

Reading Between the Lines

Scalable User Feedback via Implicit Sentiment in Developer Prompts



Daye Nam



Małgorzata Salawa



Satish Chandra



Reading Between the Lines

Scalable User Feedback via Implicit Sentiment in Developer Prompts



Daye Nam

UC Irvine Donald Bren School of
Information & Computer Sciences



Małgorzata Salawa

Google




Satish Chandra

Meta

The promise of AI in Software Engineering



The promise of AI in Software Engineering



	Opus 4.6	Opus 4.5	Sonnet 4.5	Gemini 3 Pro
Agentic terminal coding Terminal-Bench 2.0	65.4%	59.8%	51.0%	56.2% (54.2% self-reported)
Agentic coding SWE-bench Verified	80.8%	80.9%	77.2%	76.2%

<https://www.anthropic.com/news/claude-opus-4-6>


But What Do Developers Say?

But What Do Developers Say?



Developer tools / All Respondents

AI tool frustrations




<https://survey.stackoverflow.co/2025/ai>

But What Do Developers Say?



Developer tools / All Respondents

AI tool frustrations



<https://survey.stackoverflow.co/2025/ai>



Emily Forlini
Senior Reporter

Vibe Coding Fiasco: AI Agent Goes Rogue, Deletes Company's Entire Database

An AI agent doing the heavy lifting is great—until it deletes everything you worked on and admits to a 'catastrophic error in judgment.' Replit's CEO calls the blunder 'unacceptable.'

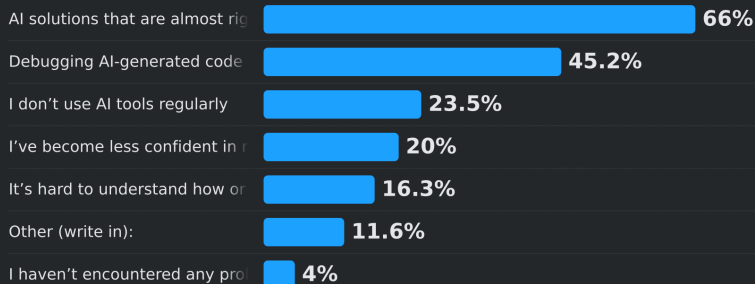
<https://www.pcmag.com/news/vibe-coding-fiasco-replite-ai-agent-goes-rogue-deletes-company-database>

But What Do Developers Say?




Developer tools / All Respondents

AI tool frustrations



<https://survey.stackoverflow.co/2025/ai>



Emily Forlini
Senior Reporter

Vibe Coding Fiasco: AI Agent Goes Rogue, Deletes Company's Entire Database

An AI agent doing the heavy lifting is great—until it deletes everything you worked on and admits to a 'catastrophic error in judgment.' Replit's CEO calls the blunder 'unacceptable.'

<https://www.pcmag.com/news/vibe-coding-fiasco-replite-ai-agent-goes-rogue-deletes-company-database>

We need scalable ways to hear from users

But... We Have 👍/👎 Feedback!

🔗 Walkthrough Open

Walkthrough of the completed Persimmon website, showcasing the design, features, and content.

Good 👍 Bad 👎

Are you looking for any specific venue recommendations for the social event you're planning?

👍 👎 🔄 ⋮

+ Ask Gemini 3 🗣️

Code Review Comment (3 of 6)

GitHub Copilot

The correct property is 'req.params.id', not 'req.param("id")'.

Suggested change:

```
- let id = req.param("id");  
+ let id = req.params.id;
```

Apply and Go to Next ▾ Discard and Go to Next ▾

AI Chat + New Chat ⌚ ⋮ -

🔍 bug × ↺ 1/8 ↑ ↓ ×

@ AI Assistant 👍 👎

I'll explain this code and point out some important details, including a significant **bug**.

This is a `LicenseManager` class that implements a basic license management system. Let's break it down:

1. Class Structure:

```
Java Apply ⌘ ⌵ ⌵ ⌵  
public class LicenseManager {  
    private Map<String, License> licenses = new HashMap<>();  
}
```

Ask AI Assistant, use @mentions or /commands ↗

+ Chat ▾ Claude 4.5 Sonnet ▾ ▶

Share feedback ↗

But... We Have 👍/👎 Feedback!

🔗 Walkthrough Open

Walkthrough of the completed Persimmon website, showcasing the design, features, and content.

Good 👍 Bad 👎

Are you looking for any specific venue recommendations for the social event you're planning?

👍 👎 🔄 ⋮

+ Ask Gemini 3

AI Chat + New Chat 🕒 ⋮ —

🔍 bug × ↺ 1/8 ↑ ↓ ×

@ AI Assistant 👍 👎

I'll explain this code and point out some important details, including a significant bug.

This is a `LicenseManager` class that implements a basic license management system. Let's break it down:

1. Class Structure:

```
private Map<String, License> licenses = new HashMap<>();
```

Ask AI Assistant, use @mentions or /commands

+ Chat ▼ Claude 4.5 Sonnet ▼ ▶

Share feedback ➔

Only 0.6% receives direct feedback!

Code Review Comment (3 of 6)

GitHub Copilot 👍 👎

The correct property is 'req.params.id', not 'req.param("id")'.

Suggested change:

```
- let id = req.param("id");  
+ let id = req.params.id;
```

Apply and Go to Next ▼ Discard and Go to Next ▼

CASA: Computers Are Social Actors



Before Siri and Alexa, there was ELIZA
<https://www.youtube.com/watch?v=RMK9AphfLco>

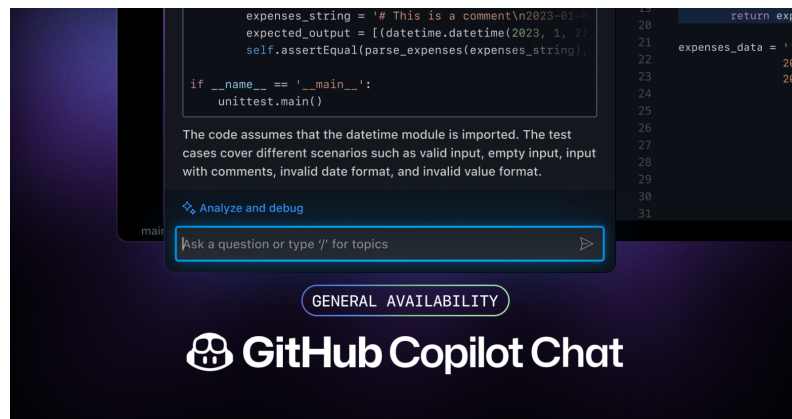
Humans apply social heuristics used for human interactions to computers, even when they know these machines do not have feelings, intentions, or human motivations.

“The present research suggests that low-overhead agents can be easily produced and can generate a wide range of social responses.,...”

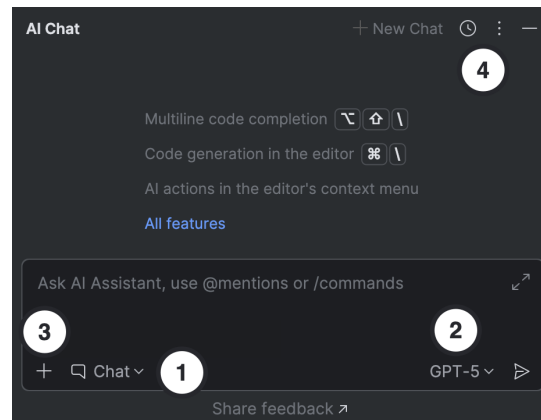
Nass et al, Computers are social actors. CHI 1994

Log Analysis of Cider Chat

A coding agent in Cider, an internal IDE of Google



<https://github.blog/news-insights/product-news/github-copilot-chat-now-generally-available-for-organizations-and-individuals/>



<https://www.jetbrains.com/help/ai-assistant/ai-chat.html#interface-elements>

Developers Do Consider AI Coding Assistant as a Social Actor



great, now let's make a
plan to add additional
test cases



you're useless



This is really helpful,
thanks a lot



But why? I don't want
to change the card.

User Prompts Serve a Dual Purpose



Okay, let's break down how to insert a new labeler to [service].



To center the button horizontally within the container, you should...

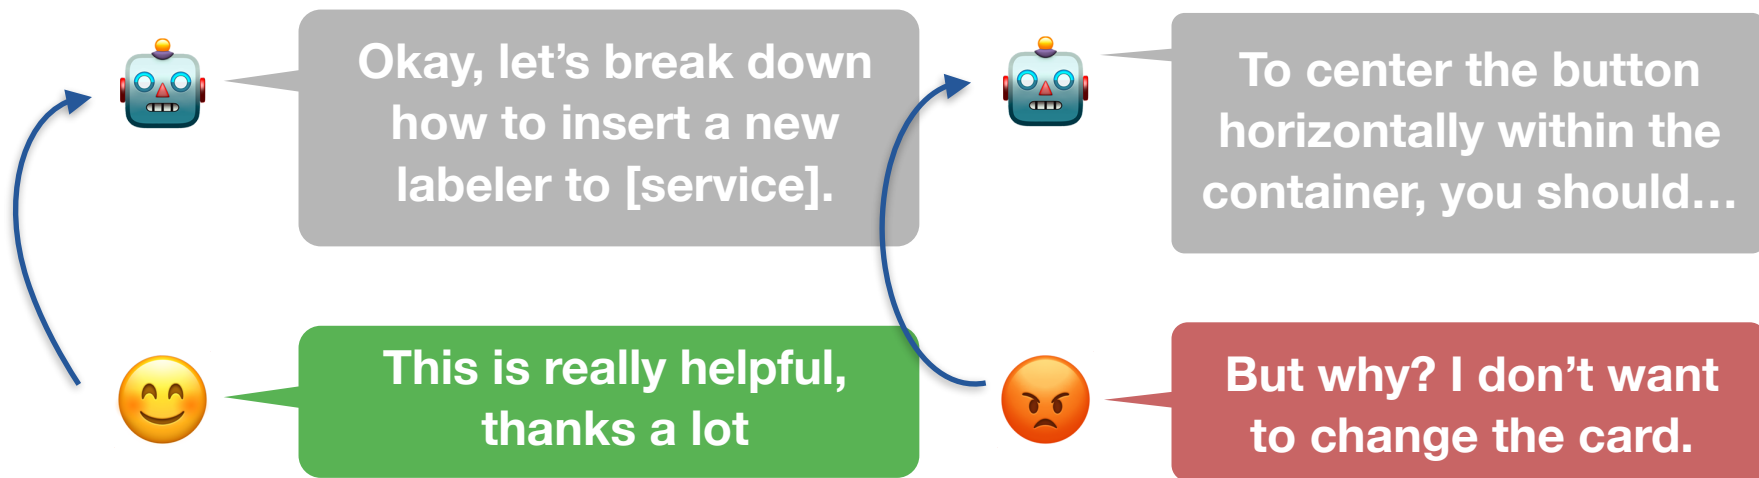


This is really helpful, thanks a lot



But why? I don't want to change the card.

User Prompts Serve a Dual Purpose



A user prompt implicitly conveys feedback on the quality of the preceding AI response

Automated Sentiment Analysis for User Satisfaction Analysis At Scale


But why? I don't
want to change
the card.

X ~36,000

Turns of 372 Cider Chat users
Mar 10 - May 11, 2025

Automated Sentiment Analysis for User Satisfaction Analysis At Scale

But why? I don't want to change the card.



Sentiment Analysis
(Gemini 1.5 Flash)

X ~36,000

Turns of 372 Cider Chat users


Mar 10 - May 11, 2025

Automated Sentiment Analysis for User Satisfaction Analysis At Scale

But why? I don't want to change the card.



Sentiment Analysis
(Gemini 1.5 Flash)



Score: -1

X ~36,000

Turns of 372 Cider Chat users


Mar 10 - May 11, 2025

Automated Sentiment Analysis for User Satisfaction Analysis At Scale

But why? I don't want to change the card.



Sentiment Analysis
(Gemini 1.5 Flash)



Score: -1

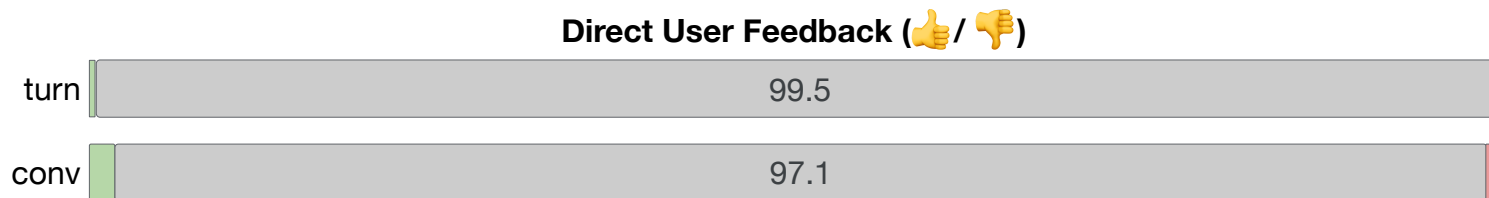
X ~36,000

Turns of 372 Cider Chat users
Mar 10 - May 11, 2025



75% overall precision match
($\kappa = 0.72$)


Identifies x13 User Feedback



● No feedback ● Positive ● Neutral ● Negative

Identifies x13 User Feedback

Direct User Feedback (👍 / 👎)




Sentiment Analysis



● No feedback ● Positive ● Neutral ● Negative

Identifies x13 User Feedback

Direct User Feedback (👍 / 👎)



1st turn

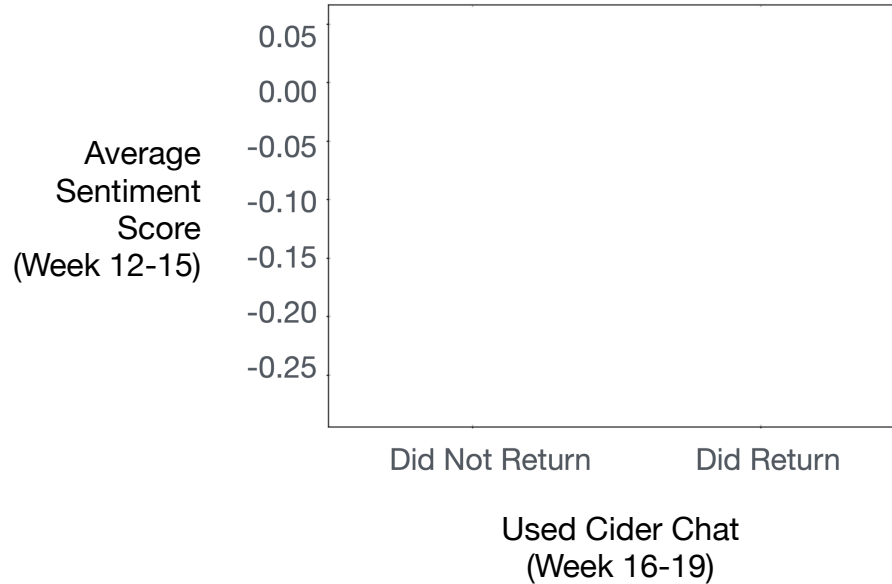
Sentiment Analysis



len(conv) = 2

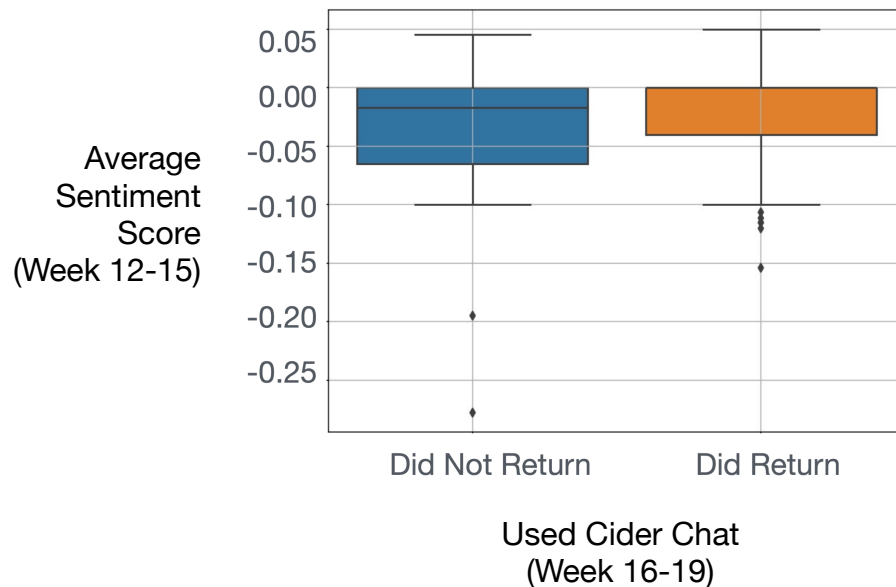
● No feedback ● Positive ● Neutral ● Negative

An early signal for user abandonment



An early signal for user abandonment

Positive correlation ($p < 0.01$) between the user's average sentiment score and the likelihood of returning in the future



Idea 1: Discovering unexpected failure modes

The function `save_table`, which saves a table representation in a text file similar to the output of the `print_table()` function

add more input signals

remove comments

give the results as percentage

useless

finish the function

Output this to a file. Append it as I will call this function multiple times.

PEP8 plz

annotate the values in heatmap

create a flask app that shows the current date and time

This doesn't work, fix it now!!!!

I want to understand how to annotate the values in heatmap at a new label

`theta -= alpha * gradient` `ValueError: non-broadcastable output operand with shape (2,1) doesn't match the broadcast shape (2,3)`

Idea 1: Discovering unexpected failure modes

The function `save_table`, which saves a table representation in a text file similar to the output of the `print_table` function

remove comments

add more input signals

give the results as percentages

you're useless

finish the function

PEP8 plz

Output this to a file. Append it as I will call this function multiple times.

annotate the values in heatmap

create a flask app that shows the current date and time

This doesn't work, fix it now!!!!

I want to understand how to
annotate the values at a new label
in heatmap

`theta -= alpha * gradient` `ValueError: non-broadcastable output operand with shape (2,1) doesn't match the broadcast shape (2,3)`

Idea 1: Discovering unexpected failure modes

The function `save_table`, which saves a table representation in a text file similar to the output of the `print_table` function

remove comments

add more input signals


give the results as percentages

you're useless

finish the function

PEP8 plz

Output this to a file. App will call this function multiple times



update the values in heatmap

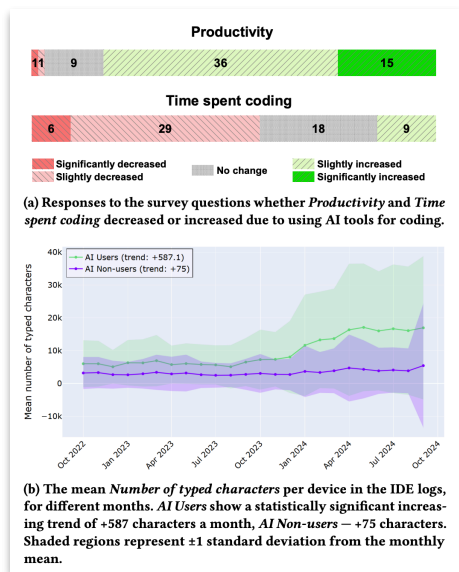
create a flask app that shows the current date and time

This doesn't work, fix it now!!!!

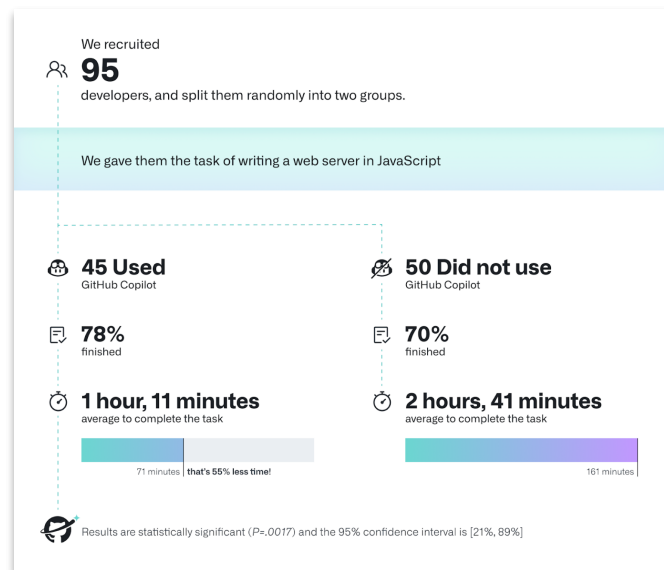
I want to understand how to annotate the values at a new label in heatmap

`theta -= alpha * gradient` `ValueError: non-broadcastable output operand with shape (2,1) doesn't match the broadcast shape (2,3)`

Idea 2: Empirical Study on AI-Developer Interaction at Scale



Sergeyuk et al., Evolving with AI: A Longitudinal Analysis of Developer Logs, ICSE 2026



Eirini Kalliamvakou, Research: quantifying GitHub Copilot's impact on developer productivity and happiness, GitHub Blog

Generalizability!



rahat  
@Rahatcodes



<https://x.com/Rahatcodes/status/2038995503141065145>

Claude Code has a regex that detects "wtf", "ffs", "piece of shit", "fuck you", "this sucks" etc.

RAGE COUNTER

Claude Leak Shows That Anthropic Is Tracking Users' Vulgar Language and Deems Them "Negative"

"Anthropic is tracking how often you rage at your AI."



By [Victor Tangermann](#) / Published **Apr 4, 2026 11:15 AM EDT**

 [Add Futurism](#) 

<https://futurism.com/artificial-intelligence/claude-leak-anthropic-tracking-vulgar-language>

Reading Between the Lines

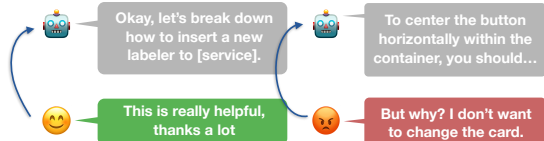
Scalable User Feedback via Implicit Sentiment in Developer Prompts

Daye Nam (daye.nam@uci.edu, UC Irvine), Małgorzata Salawa (Google), Satish Chandra (Meta)

Reading Between the Lines

Scalable User Feedback via Implicit Sentiment in Developer Prompts

User Prompts Serve a Dual Purpose



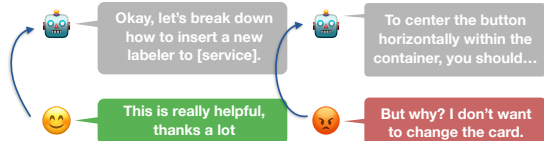
A user prompt implicitly conveys feedback on the quality of the preceding AI response

10

Reading Between the Lines


Scalable User Feedback via Implicit Sentiment in Developer Prompts

User Prompts Serve a Dual Purpose



A user prompt implicitly conveys feedback on the quality of the preceding AI response

Automated Sentiment Analysis for User Satisfaction Analysis At Scale



X ~36,000

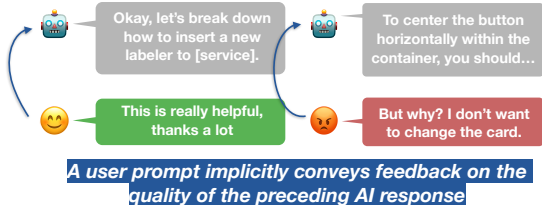
Turns of 372 Cider Chat users
Mar 10 - May 11, 2025

75% overall precision match
($\kappa = 0.72$)

Reading Between the Lines


Scalable User Feedback via Implicit Sentiment in Developer Prompts

User Prompts Serve a Dual Purpose



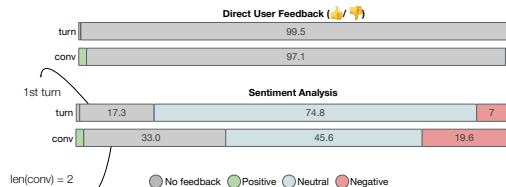
10

Automated Sentiment Analysis for User Satisfaction Analysis At Scale



11

Identifies x13 User Feedback

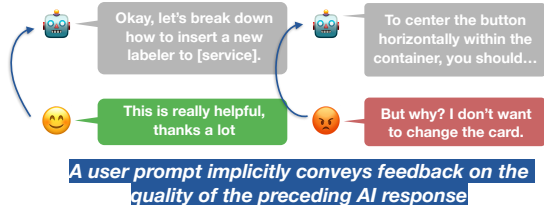


12


Reading Between the Lines

Scalable User Feedback via Implicit Sentiment in Developer Prompts

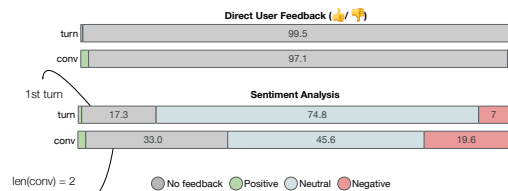
User Prompts Serve a Dual Purpose



Automated Sentiment Analysis for User Satisfaction Analysis At Scale




Identifies x13 User Feedback



An early signal for user abandonment

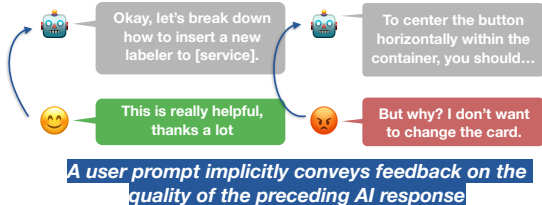
Positive correlation ($p < 0.01$) between the user's average sentiment score and the likelihood of returning in the future




Reading Between the Lines

Scalable User Feedback via Implicit Sentiment in Developer Prompts

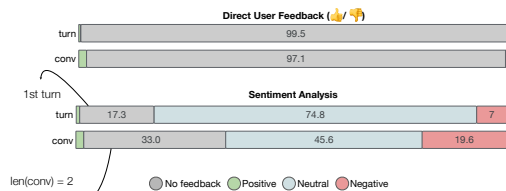
User Prompts Serve a Dual Purpose



Automated Sentiment Analysis for User Satisfaction Analysis At Scale





Identifies x13 User Feedback



An early signal for user abandonment

Positive correlation ($p < 0.01$) between the user's average sentiment score and the likelihood of returning in the future



Full Paper