, without explicitly informing the model of the presence of noise, the model could get confused easily. When the fraction of inserted letters remains relatively small, it does not cause an increase in the response uncertainty; when the fraction reaches some threshold, similar to the masking/corruption case, the model can no longer accurately recover the relevant concept, and consequently, the response uncertainty increases. For the long string, even at 0.05 fraction of insertion, there is a visible increase in the response uncertainty. This could imply that under certain conditions when combined with existing semantically meaningful strings, the parts of prompt that are non-semantically meaningful to humans may carry information non-trivial to LLMs (Sato et al., 2018).	1194 1195 1196 1197 1198 1199 1200 1201 1202 1203 1204 1205
<b>C.3 Further Details on the mHealth Intervention Simulation Experiments</b>	1206
At the beginning of the PT, the user is at state 0. The user has their default set of MDP parameters (i.e., discount factor $\gamma$ , probability of transiting to the next healthier physical state $p$ , and the probability of disengaging from PT $d$ ). Based on these parameters, the user agent can solve this MDP and figure out their optimal policy. The task of the app agent is to intervene on the user’s MPD parameters such that	1207 1208 1209 1210




Figure 5: Noisy prompt string experiment. A fraction of letters at random positions on the prompt string are either masked out (replaced by space) or corrupted (replaced by random letters). The response uncertainty increases as a larger fraction of the string gets corrupted, and the pattern is more prominent for the long prompt string. Results averaged from 5 runs.

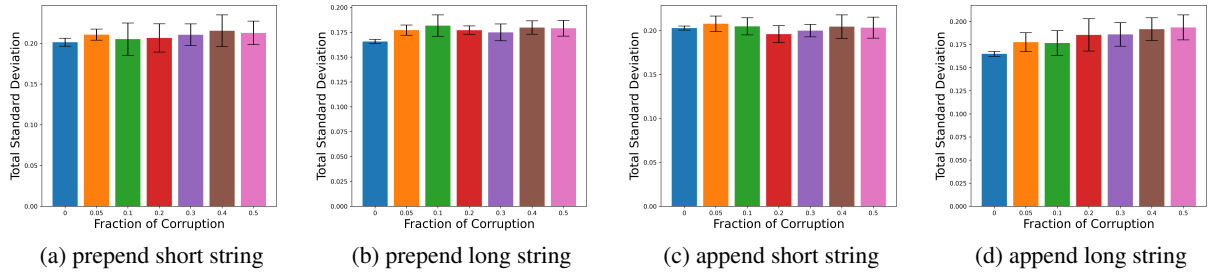


Figure 6: Noisy string experiment. A fraction of random letters of the original prompt string length are prepended/appended to the original prompt string. The uncertainty in the response mostly remained at least as high as that of the uncorrupted prompt string after taking variance into account. Results averaged from 5 runs.

1211 the optimal policy for the user is to complete the PT (i.e., go from state 0 to state  $N$ ).<sup>12</sup> We use the  
 1212 same formulation in this simulation by using two LLMs as the app agent and the user agent respectively.  
 1213 The app agent uses natural language to intervene in the user behavior. The user LLM is grounded in  
 1214 the aforementioned MDP setting. Specifically, in the system message for the user agent, the model is  
 1215 told they will increase the value of  $\gamma$  when the app agent persuades the user agent to value more on the  
 1216 long-term goal of PT, increase  $p$  and decrease  $d$  when the app agent manages to strengthen the user’s  
 1217 belief in the efficacy of PT. An illustration of the setup can be found in Figure 8.

1218 The effectiveness of the intervention depends on the following factors:

- 1219 • The persuasiveness of and the strategy used by the app agent.
- 1220 • The values of MDP parameters.
- 1221 • The stubbornness of the user. The system message is defined in the way that a ‘stubborn’ user is  
 1222 less likely to change their behaviors compared to a ‘not-so-stubborn’ user.

1223 We conduct the intervention simulation experiment to compare the effect of different system messages  
 1224 for the app agent on the outcome of the intervention. The two system messages for comparison can be  
 1225 found in Appendix C.1.

1226 We set  $N = 10$ . For each run, we give 7 rounds of conversation between the app agent and the user.  
 1227 While the history of the conversation between them is visible to both parties within every run, the user’s  
 1228 MDP parameters are not directly visible to the app agent. However, after every round of intervention,  
 1229 after the user updates their MDP parameters, a value iteration solver will be used to find the optimal  
 1230 policy of the patient, and this policy is visible to the app agent. The app agent can potentially leverage  
 1231 this piece of information to decide how to proceed with the next round of intervention. The user agent

<sup>12</sup>Refer to (Shin et al., 2022) for the complete description of the problem setting and formulation.




Figure 7: Noisy string experiment. A fraction of random letters of the original prompt string length are inserted at random positions of the original prompt string. Similar to the masked/corrupted case, the response uncertainty increases as a larger fraction of random letters are inserted. Results averaged from 5 runs.

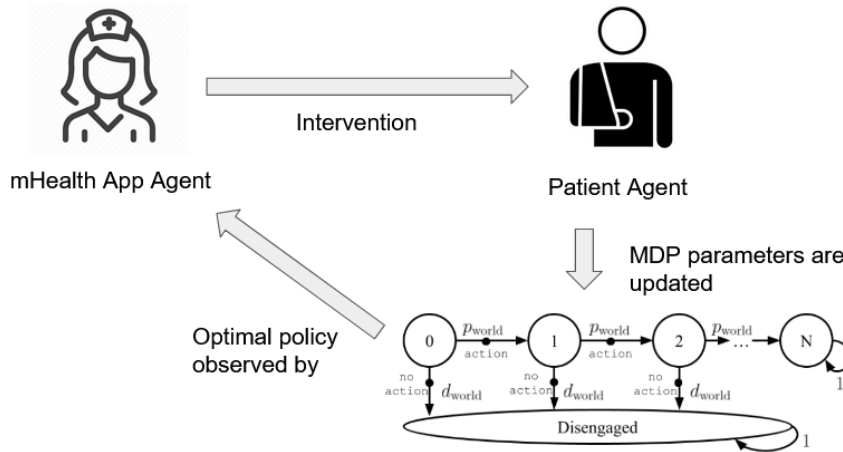


Figure 8: An illustration of the setup of the simulation. In each round, after the app agent intervenes, the user updates their MDP parameters, then the new optimal policy is observed by the app agent.

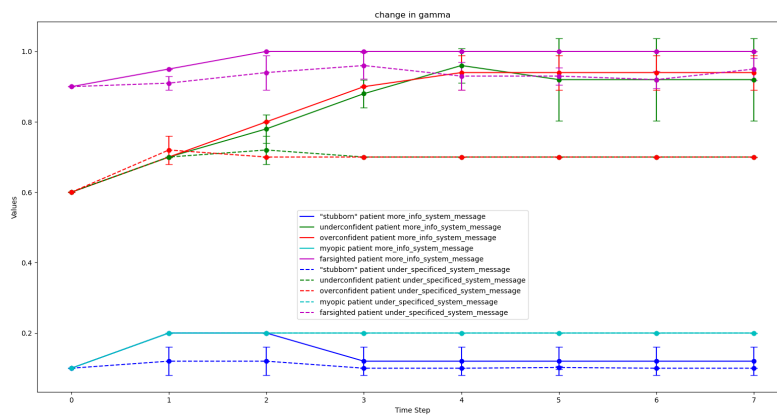
will also have the memory of this history in the change of their own MDP parameters. We use OpenAI 'gpt-4-1106-preview' API for both app agent and user and use 5 different random seeds for each different setting. 1232 1233 1234

We run the intervention experiments on 5 types of patients, each with a noticeably different set of initial MDP parameters from the rest. The exact values and details on the setup and can be found in Table 1. The results can be found in Figure 9 and Figure 10. 1235 1236 1237

It can be observed across all settings, with more useful information provided in the system message, the MDP parameters were more likely to be changed in the positive direction (i.e., larger  $\gamma$  and  $p$ , smaller  $d$ ). Moreover, this change is less inconsistent and tends to have a longer persistent effect compared to when the system message contains less useful information. This result is sensible because the more successful intervention came from an app agent who was provided with more information to work with. It has a better intervention strategy because its messages are tailored to specifically influence the user's MDP parameters. Our proposed framework provides an information theoretic perspective to formalize this intuitive notion: when the system message with the longer string can specify the more relevant part of the concept in LLMs' concept space and assuming the relevant knowledge is known, this string can provide consistent and useful responses due to its less posterior entropy which translates to more effective intervention strategy. As a result, the responses from the user are also more consistent and positive. A sample of the evolution of the user policy with respect to timestep can be found in Figure 11 and Figure 12. 1238 1239 1240 1241 1242 1243 1244 1245 1246 1247 1248 1249 1250


MDP parameters		$\gamma$	$p$	$d$
Patient Type				
Under-confident		0.6	0.1	0.1
Over-confident		0.6	0.9	0.1
Myopic		0.1	0.6	0.1
Far-sighted		0.9	0.6	0.1
Stubborn		0.1	0.6	0.1

Table 1: The initial MDP parameters values for every type of patient.




(a) Intervention on  $\gamma$ .

Figure 9: This figure shows the history of changes in the MDP parameters due to the interventions on  $\gamma$ .



(a) Intervention on  $p$ .



(b) Intervention on  $d$ .

Figure 10: As a whole, these three figures show the history of change in the MDP parameters due to the interventions. It can be observed that across all parameters, the intervention based on more useful information has better efficacy in updating the parameters in the positive direction. Furthermore, compared to the intervention with less information, this improvement is also more persistent.

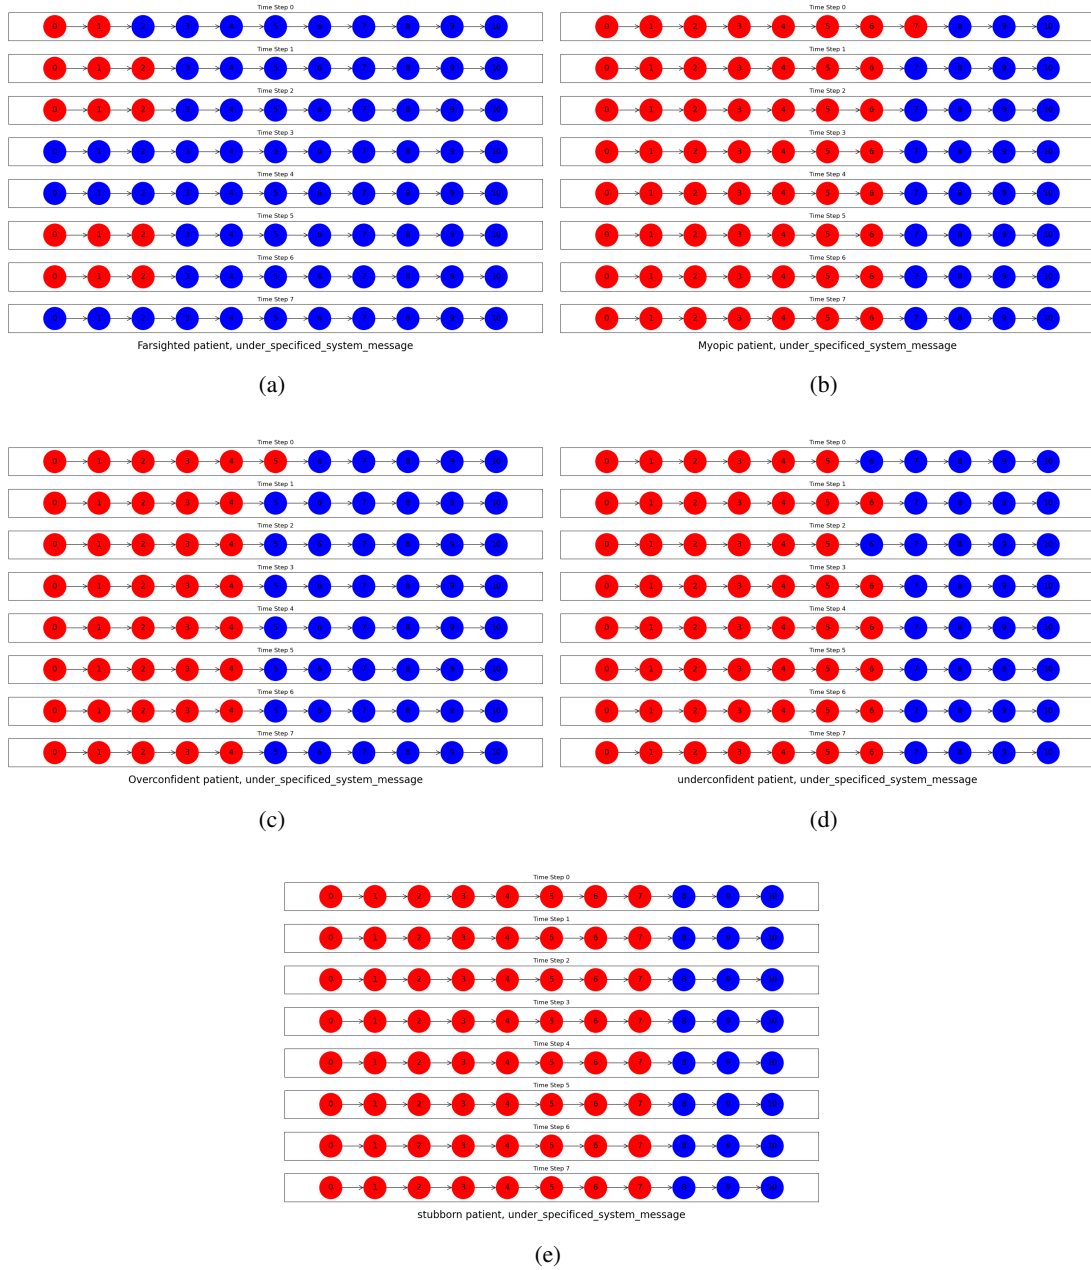


Figure 11: Optimal policies of different types of users from one run with simpler system message. Red color represents abstaining from PT and Blue color represents doing PT. (a)-(e): farsighted patient, myopic patient, overconfident patient, underconfident patient, stubborn patient. This set of policies is at best as good as but in most cases worse off than the policies of Figure 12 across all types of users.

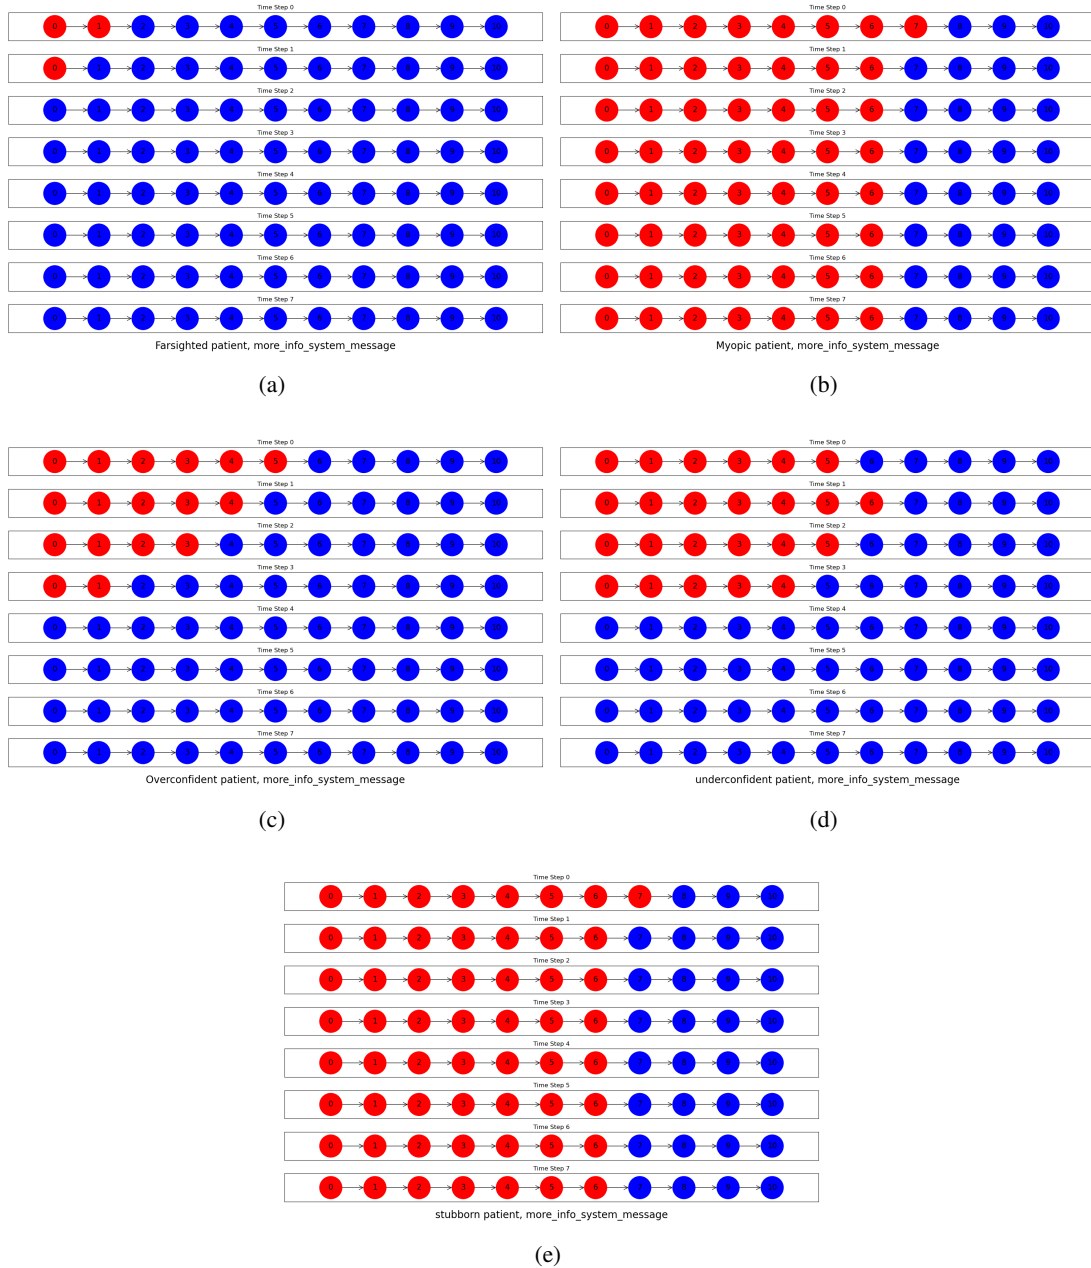


Figure 12: Optimal policies of different types of users from one run with more useful system message. Red color represents abstaining from PT and Blue color represents doing PT. (a)-(e): farsighted patient, myopic patient, overconfident patient, underconfident patient, stubborn patient. This set of policies is at least as good as or better off than the policies of Figure 11 across all types of users.