

Promise: Model-driven Stateful Prompt Orchestration for Persuasive Conversational Interactions

Wenyuan Wu and **Jasmin Heierli** and **Max Meisterhans**
Adrian Moser and **Andri Färber** and **Mateusz Dolata**
Elena Gavagnin and **Alexandre de Spindler** and **Gerhard Schwabe**
wenyuan@ifi.uzh.ch

Abstract

Natural language-based interactions are gaining importance within health information systems, showing promise in achieving key medical objectives. For example, they support a more patient-centered approach in healthcare by adapting to the unique needs of each patient. When properly integrated, these interactions can streamline patient-physician collaboration, leading to more targeted and effective consultations and treatments. Such consultations involve conversational interactions that usually comprise several consecutive, alternating, possibly nested strands of conversation, and in which persuasive strategies are used in a targeted manner. Consequently, if a system is to have a similar treatment benefit between consultations, it must be able to handle such complex and persuasive interactions.

While the development of increasingly powerful language models (LMs) raises expectations for more sophisticated conversational behaviors, directing LM behavior remains challenging, making fast, iterative experimentation difficult. Despite recent advances in prompting techniques that provide a more efficient method of controlling behavior, prompts alone cannot ensure consistent LM behavior in complex interactions.

To facilitate the investigation of the feasibility and utility of complex and persuasive conversational interactions, we developed PROMISE (PRompt-Orchestrating Model-driven Interaction State Engineering). This framework provides support for the rapid design, implementation, and experimental validation of such interactions by breaking down complex interactions into more specific prompts that increase the predictability of LMs and leverage their conversational capabilities.

PROMISE has been effectively applied in projects such as to improve health literacy or therapy adherence. It enables personalized conversational interactions by integrating patient data from electronic health records, facilitating empathetic and persuasive conversations, and dynamically adapting conversational behaviours to achieve intended interaction outcomes. It is available at GitHub (<https://github.com/zhaw-iwi/promise>) and includes an API supporting the design and implementation of conversational interactions. It also provides support for web-based interactions, including transparent persistence, a REST API, and a ready-to-use default chat interface. More recently, we added the means to dynamically adapt persuasion strategies throughout complex interactions.