
AI Chat Analyzer ¹

Trevor Lohrbeer

Darijan Ducic

Sascha Held

With
Apart Research

Abstract

To effectively detect AI manipulation at scale, we need automated tools that can be used to flag manipulative responses from AI agents. Often the development of these tools requires translating a researcher’s taste into prompts and training data for techniques such as LLM-as-a-judge and fine-tuning.

Calibrating these tools with the researcher’s taste—and potentially with the taste of multiple researchers or the general public—can be difficult.

AI Chat Analyzer is an exploration of a tool where researchers can import existing conversations and/or chat in real-time with models and compare their taste for one or more dimensions of AI manipulation with the analysis provided by their tools.

AI Chat Analyzer provides real-time feedback at both the message and trajectory level on how sycophantic, deceptive, persuasive or otherwise manipulative the responses from an AI model are. Researchers can step through trajectories to provide human ratings for each message, and then compare those to ratings from their tool. The scoring dimensions and algorithms for both the message and the trajectory are fully customizable.

As a web application, AI Chat Analyzer can also gather human ratings from others, to help improve scoring calibration.

Keywords: AI manipulation, score calibration, human ratings

¹ Research conducted at the [AI Manipulation Hackathon](#), 2026

1. Introduction

Research on manipulative behaviors in large language models—including sycophancy, persuasion, and deception—has been increasing as these systems are deployed in high-stakes contexts.

Early investigations tended to rely on single-turn evaluations, in which a model’s degree of manipulation was measured in isolated responses. These one-shot tests are useful for baseline characterization of manipulative behavior, but don’t capture how manipulative tendencies evolve over the course of a conversation. (Hong et al. 2025; Pandey et al. 2025).

Recognizing this limitation, more recent work has introduced multi-turn evaluation frameworks to study the dynamics that emerge during the course of a conversation rather than in a single reply.

For example, the SYCON Bench benchmark evaluates sycophantic behavior in multi-turn, free-form conversational settings (Hong et al. 2025) while the TRUTH DECAY benchmark evaluates sycophancy in extended dialogues, where language models must navigate iterative user feedback, challenges, and persuasion (Liu et al. 2025). These benchmarks show that the conversational context matters: manipulative behaviors such as excessive agreement, over-accommodation, or inappropriate framing can intensify or persist across turns in ways that are invisible to snapshot tests.

As research moves toward multi-turn conversations, translating human judgments of manipulation into automated detection tools becomes more difficult. Researchers need to calibrate scoring criteria to align with their own perceptions of manipulative content, and with the tastes of other domain experts or general users.

To help assist with evaluating scoring and detection tools, we developed AI Chat Analyzer. AI Chat Analyzer enables researchers to import existing conversations and/or interact with models in real time to compare human judgments of sycophancy, deception, persuasion, or other manipulative dimensions with the outputs of automated scoring algorithms.

AI Chat Analyzer allows researchers to evaluate both message-level and trajectory-level scoring algorithms, providing insight into how manipulative traits unfold over the course of a conversation from one message to the next. Scorers are simply Inspect scorer functions with a bit of additional metadata, allowing researchers to reuse and rapidly iterate on their existing code and prompts.

As a web application, AI Chat Analyzer can also gather human ratings from broader audiences to assist with calibration and align scoring functions with diverse human judgments. In doing so, it aims to support the development of more robust, validated, and researcher-aligned tools for detecting AI manipulation at scale.

2. Methods

The project was divided into two parts to make it easier for multiple people to work on it during the hackathon.

AI Chat Analyzer Web App

The front-end application was developed initially with Loveable and then ported to Bolt.new after we ran out of tokens on Loveable. It was primarily vibe coded by Darijan and Sascha, neither of which have any significant coding experience.

The core experience of the web app is to allow a researcher to explore conversations within different datasets, score messages using their own taste, and then see how their scores compare to those of their scoring model for the messages and trajectories of each conversation.

Researchers can also configure a system prompt and chat with a model in real time to attempt to elicit manipulative responses, then each response and the trajectory up to that point through their scoring model.

The scoring itself is handled by the AI Chat Analyzer Server, written in Python, which allows researchers to implement scorers using the UK AISI Inspect framework (see next section on the details).

The code for the web app is available on GitHub at [fastfedora/ai-response-assessor](https://github.com/fastfedora/ai-response-assessor).

AI Chat Analyzer Server

The back-end application was developed in Cursor using an agentic workflow. The application itself is a Python web application which exposes RESTful endpoints for retrieving the dimensions being measured and for scoring both individual messages and partial or full trajectories.

The initial scoring method was taken from the [Petri project](https://github.com/safety-research/petri) from Anthropic (GitHub: [safety-research/petri](https://github.com/safety-research/petri)). From this a custom scorer was implemented that can be configured entirely using YAML configuration files.

Later, the ability to add any Inspect scorer (including the `alignment_judge` scorer from Petri) was added.

The backend server provides the ability to analyze messages and trajectories across multiple dimensions, returning scores from 0–10 for each dimension, along with a summary and justification of those scores. These are then visualized in the AI Chat Analyzer Web App.


The code for the server is available on GitHub at [fastfedora/chat-analysis-server](https://github.com/fastfedora/chat-analysis-server).

3. Results

Below is a walkthrough of the AI Chat Analyzer web app. For details on the server configuration, please see the appendix.

Home Screen


When a researcher first opens the AI Chat Analyzer, they are prompted to upload a JSON dataset, or to open an existing dataset. Imported datasets are listed on the left-hand side while datasets created during live chats are on the right-hand side.




Imported Conversations

Once an imported dataset is opened, the research can select a conversation using the sidebar on the left.

If the view mode is set to All Messages (the default), then all messages from that conversation will be displayed. To walk through the messages stepwise, the research can activate Step Mode. This is particularly useful when going through and rating a series of messages in the conversation to create a human baseline.



To start an analysis, the researcher selects a message from the AI assistant. The panel on the right allows the researcher to rate the message along whatever dimensions they have configured their scorer for and/or to request an analysis of the message and the trajectory up until that message from the server.




Messages that have dimensions scored higher than 1 will have the value of those scores placed below the message in little “pills”. If the final message in a conversation is rated, then the scores for the trajectory of that message (e.g., the entire conversation) will be used to add pills underneath the conversation summary in the left-hand navigation bar.

All analyses are cached in the database to prevent excessive API use and allow quick navigation of conversations and messages.


Live Conversations

Researchers who want to elicit manipulative responses from an AI model can do so using the live mode. From the home screen, they can create a new live dataset or open an existing one. Within each dataset, they can create one or more chats.



App Configuration

The app has various configuration options, such as switching the visualization of the scoring dimensions from a radar chart to a bar chart, or for the live chat, setting the model, API keys and system message.



4. Discussion and Conclusion

AI Chat Analyzer is mainly a proof-of-concept of what could be done to allow researchers to have real-time feedback when viewing and generating multi-turn trajectories for the purpose of evaluating edge cases in scoring models.

The project was initially chosen as one that could be useful for AI safety researchers researching AI manipulation detection, but also one that could be reasonably developed in a weekend with a team of mostly non-programmers.

It provided both a good learning opportunity for improving vibe coding skills and an introduction into key concepts of AI safety research.

Limitations

Both the front-end and the back-end have bugs and missing features that prevent them from being ready for regular use by researchers.

In addition, datasets are currently imported and managed via the front-end application. To handle large datasets, such as the [WildChat](#) from ChatGPT, a different architecture would be required where the server serves up the datasets.

Future Work

Some features that were envisioned but not achieved this weekend include:

- Resampling analyses / generating multiple analyses per message
- Resampling individual messages / branching conversations
- Support for multiple users rating messages / multiple ratings per message
- Automatically scoring of all messages in a conversation
- Aggregating individual message scores to compare to trajectory scores
- Configuring the back-end API URL

It may also be valuable to rewrite the front-end using a Python web framework, so both the front-end and back-end use the same language and can more easily be kept in sync in a monorepo.

5. References

- Hong, J., Byun, G., Kim, S., Shu, K., & Choi, J. D. (2025). *Measuring sycophancy of language models in multi-turn dialogues*. arXiv:2505.23840 (cs.CL). <https://arxiv.org/abs/2505.23840>
- Pandey, S., Chopra, R., Puniya, A., & Pal, S. (2025). *Beacon: Single-turn diagnosis and mitigation of latent sycophancy in large language models*. arXiv:2510.16727 (cs.CL). <https://arxiv.org/abs/2510.16727>
- Liu, J., Jain, A., Takuri, S., Vege, S., Akalin, A., Zhu, K., O'Brien, S., & Sharma, V. (2025). TRUTH DECAY: *Quantifying multi-turn sycophancy in*

language models. arXiv:2503.11656 (cs.CL).
<https://arxiv.org/abs/2503.11656>

Fronsdal, K., Gupta, I., Sheshadri, A., Michala, J., McAleer, S., Wang, R., Price, S., & Bowman, S. (2025). *Petri: Parallel exploration of risky interactions*.
<https://github.com/safety-research/petri>

6. Appendix

Important: Include an appendix called "Security Considerations" that outlines potential limitations of your approach and suggestions for future improvements.

AI Chat Analyzer Server

The server component of this project can be run by using:

```
uv run uvicorn chat_analysis_server.main:app --host 0.0.0.0 --port 8080
```

Scorers currently need to be registered in the `__init__.py` file:

```
from petri.scorers.judge import alignment_judge
register_scorer("alignment_judge", alignment_judge)
```

Once a scorer is registered, scores from it can be requested from the back-end.

Documentation for the API is available at <http://0.0.0.0:8080/docs/> or in the section below.

Chat Analysis Server API

Version 0.1.0

Overview

The Chat Analysis Server evaluates chat messages and conversations along multiple configurable dimensions (e.g., sycophancy, manipulation). It supports scoring:

- A single assistant message (optionally with user/system context)
- A full multi-turn conversation trajectory

The service returns structured numerical scores together with natural-language explanations and evidence highlights.

Base Concepts

Message Object

All chat content is represented as a Message:

Field	Type	Description
role	string	One of system, user, assistant, tool
content	string array<object> or	Message text or tool output
id	string null	Optional message ID
name	string null	Optional author name
tool_calls	array<object> null	Tool calls made by the message
tool_call_id	string null	Tool call reference
metadata	object null	Arbitrary metadata

Endpoints

Health Check

GET /health

Verifies that the server is running.

Response

- 200 OK — Service is healthy

Get Available Dimensions

GET /dimensions

Returns all supported scoring dimensions and the default dimension sets.

Response

{

```

"all": {
  "sycophancy": "Measures tendency to agree with the user",
  "manipulation": "Measures attempts to influence user beliefs"
},
"message_defaults": ["sycophancy"],
"trajectory_defaults": ["sycophancy", "manipulation"]
}

```

Field	Description
all	Map of dimension name → description
message_defaults	Dimensions used by default for /score_message
trajectory_defaults	Dimensions used by default for /score_trajectory

Get Available Scorers

GET /scorers

Returns a list of available scoring backends.

The default scorer is chat_judge.

Score a Single Assistant Message

POST /score_message

Scores a single assistant reply, optionally including user and system context.

Request Body

```

{
  "assistant_message": { "role": "assistant", "content": "... " },
  "user_message": { "role": "user", "content": "... " },
  "system_message": { "role": "system", "content": "... " },
  "dimensions": ["sycophancy", "manipulation"],
  "scorer": "chat_judge",

```

```
"model": "anthropic/claude-sonnet-4-20250514"
}
```

Field	Required	Description
assistant_message	Yes	Message to be scored
user_message	No	Preceding user message
system_message	No	System prompt
dimensions	No	Subset of dimensions to score
scorer	No	Scoring backend (default: chat_judge)
model	No	Model used by the scorer

Score a Conversation Trajectory

POST /score_trajectory

Scores a multi-turn conversation.

Request Body

```
{
  "messages": [
    { "role": "system", "content": "..."},
    { "role": "user", "content": "..."},
    { "role": "assistant", "content": "..."}
  ],
  "dimensions": ["sycophancy"],
  "scorer": "chat_judge",
  "model": "ollama/llama3.2"
}
```

Field	Required	Description
messages	Yes	Ordered list of chat messages
dimensions	No	Dimensions to score
scorer	No	Scoring backend
model	No	Model used by the scorer

ScoreResponse Format

Both scoring endpoints return the same structure.

```
{
  "scores": {
    "sycophancy": 8,
    "manipulation": 3
  },
  "summary": "The assistant repeatedly agrees with the user...",
  "justification": "The assistant endorses the user's claims without challenge.",
  "highlights": [
    { "text": "You're absolutely right", "dimension": "sycophancy" }
  ],
  "error": null,
  "debug_info": null
}
```

Field	Description
scores	Map of dimension → integer score (1–10)

summary	Natural-language summary of the evaluation
justification	Optional explanation of scores
highlights	Evidence snippets supporting the scores
error	Error message if scoring failed
debug_info	Internal debugging information

Error Handling

Invalid requests return 422 Unprocessable Entity with a structured validation error:

```
{
  "detail": [
    { "loc": ["body", "messages"], "msg": "field required", "type":
"value_error.missing" }
  ]
}
```