



Code Generation, guest lecture

Human Interactions with Code Gen Models

Sherry Tongshuang Wu

HCI/LTI

@tongshuangwu / sherryw@cs.cmu.edu

Outline

Evaluation:

Metrics inspired by human-human interactions
Quantitative and qualitative user modeling

Design and implementation:

The impact of UI
The focus on user needs

How do you know a code gen model is **good?**

Metrics evaluating code gen models

Outcomes	Human-AI (Copilot)
Quality	<p>🔴 vs. Human-Human: more lines of code deleted in next session (lower quality) [9]</p> <p>🟢 vs. Human Solo: significantly improve correctness score and reduce encountered errors for novice students [10]</p> <p>🟡 vs. Human Solo: no significant difference in task success [11] or task success rate in given time [12]</p>
Productivity	<p>🟢 vs. Human-Human: more lines of added code [9]</p> <p>🟢 vs. Human Solo: 55.8% reduction in completion time [11]</p> <p>🟢 vs. Human Solo: significantly increase task completion and reduce task completion time for novice students [10]</p> <p>🟡 vs. Human Solo: no significant difference in the task completion rate in given time [12]</p>
Satisfaction	<p>🟢 vs. Human Solo: higher self-ratings of satisfaction [12, 16, 17]</p>
Learning	<p>🟡 vs. Human Solo: no significant difference in immediate and retention post-test performance of novices, students with more prior experiences have more learning gains from AI code generator [10]</p>
Cost	<p>No experiment yet. Vaithilingam et al. [12], Bird et al. [16] hypothesized that human-AI may lead to more unnecessary debugging vs. Human Solo</p>

How do you know a code gen model is **useful?**


A: define metrics more specific for usability tests!

Human-Human vs. Human-AI pAIr Programming

Human-AI pAIr Programming:

Programmer and LLM work together at the same computer, solving the same task.

Copilot, an LLM-powered programming assistance tool, advertises itself as **"your AI pair programmer"**




Human-Human Pair Programming

Two programmers work together on the same task using a single device. [Beck, 1999]

Driver: performs the coding

Navigator: aids in planning, reviewing, debugging



Similar to human-human co-programming, human-AI pair programming involves a lot of study and metrics that should capture the interaction aspect.

Qianou Ma, Tongshuang Wu, and Kenneth Koedinger. "Is AI the better programming partner? Human-Human Pair Programming vs. Human-AI pAIr Programming." *AIED 2023 workshop*

Metrics for human-AI interaction

Outcomes	Human-AI (Copilot)
Quality	<p>😞 vs. Human-Human: more lines of code deleted in next session (lower quality) [9]</p> <p>😊 vs. Human Solo: significantly improve correctness score and reduce encountered errors for novice students [10]</p> <p>😐 vs. Human Solo: no significant difference in task success [11] or task success rate in given time [12]</p>
Productivity	<p>😊 vs. Human-Human: more lines of added code [9]</p> <p>😊 vs. Human Solo: 55.8% reduction in completion time [11]</p> <p>😊 vs. Human Solo: significantly increase task completion and reduce task completion time for novice students [10]</p> <p>😐 vs. Human Solo: no significant difference in the task completion rate in given time [12]</p>
Satisfaction	<p>😊 vs. Human Solo: higher self-ratings of satisfaction [12, 16, 17]</p>
Learning	<p>😐 vs. Human Solo: no significant difference in immediate and retention post-test performance of novices, students with more prior experiences have more learning gains from AI code generator [10]</p>
Cost	<p>No experiment yet. Vaithilingam et al. [12], Bird et al. [16] hypothesized that human-AI may lead to more unnecessary debugging vs. Human Solo</p>

Outcomes	Human-Human vs. Human Solo	Human-AI (Copilot)
Quality	<ul style="list-style-type: none"> 😊 significantly lower defect density for complex code [7] 😐 no difference for simpler code [7] 😊 significantly higher percentage of test cases passed [8] 	<ul style="list-style-type: none"> 😞 vs. Human-Human: more lines of code deleted in next session (lower quality) [9] 😊 vs. Human Solo: significantly improve correctness score and reduce encountered errors for novice students [10] 😐 vs. Human Solo: no significant difference in task success [11] or task success rate in given time [12]
Productivity	<ul style="list-style-type: none"> 😞 significantly fewer lines of code per person hour writing simpler code [7] 😐 no significant difference writing more complex code [7] 😊 29% shorter time to complete task (pair speed advantage = 1.4) [13] 	<ul style="list-style-type: none"> 😊 vs. Human-Human: more lines of added code [9] 😊 TL;DR: there are many metrics that we can adapt from human-human interactions! 😊 vs. Human Solo: significantly increase task completion and reduce task completion time for novice students [10] 😐 vs. Human Solo: no significant difference in the task completion rate in given time [12]
Satisfaction	<ul style="list-style-type: none"> 😊 higher self-ratings of satisfaction [14] 😞 students with greater self-confidence and self-efficacy less enjoy the pair programming experience [15] 	<ul style="list-style-type: none"> 😊 vs. Human Solo: higher self-ratings of satisfaction [12, 16, 17] Multi-dimension reflection on quality Connection to real-world impact (exam) Reflection on real user benefit (teaching staff workload!)
Learning	<ul style="list-style-type: none"> 😊 higher grades, exam scores [18], and retention [19] 😊 significantly higher gains in exam performance in female students than male students [20] 	<ul style="list-style-type: none"> 😐 vs. Human Solo: no significant difference in immediate and retention post-test performance of novices, students with more prior experiences have more learning gains from AI code generator [10]
Cost	<ul style="list-style-type: none"> 😞 increased management workload to match, schedule a pair, resolve collaboration conflict, assess individual contributions, etc. [21] 😊 reduced teaching staff workload (grading one assignment from a pair) [8] 	<ul style="list-style-type: none"> No experiment yet. Vaithilingam et al. [12], Bird et al. [16] hypothesized that human-AI may lead to more unnecessary debugging vs. Human Solo

Example Usability Metric: Syntactic similarity metric

Reference Code Snippet

```
def even_odd_count(num):  
    even_count = 0  
    odd_count = 0  
    for i in str(abs(num)):  
        if int(i)%2==0:  
            even_count +=1  
        else:  
            odd_count +=1  
    return (even_count, odd_count)
```

Generated Code Snippet

```
def even_odd_count(num):  
    even_count = 0  
    odd_count = 0  
    for i in str(num):  
        if int(i) % 2 == 0:  
            even_count += 1  
        else:  
            odd_count += 1  
    return even_count, odd_count
```

Functional Metric

pass = **0**

Similarity Metric

edit similarity = **0.93**

Human preference

preference = **0.9**

Figure 1: In the example above (counting even and odd numbers), code suggested by a model fails unit tests but is deemed useful by programmers because adding a short check (*abs* value) fixes the generation.

“While correctness captures high-value generations, programmers still rate **code that fails unit tests as valuable if it reduces the overall effort needed to complete a coding task**. Finally, we propose a hybrid metric that combines functional correctness and syntactic similarity and show that it achieves a 14% stronger correlation with value and can therefore better represent real-world gains when evaluating and comparing models.”

How do you know a code gen model is **useful for...**

A: Quantitative and qualitative modeling!

Researchers

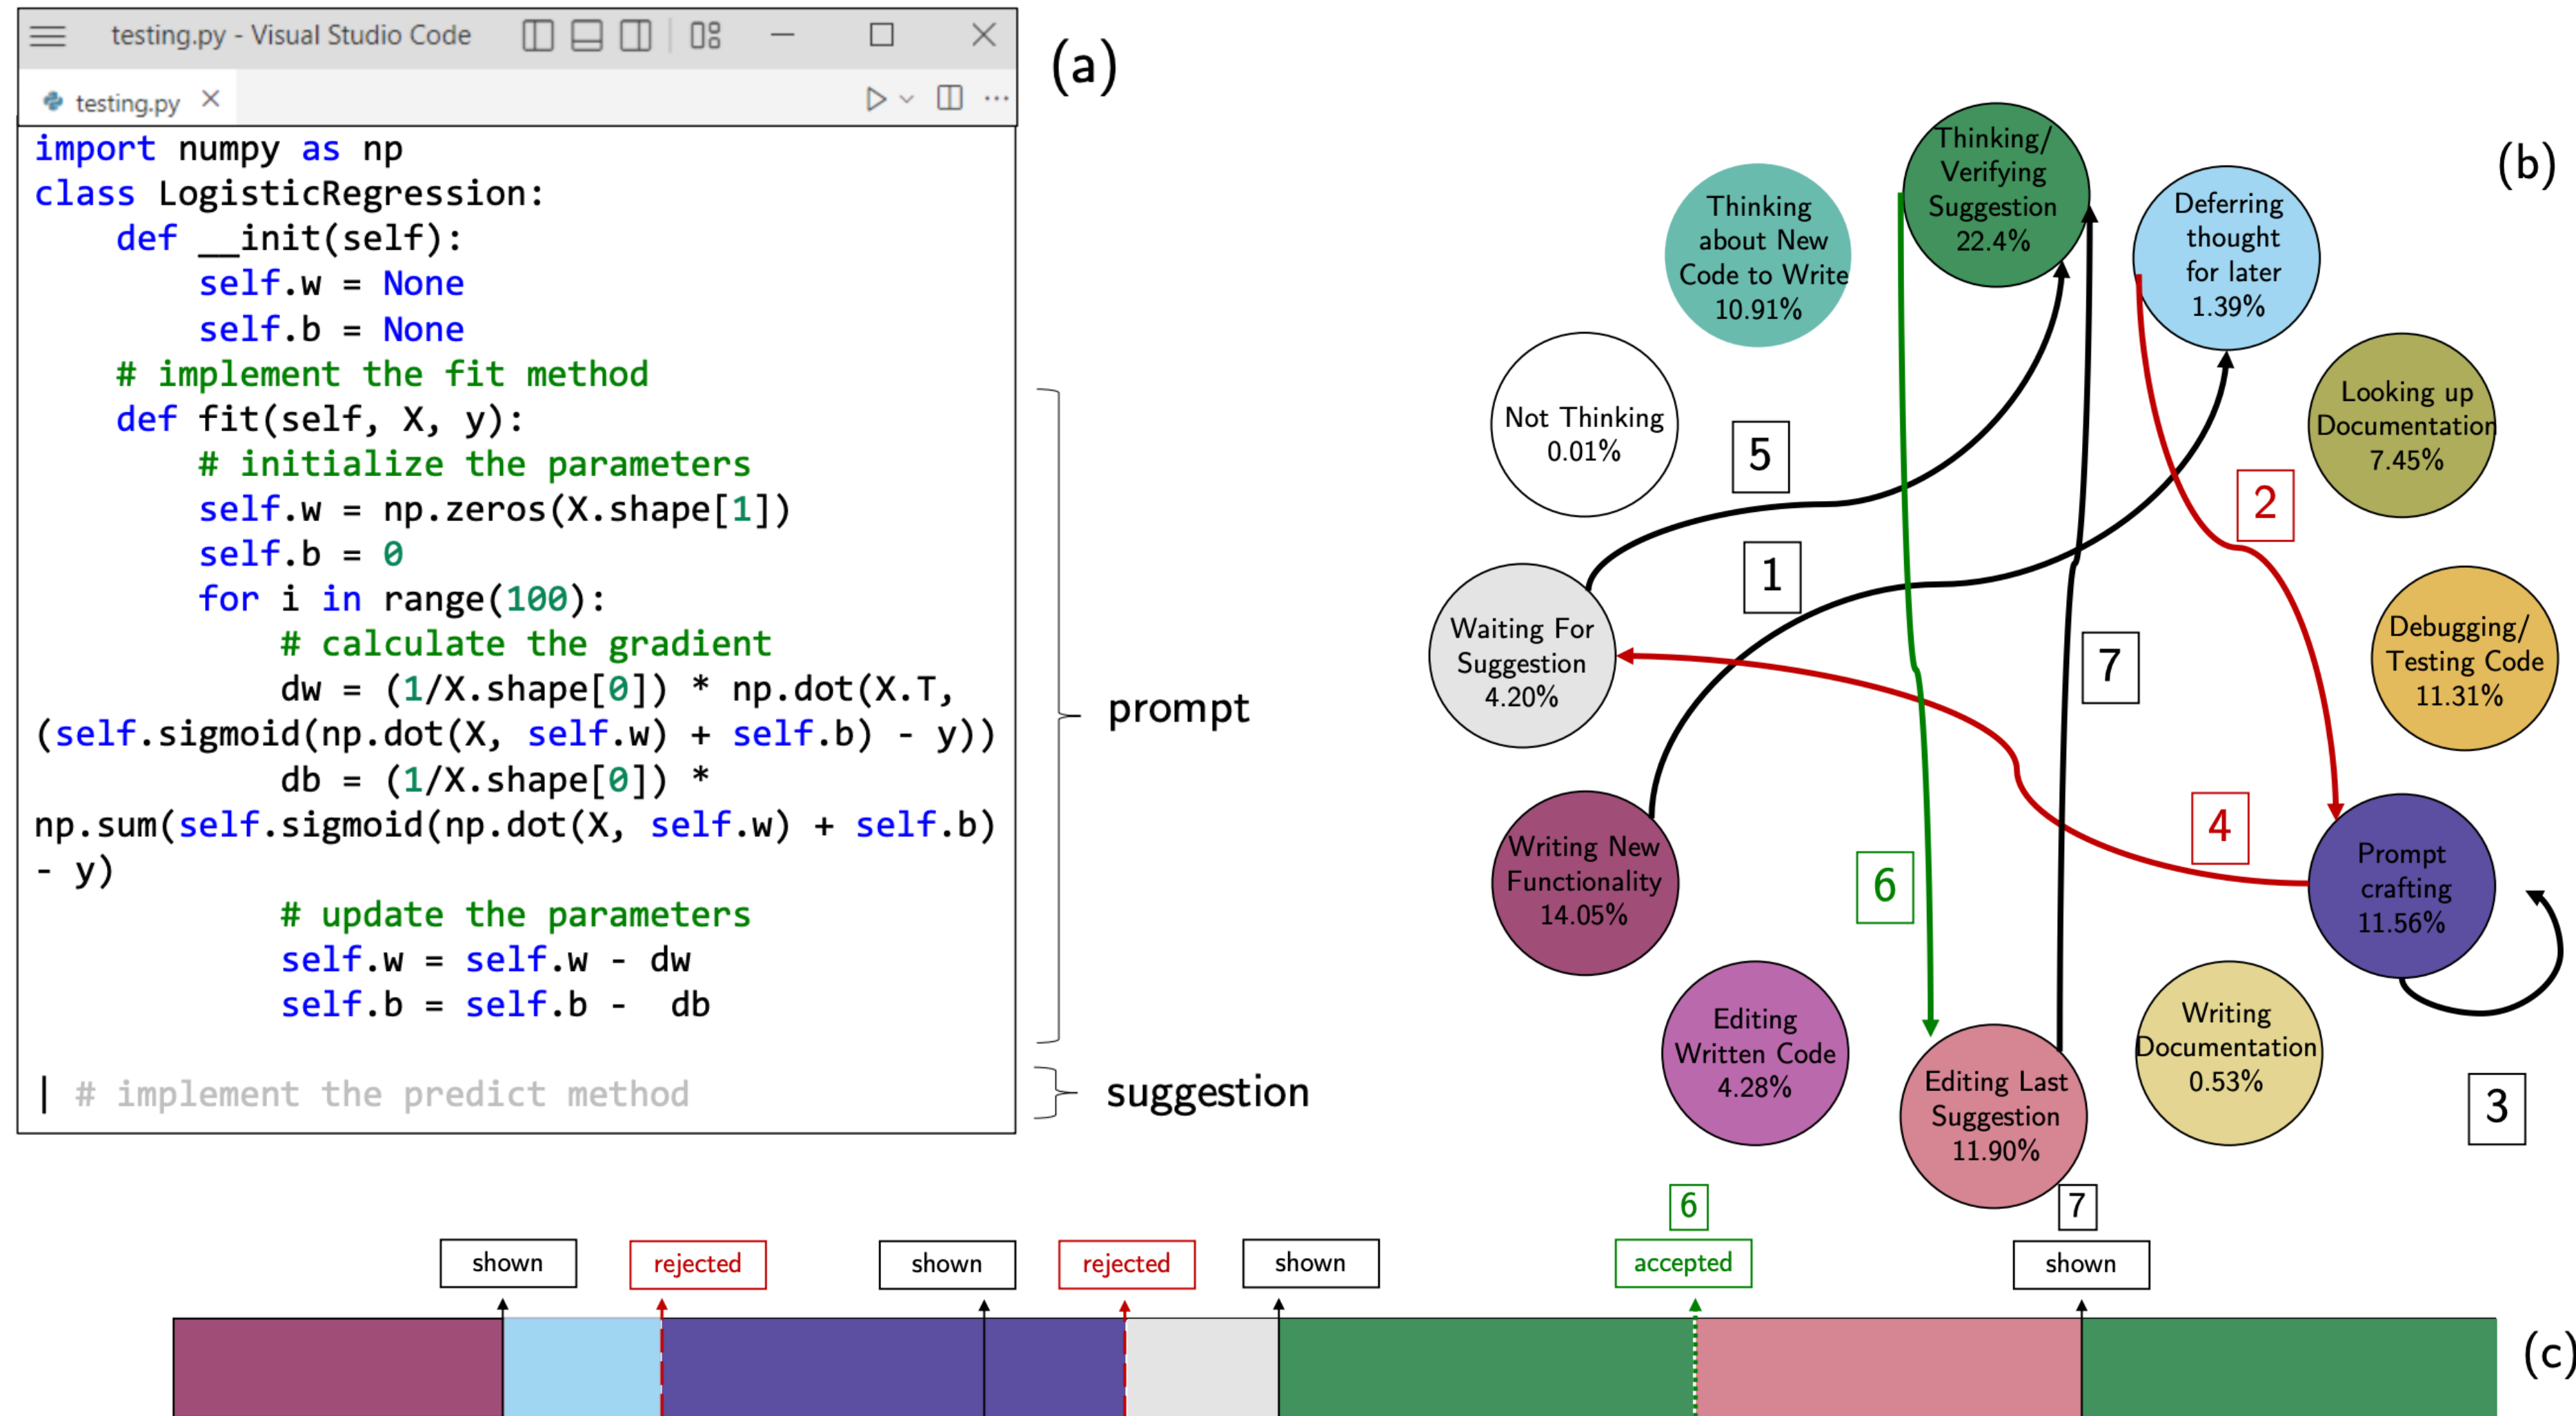
CS1 students

Junior engineers

Senior engineers

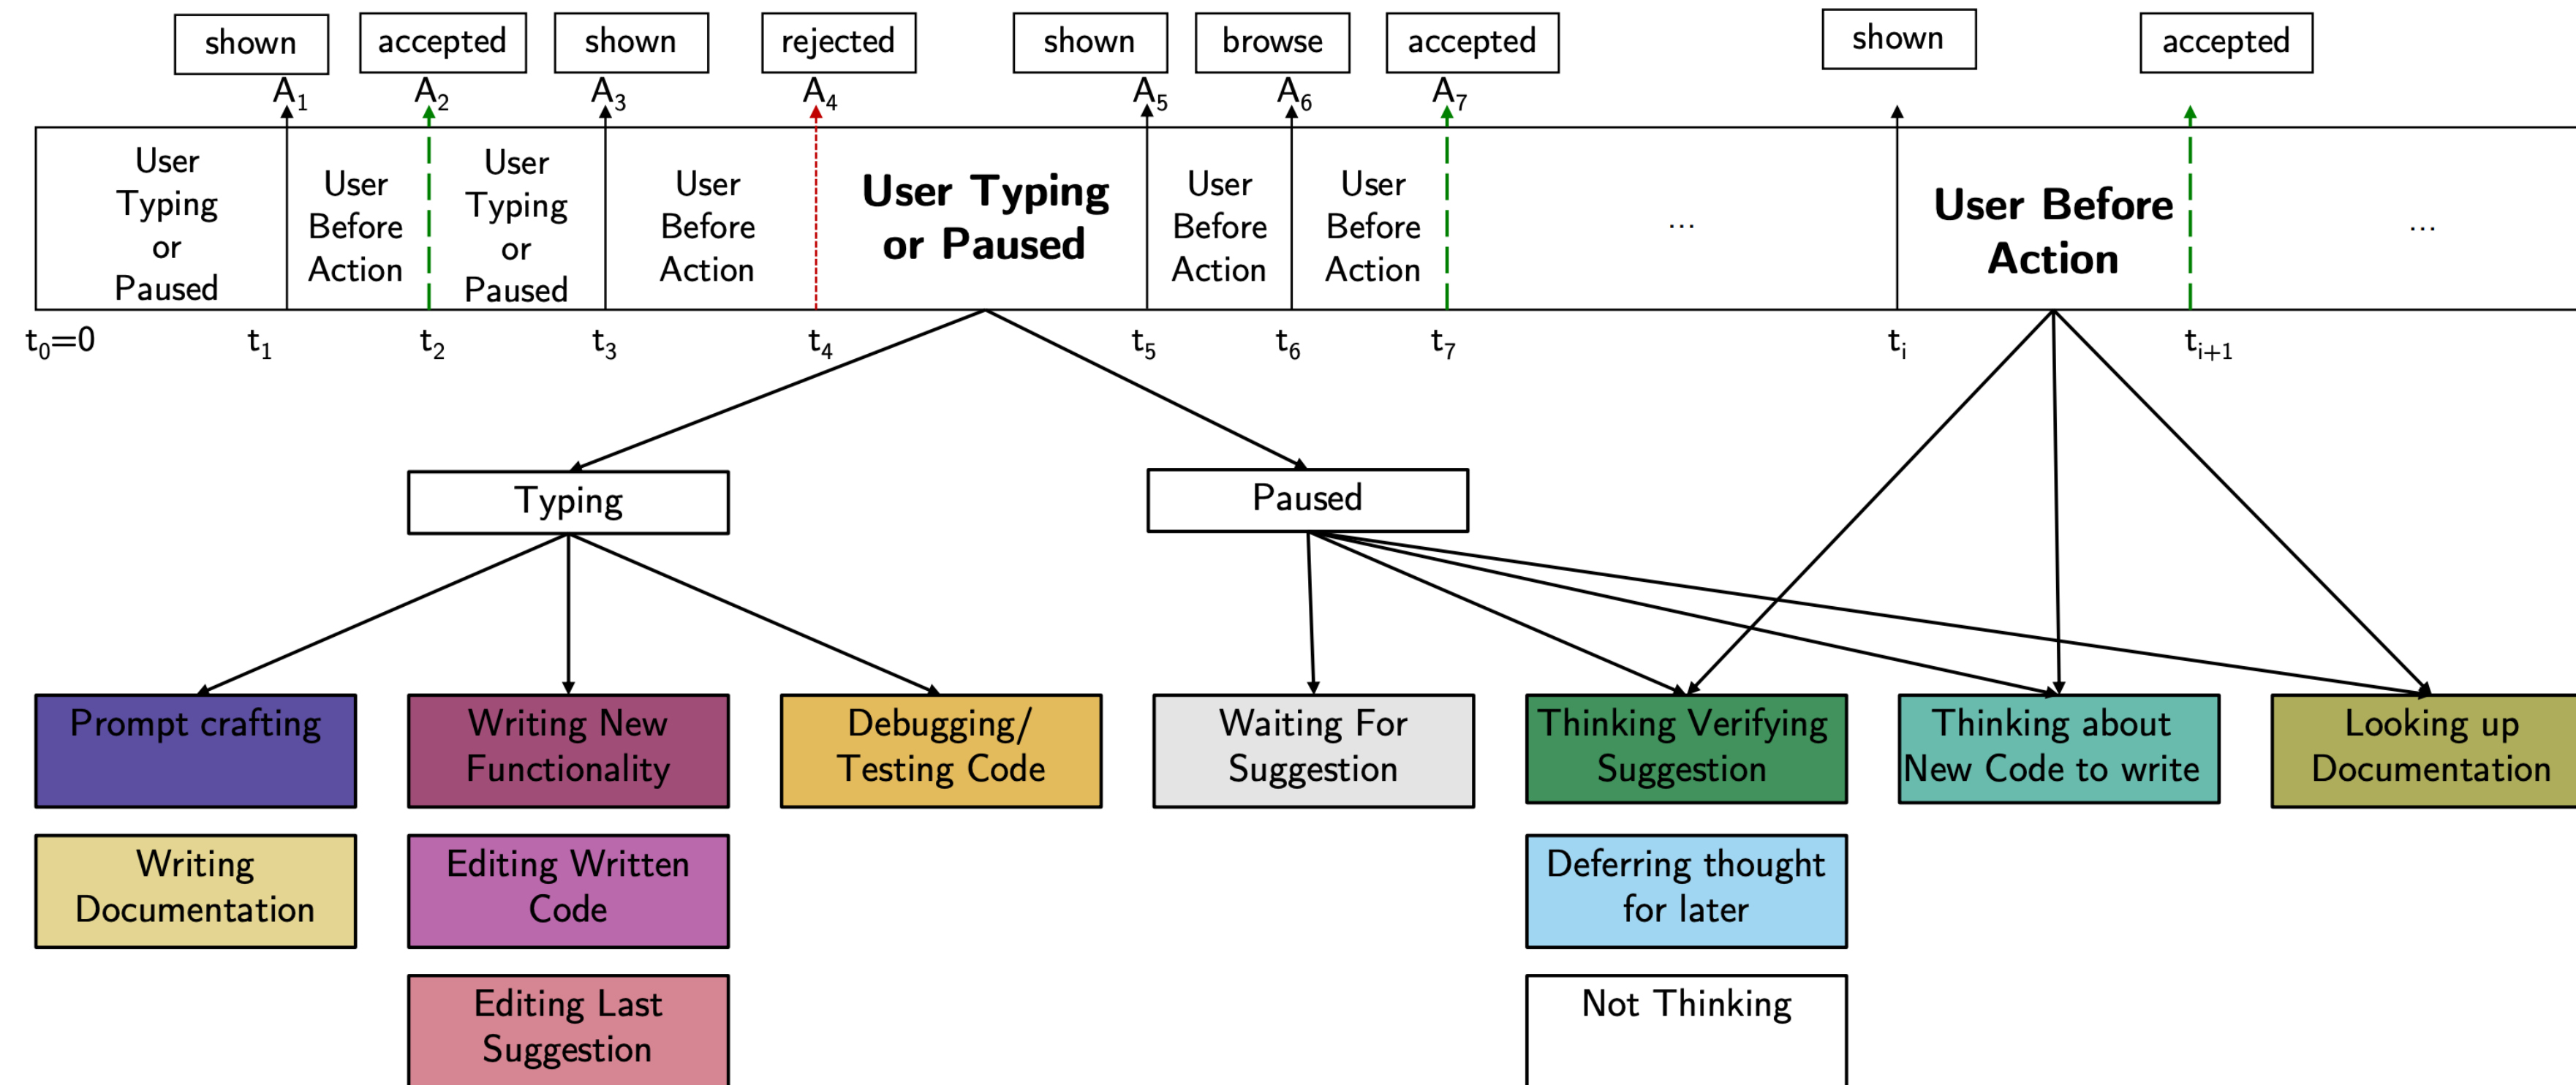
...

Study humans quantitatively: User modeling on Clickstream



Mozannar, Hussein, et al. "Reading between the lines: Modeling user behavior and costs in AI-assisted programming." *CHI 2024*


Study humans quantitatively: User modeling on clickstream



Can define and classify what happens in programmer actions (here idle times)

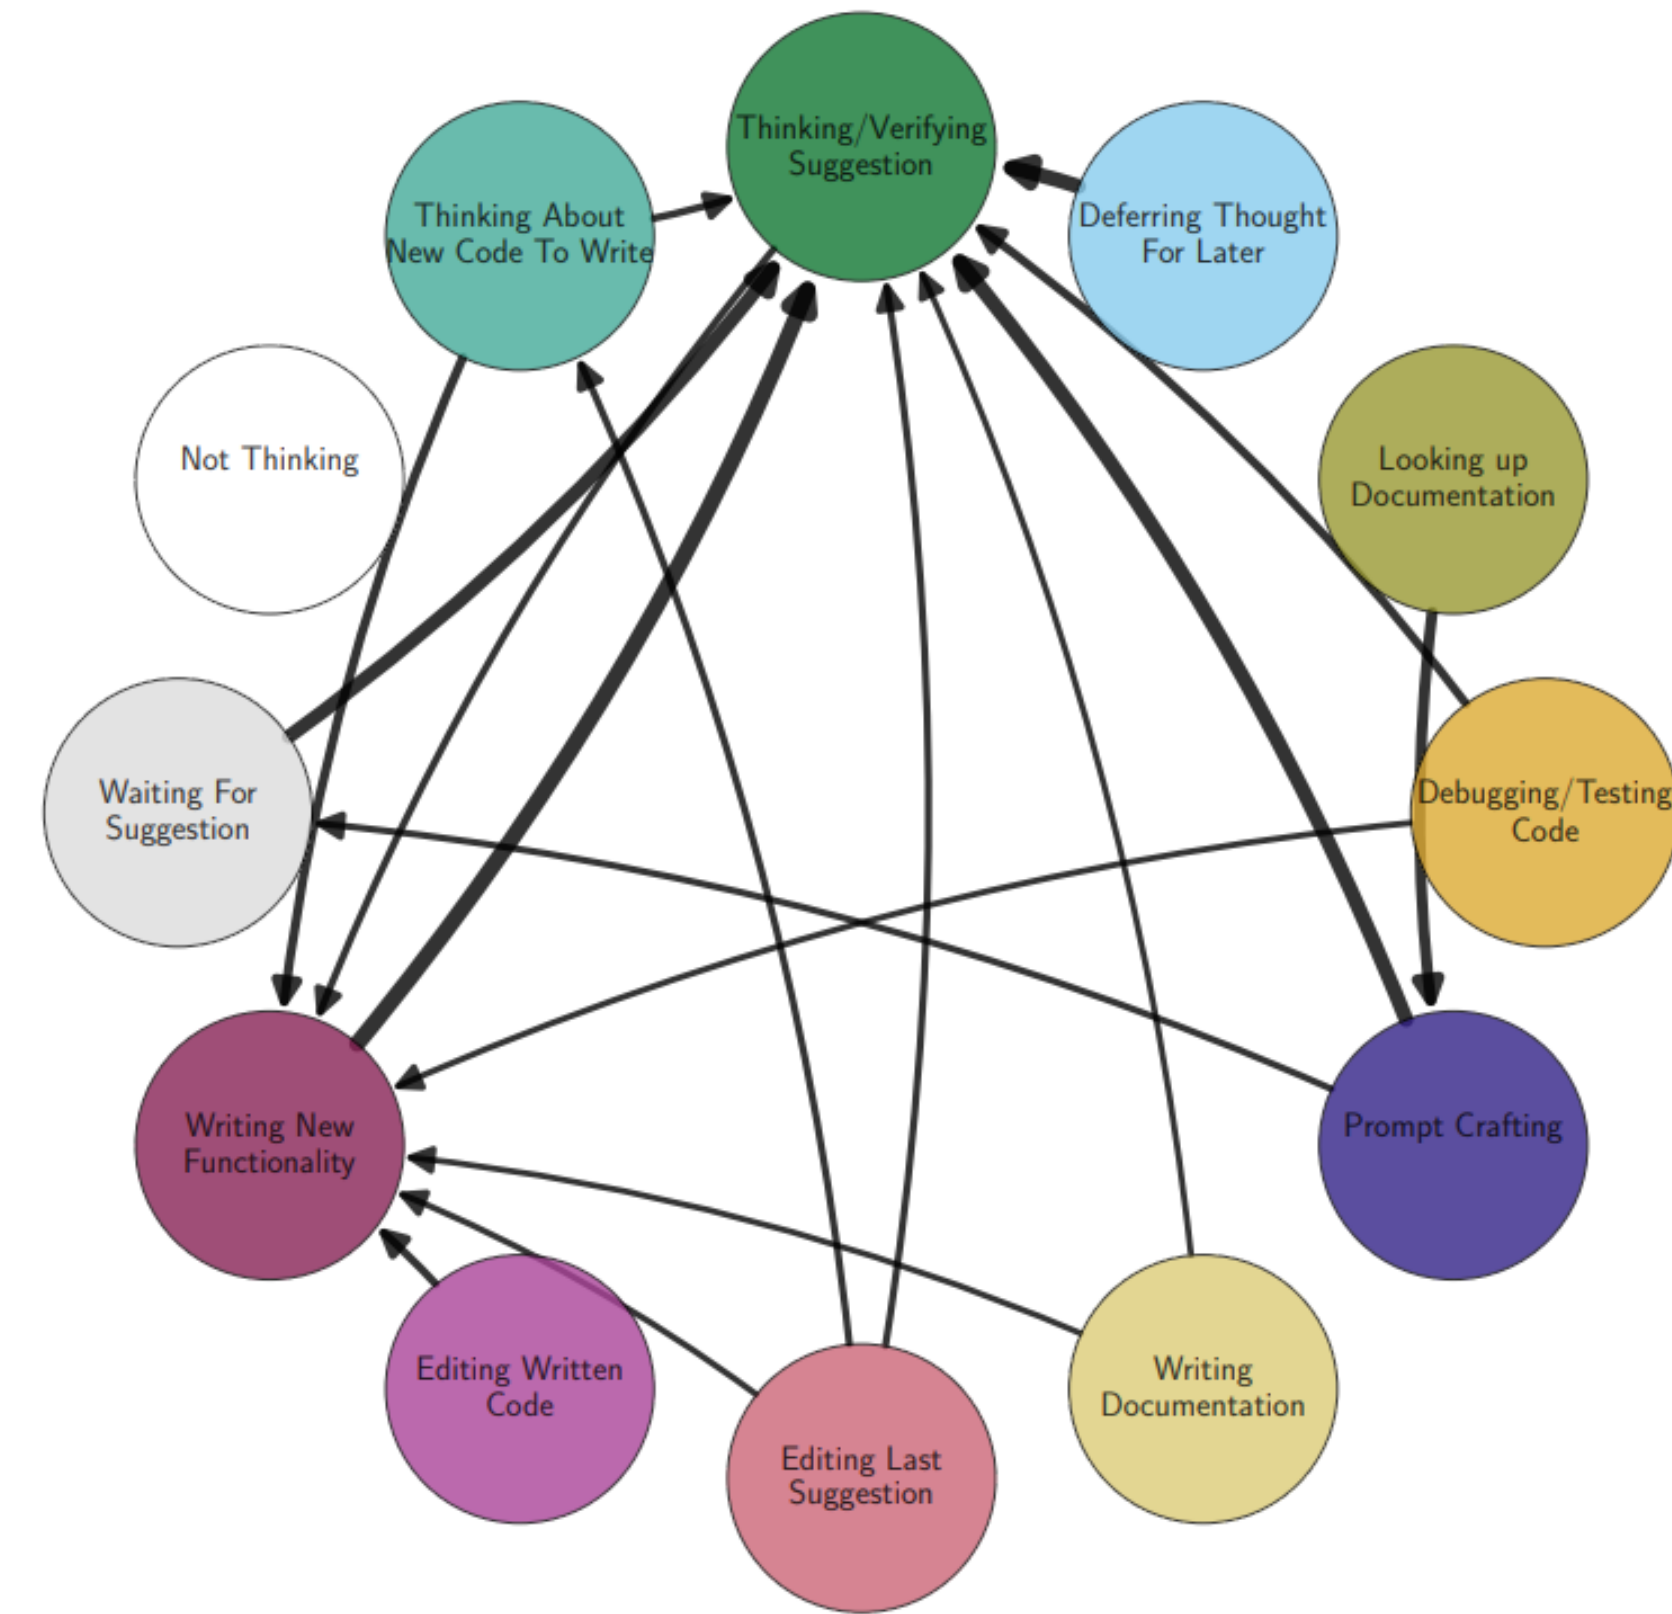
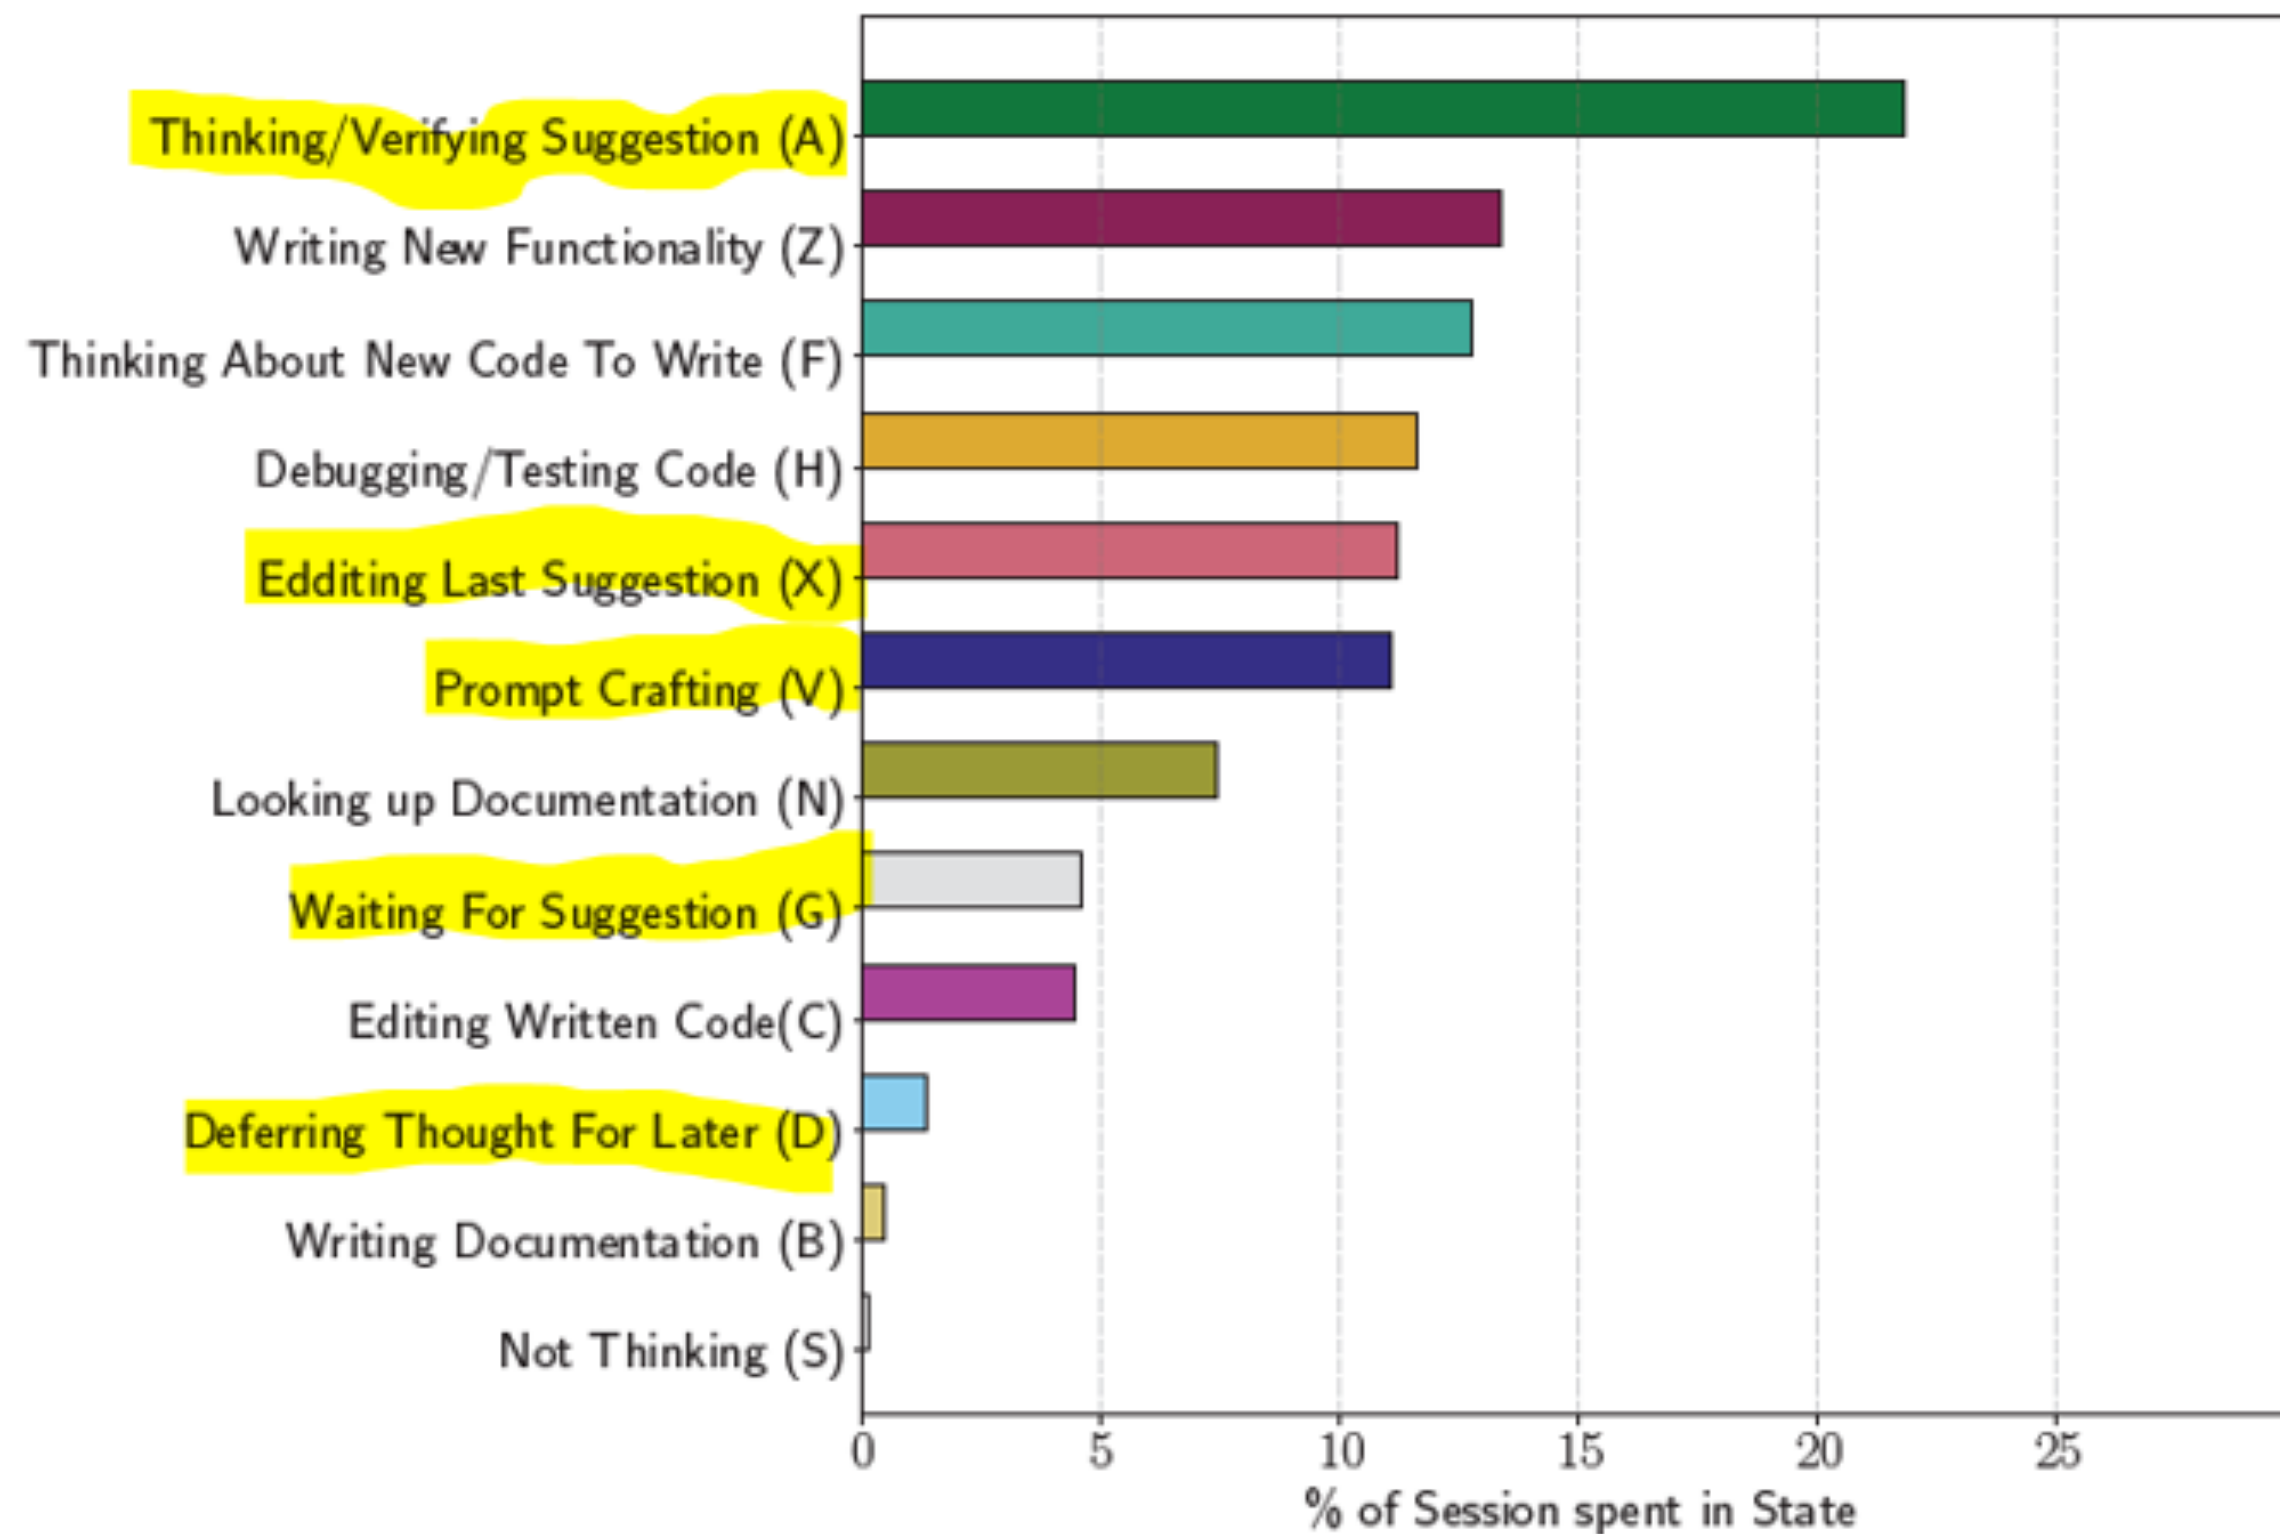
Study humans quantitatively: User modeling on clickstream

Can uncover interesting patterns for each individual people!



(a) Individual CUPS timelines for 5/21 study participants for the first 180 secs show the richness of and variance in programmer-CodeRec interaction.

Human behavior in aggregation can show avg use patterns



“In a study with 21 programmers, we saw that the most time intensive state is verifying suggestions, and Copilot related states (yellow highlight) occupy on average 51% of task time.”

Clickstream can uncover interesting patterns!

“Copilot often forces programmer to accept a sequence of suggestions in a row, teasing them to show the full function/class body, which makes them verify suggestion after accepting them (rather than before)”

```
class LogisticRegression:
    def __init__(self,X,y,alpha=0.01):
        self.X = X
        self.y = y
        self.alpha = alpha
        self.theta = np.zeros(X.shape[1])
        self.cost = []
        self.theta_history = []
```

Open left brace [indicates that suggestion is not a complete code segment

→
4 single line
Accepts later

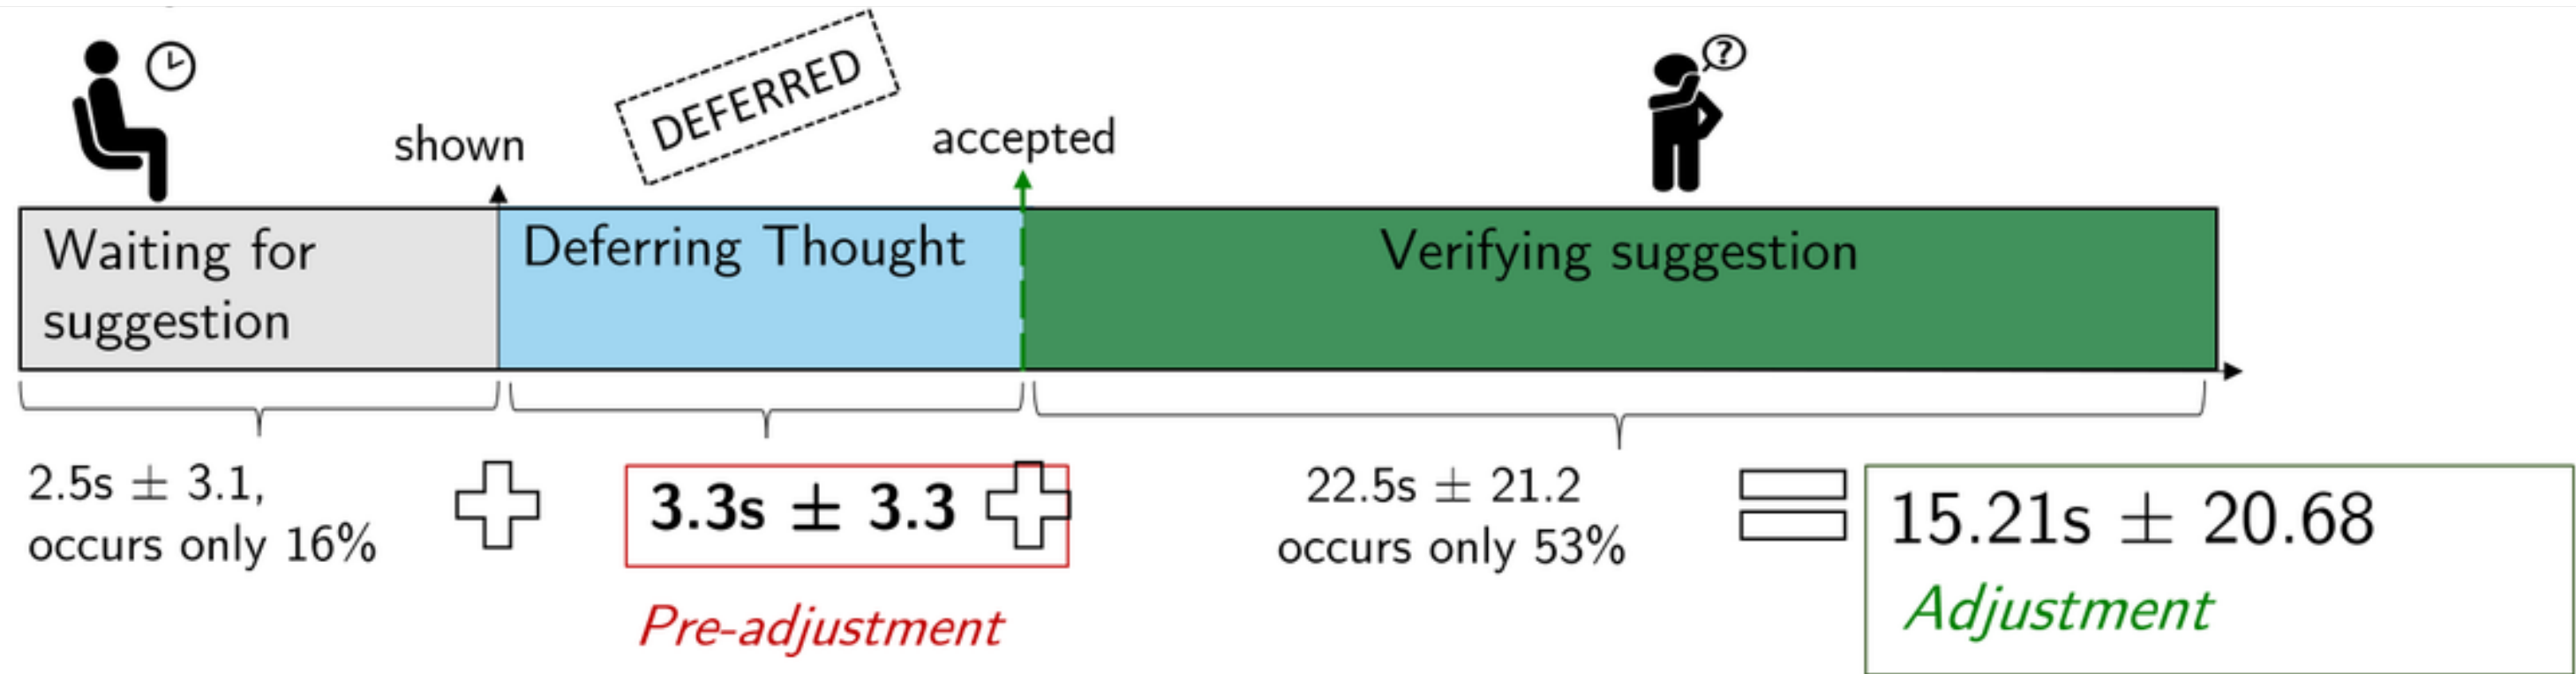
```
class LogisticRegression:
    def __init__(self,X,y,alpha=0.01):
        self.X = X
        self.y = y
        self.alpha = alpha
        self.theta = np.zeros(X.shape[1])
        self.cost = []
        self.theta_history = [
            self.theta
```

Suggestion references a method cost() not yet implemented

→
3 single line
Accepts later

```
class LogisticRegression:
    def __init__(self,X,y,alpha=0.01):
        self.X = X
        self.y = y
        self.alpha = alpha
        self.theta = np.zeros(X.shape[1])
        self.cost = []
        self.theta_history = [
            self.theta
        ]
        self.cost_history = [
            self.cost()
        ]
    def cost(self):
        return (-1 / len(self.y)) *
            np.sum(self.y * np.log(self.hypothesis())) + ...
```

The function cost references a method hypothesis() not yet implemented



Also reveal issues in metrics

Outcomes	Human-Human vs. Human Solo	Human-AI (Copilot)
Productivity	<ul style="list-style-type: none">🔴 significantly fewer lines of code per person hour writing simpler code [7]😊 no significant difference writing more complex code [7]🟢 29% shorter time to complete task (pair speed advantage = 1.4) [13]	<ul style="list-style-type: none">🟢 vs. Human-Human: more lines of added code [9]🟢 vs. Human Solo: 55.8% reduction in completion time [11]🟢 vs. Human Solo: significantly increase task completion and reduce task completion time for novice students [10]😊 vs. Human Solo: no significant difference in the task completion rate in given time [12]

Human-human - Variance in metrics!

Time and accomplishment? twice the duration, the person-hours required

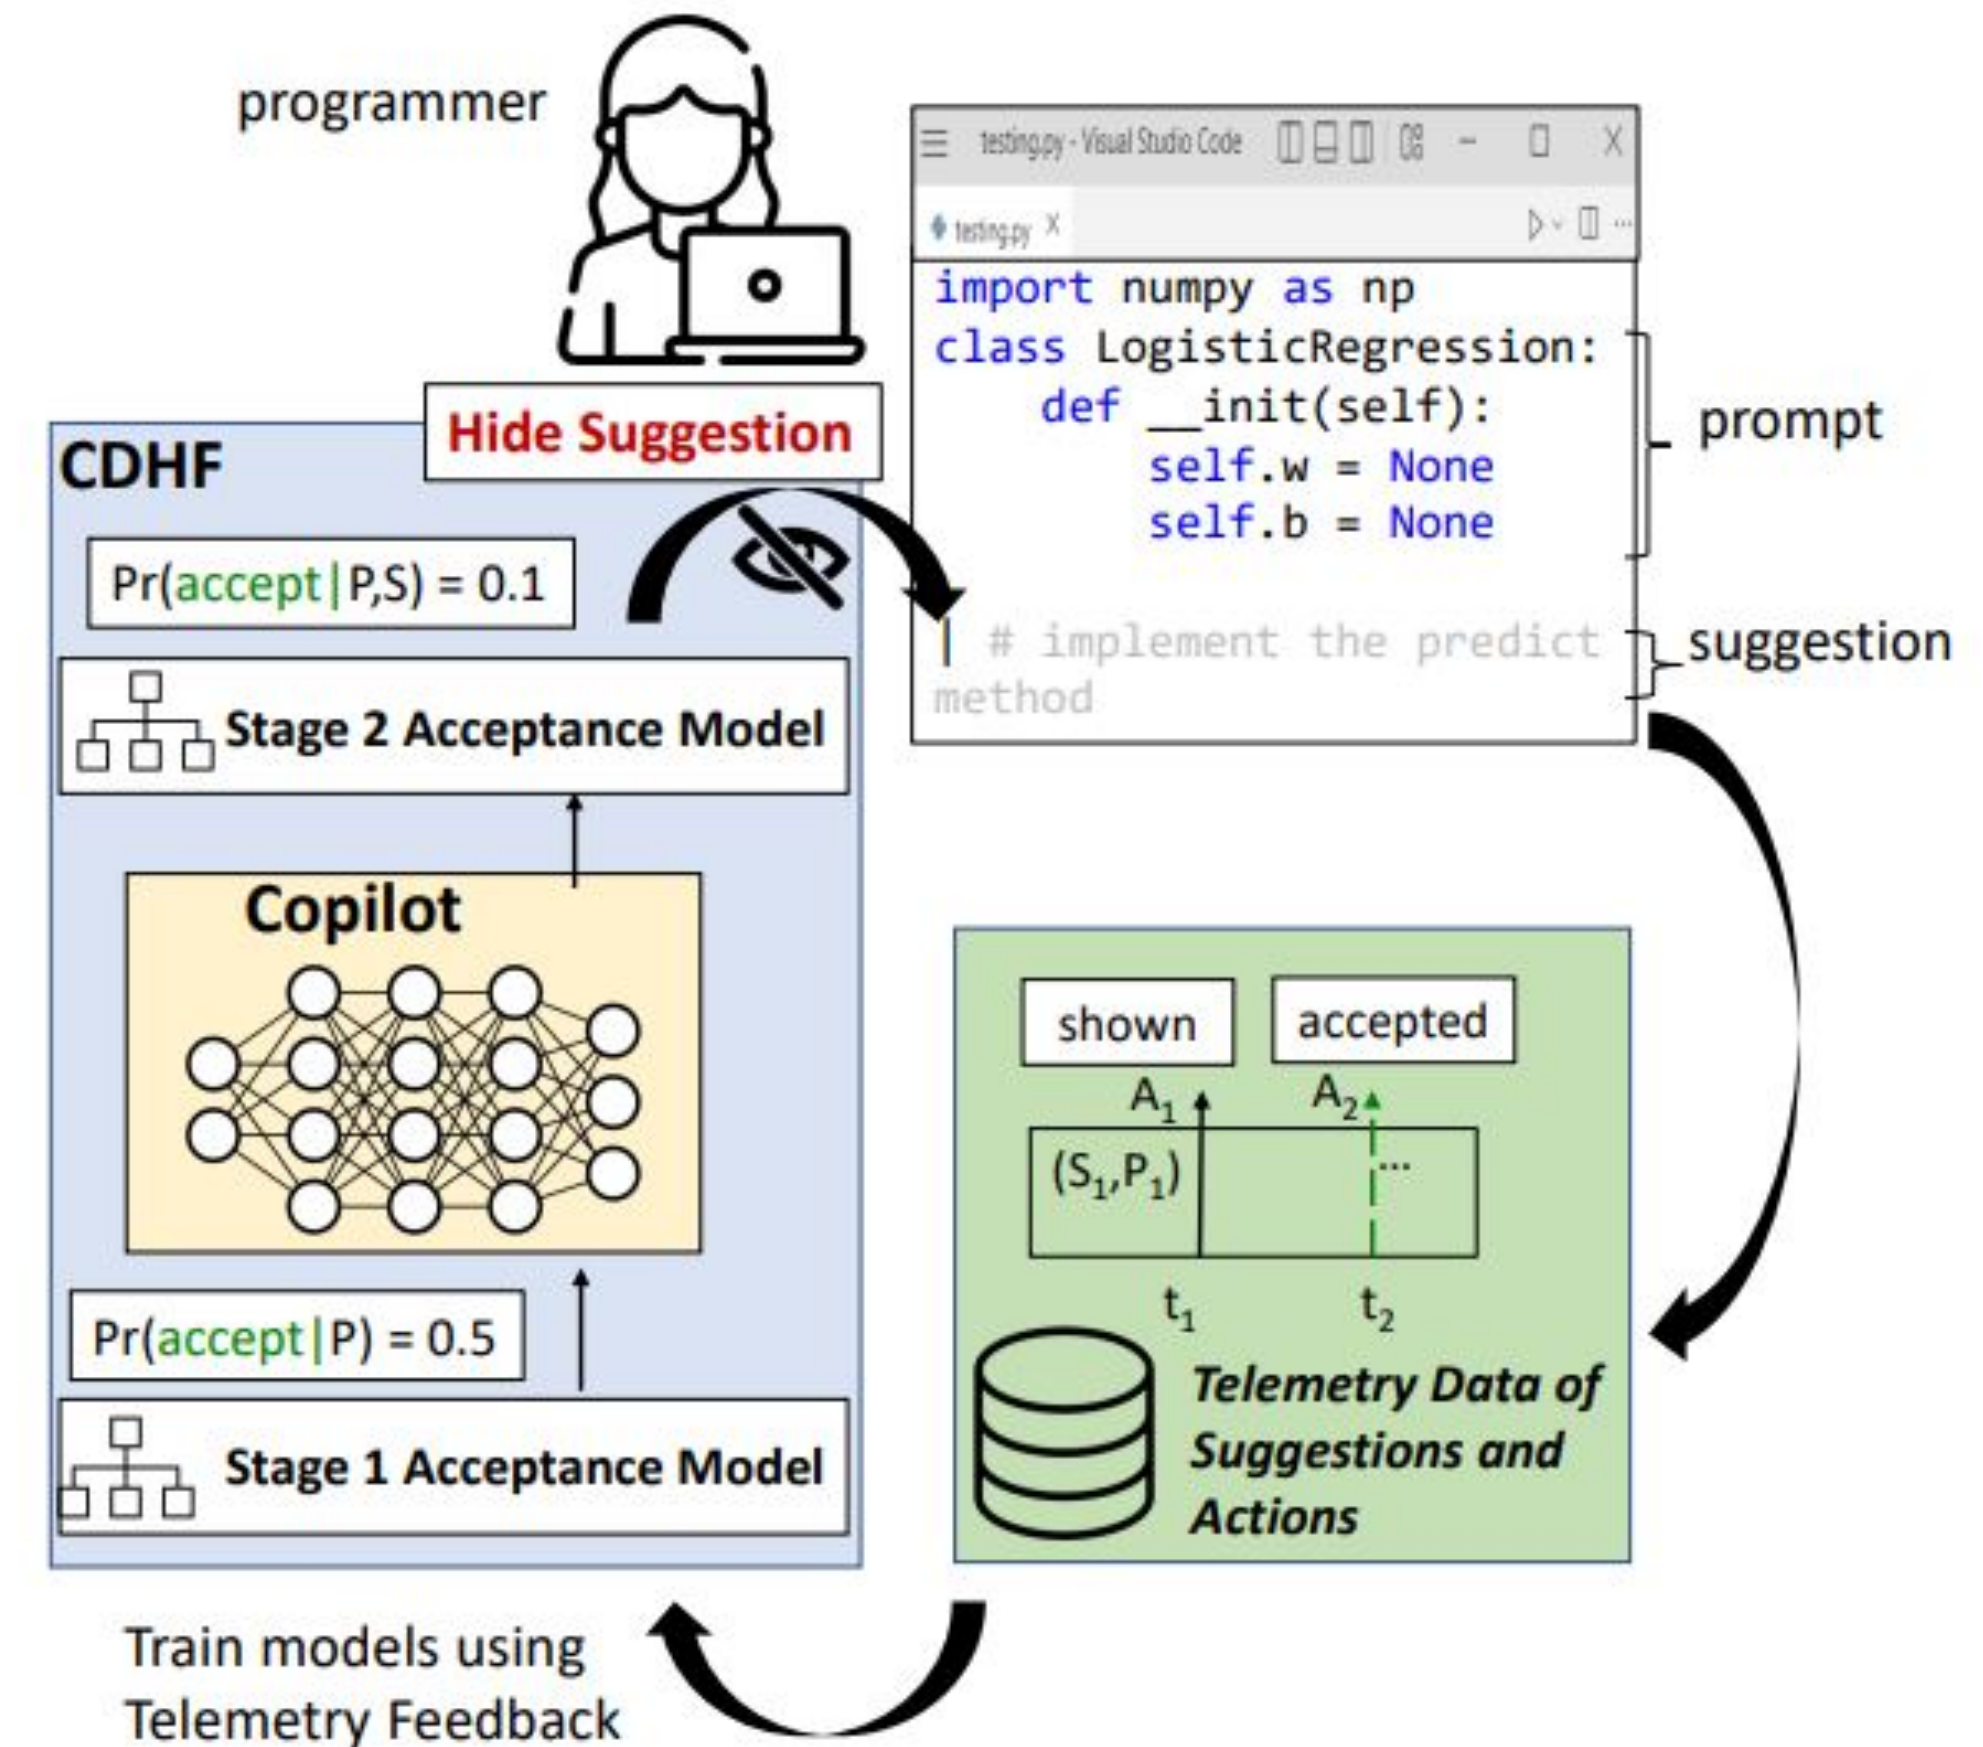
Human-AI - Too simplified metrics?

E.g. the number of lines of added code - the nature of interaction with Copilot (tab to accept suggestions) is a big factor!

Clickstream can help with design!

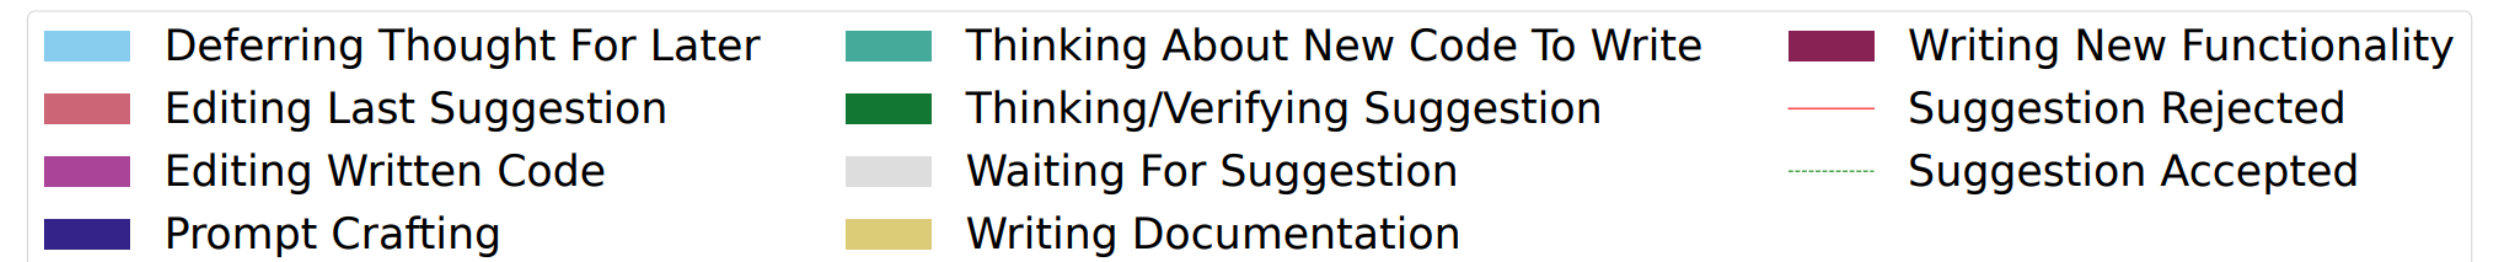

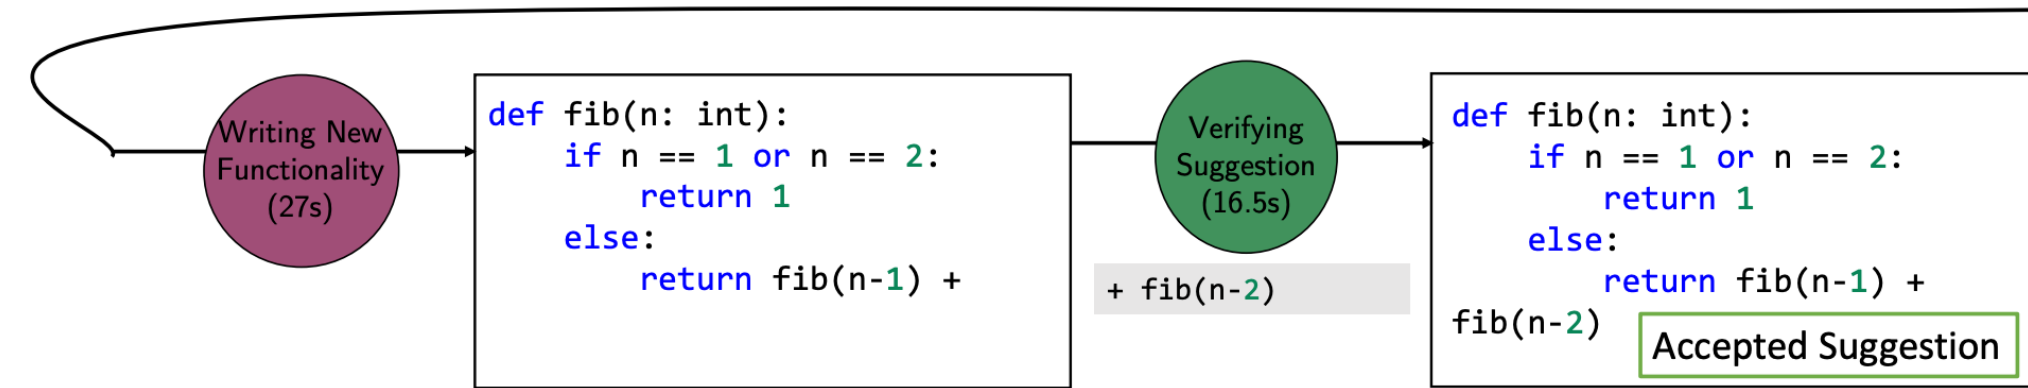
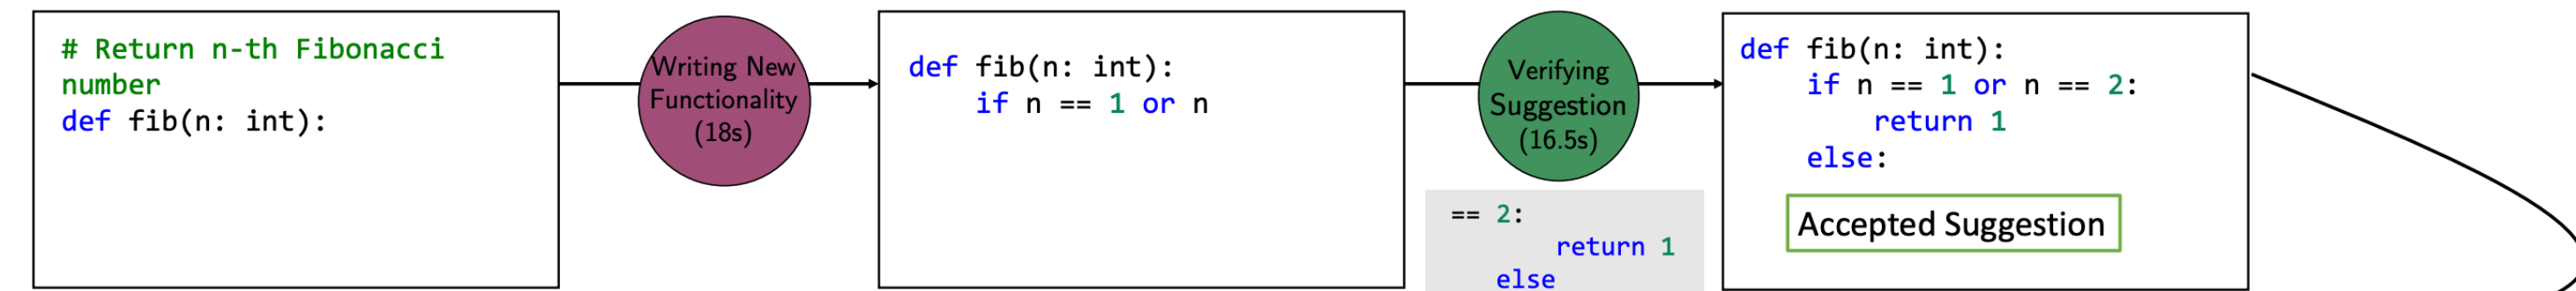
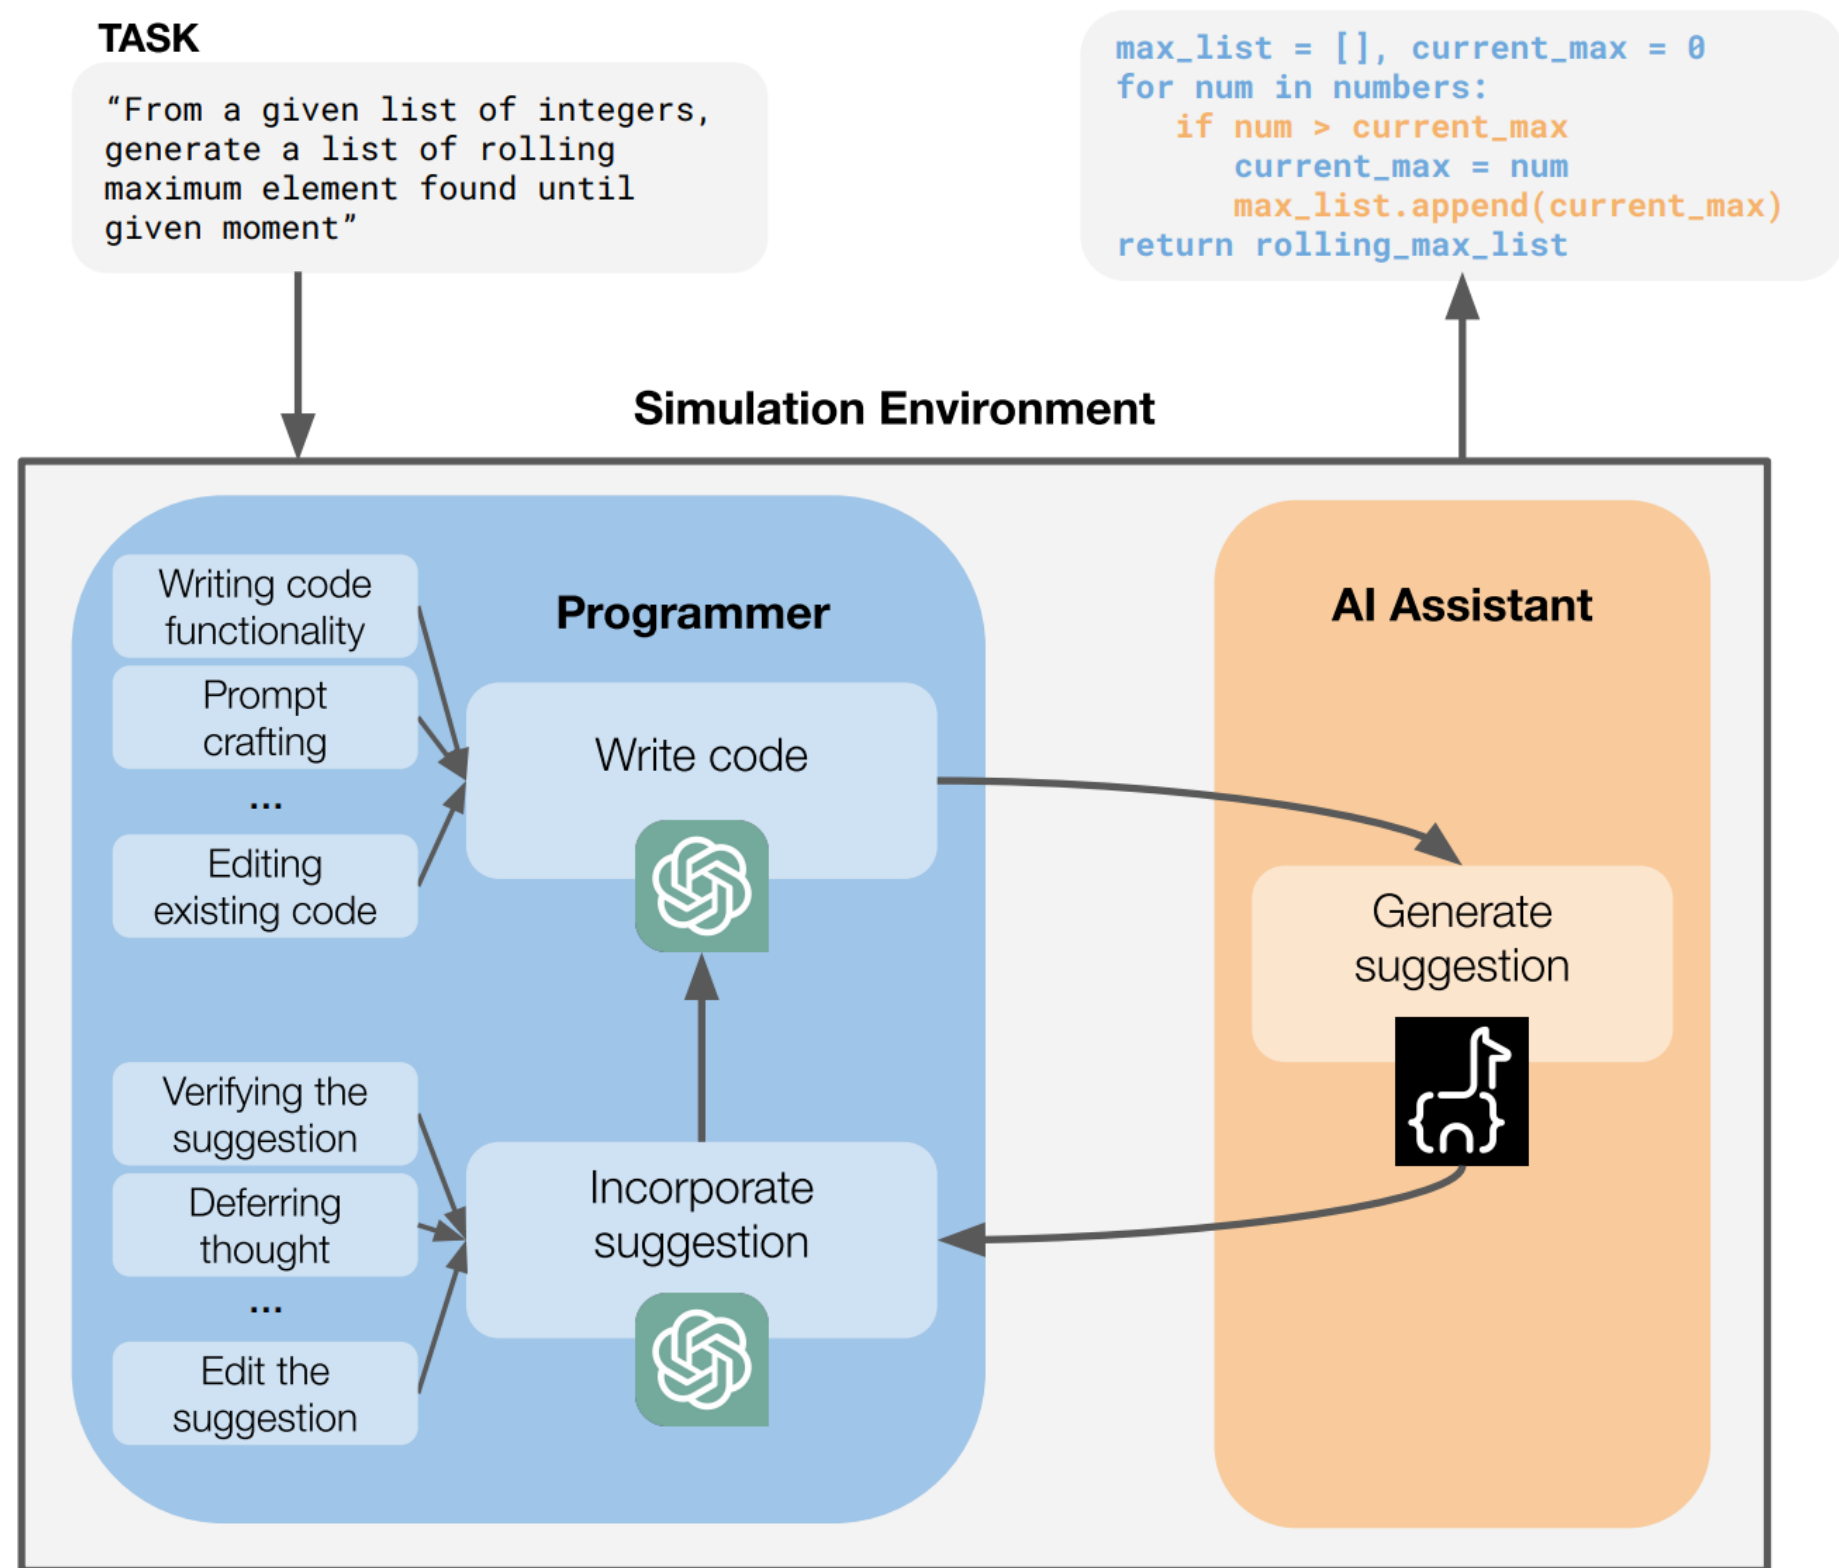
"We propose a utility theory framework, which models [when AI should make intervention to] programmers and decides which suggestions to display."

...models that predict suggestion acceptance to selectively hide suggestions reducing both latency and programmer verification time.



Mozannar, Hussein, et al. "When to show a suggestion? integrating human feedback in ai-assisted programming." *AAAI 2024*

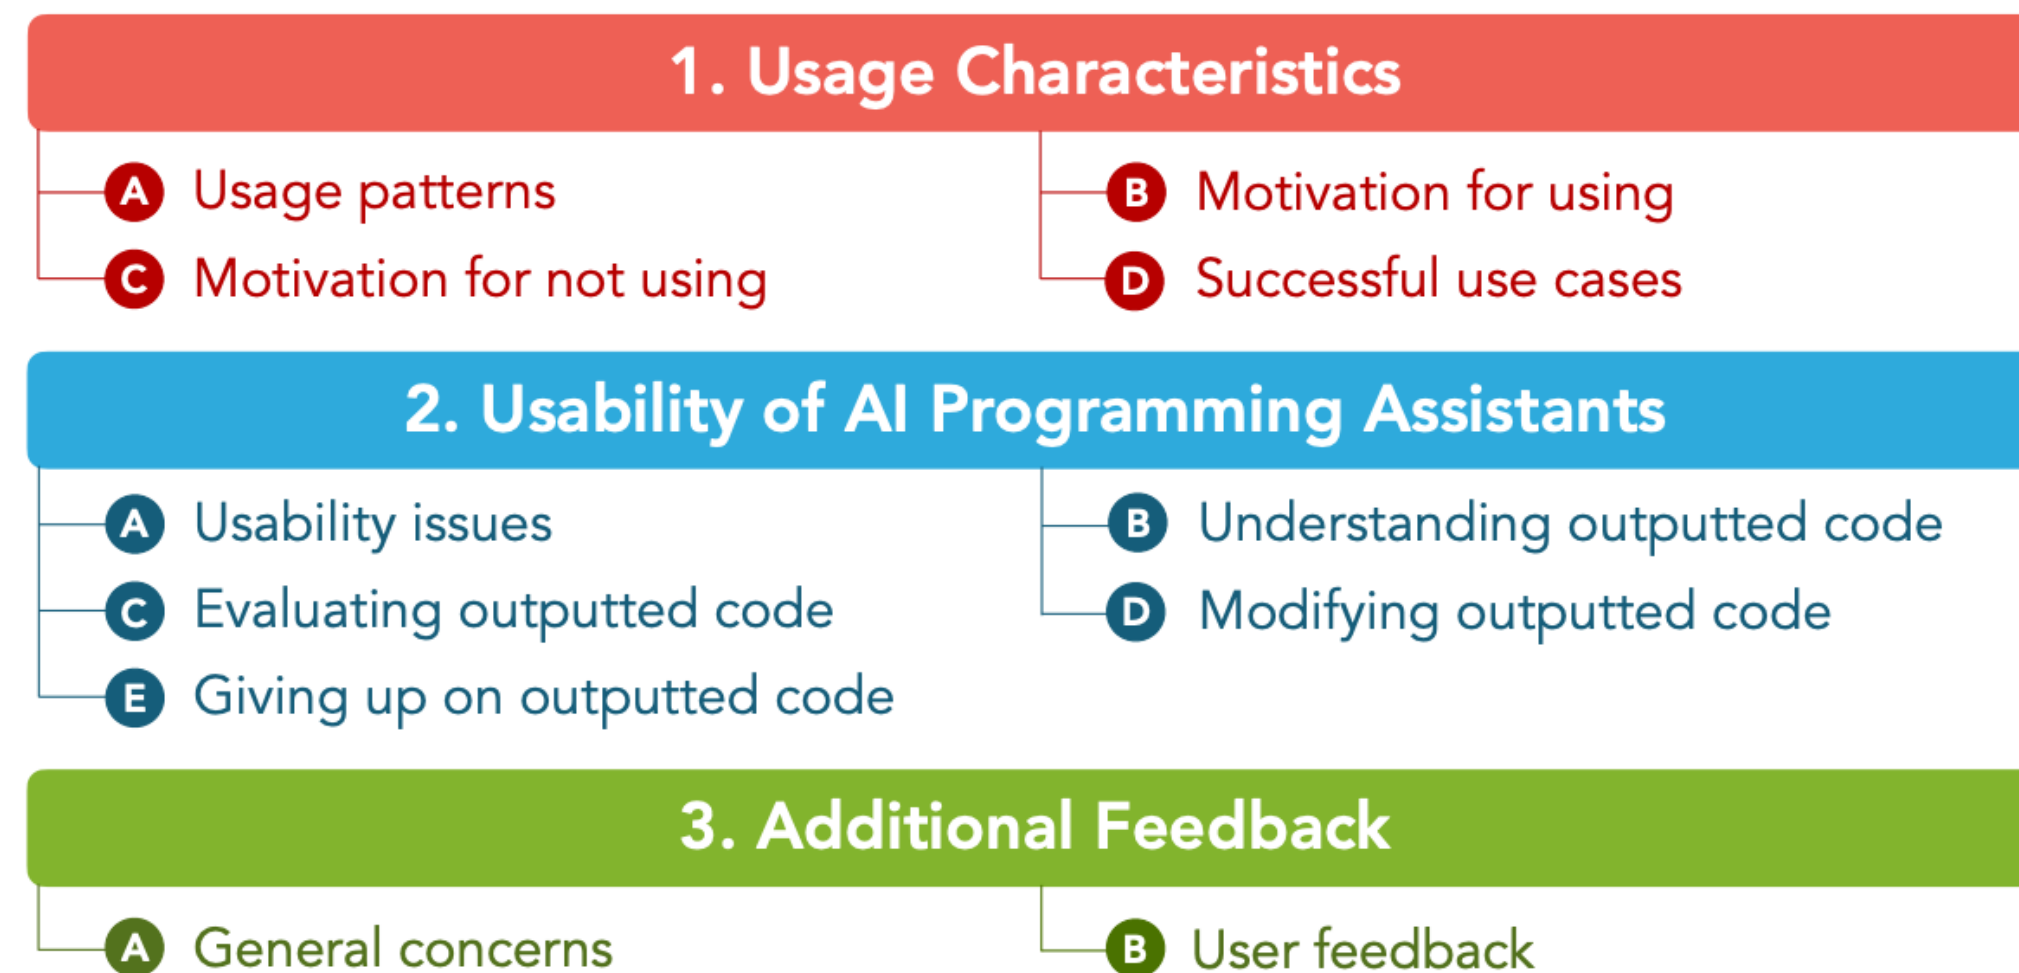
Once have some data, can simulate humans



Mozannar, Hussein, et al. "Simulating Iterative Human-AI Interaction in Programming with LLMs." *NeurIPS 2023 Workshop*

Understand humans qualitatively: Surveys

“To understand developers’ practices while using these tools and the important usability challenges they face, we administered a survey to a large population of developers and received responses from a diverse set of 410 developers.”








SURVEY QUESTIONS

- For this software project, estimate what percent of your code is written with the help of the following code generation tools.
- For each of the following reasons why you use code generation tools in this software project, rank its importance.
- For each of the following reasons why you do not use code generation tools, rank its importance.
- For your software project, estimate how often you experience the following scenarios when using code generation tools.
- For your software project, estimate how often the following reasons are why you find yourself giving up on code created by code generation tools.
- ★ What types of feedback would you like to give to code generation tools to make its suggestions better? Why?

Liang, Jenny T., Chenyang Yang, and Brad A. Myers. "A large-scale survey on the usability of ai programming assistants: Successes and challenges." *ICSE 2024*

Understand humans qualitatively: Surveys

Motivation	Distribution
A. For using	
M1 To have an autocomplete or reduce the amount of keystrokes I make.	86%  6.2%
M2 To finish my programming tasks faster.	76%  12%
M3 To skip needing to go online to find specific code snippets, programming syntax, or API calls I'm aware of, but can't remember.	68%  14%
M4 To discover potential ways or starting points to write a solution to a problem I'm facing.	50%  24%
M5 To find an edge case for my code I haven't considered.	36%  44%

Repetitive code (boilerplate code, repetitive endpoints for crud, etc.)

Code with simple logic

Autocomplete ("acceleration")

Quality assurance (e.g. log messages, test cases)

Proof-of-concepts (generate multiple implementations for a given problem)

Learning (of new libraries or programming languages)

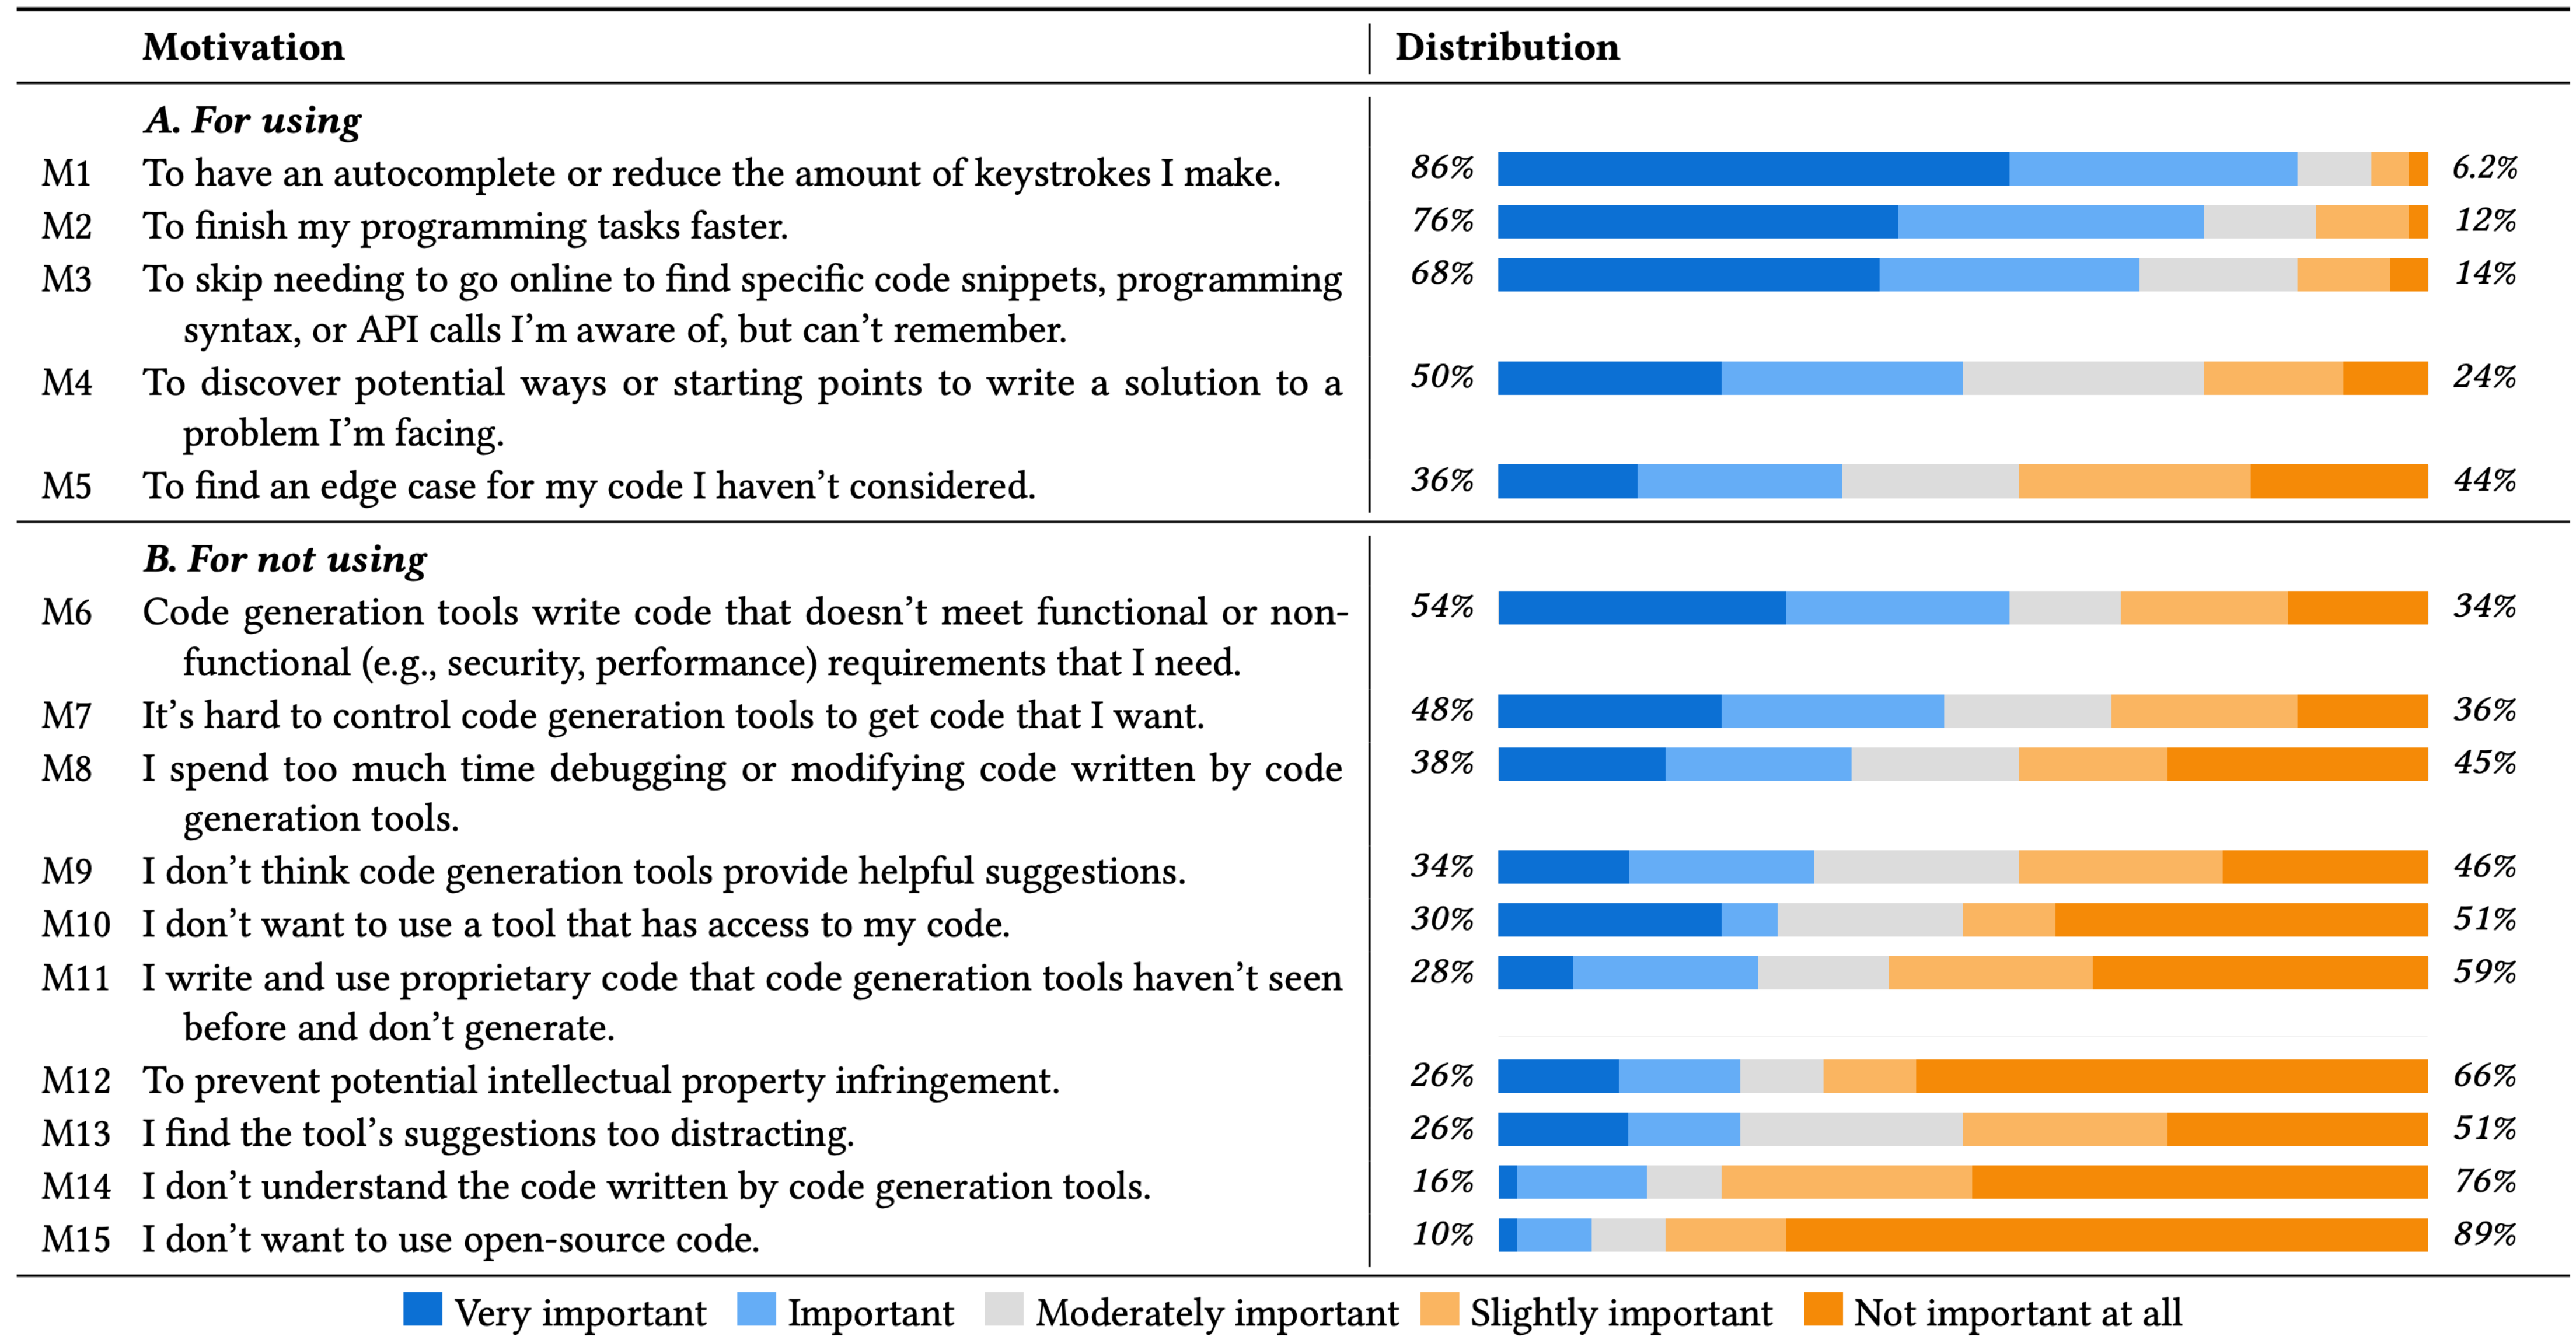
Recalling (Find syntax they were familiar with but couldn't recall)

Efficiency

Documentation

Code consistency (e.g., indentation, quickly referencing sources created within the project)

Why humans use or not use programming tools



Bonus: Taxonomy of software requirements

Functional requirements characterize units of functionality that we may want to group into coarser-grained functionalities that the software should support.

Non-functional requirements

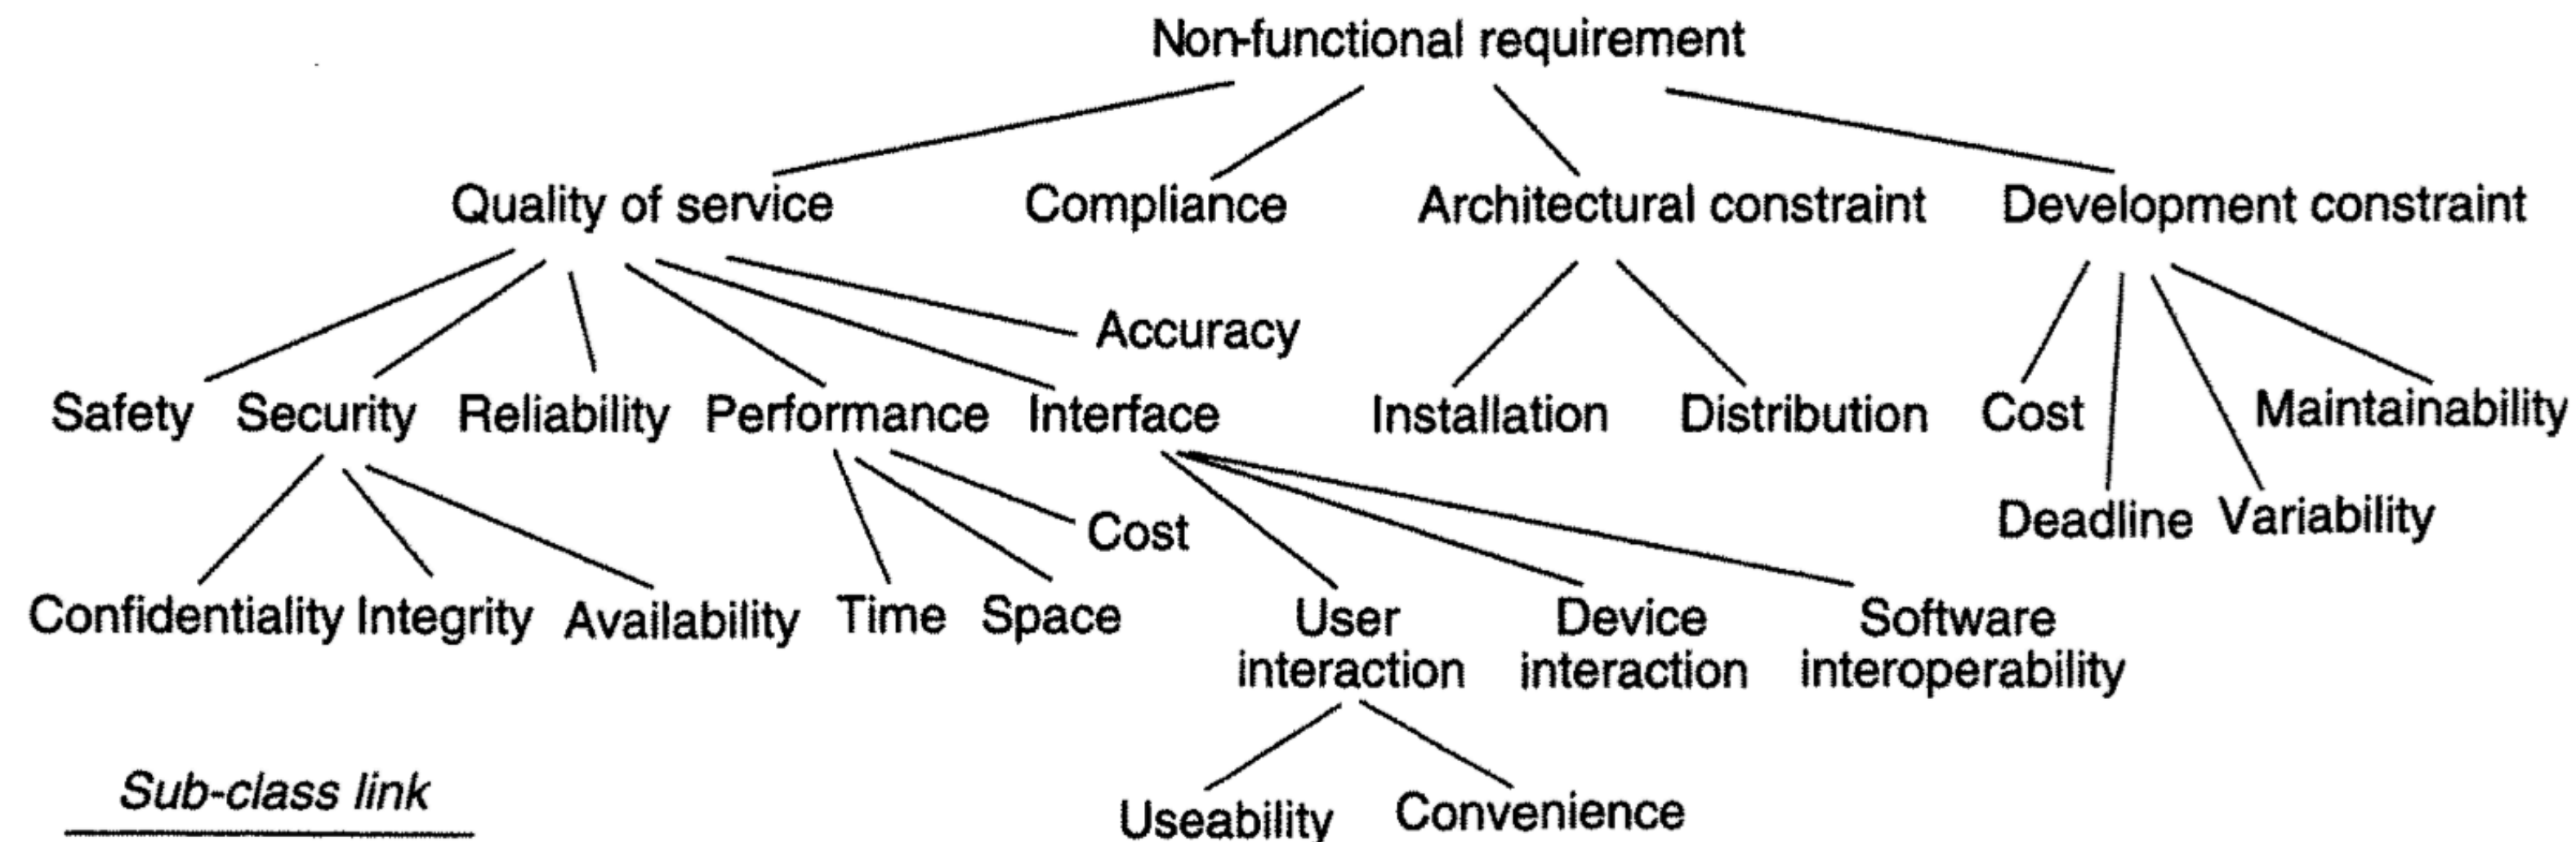


Figure 1.5 A taxonomy of non-functional requirements

More here: <https://www.cs.cmu.edu/~ckaestne/17313/2018/20180913-requirements-solicitation-and-doc.pdf>

Some other findings you might find interesting...

Engineers' prompting / input strategies (ranked by frequency)

Clear explanations, through "doctoring comments" and test cases

No strategy

Adding code

Follow conventions (e.g., well-named variables)

Break down instructions

Existing code context ("use it at advanced stages of project, for it to give better suggestions based on my project's history")

Prompt engineering

User-envisioned additional functionalities

Better understanding of code context. e.g., Code from the same workspace; Don't use deprecated API

Tool configuration: Have configurable parameters for suggestion frequency, distinguish[ing when to do] long code generation and short code [generation]

Natural language interactions

Code analysis, add annotation for functional and syntactic correctness

Explanations, e.g., link directly to documentation

More suggestions

Account for non-functional requirements

How do you know a code gen model is **useful**?

More metrics from more traditional human-human studies

Analyze actual human interactions

Hear what they have to say

How do you **make** a code gen model useful?

Again first some human-human reference...

Moderators	Human-Human vs. Human Solo	Human-AI (Copilot)
Task Types & Complexity	Complex task improve quality, simple one does not [7]; debugging is perceived as less enjoyable or effective than comprehension or refactoring [22]	N/A
Compatibility (E.g., Expertise)	Random pairing led to incompatible partners and conflicts during work [18]. Expertise: improve quality more effectively if pair is similarly skilled [14]; less-skilled students learn more and enjoy more [20, 22]; if knowledge gap is large, less-skilled programmers may tend to be more passive and disengaged [23]	N/A
Communication	Conversations with intermediate-level details contribute to pair programming success [24]; different types of discourse lead to more attempts or more debug success [25]	N/A
Collaboration	Over-reliance leads to conflicts and impedes satisfaction and learning, as work is entirely burdened on one partner [4, 18]; educators recommend regular role-switching to ensure equitable learning in collaboration [2]	N/A
Logistics	Scheduling difficulties [26], teaching & evaluating individual responsibility and accountability are important to collaboration success [27], but can lead to increased management costs [21, 28]	N/A

Qianou Ma, Tongshuang Wu, and Kenneth Koedinger. "Is AI the better programming partner? Human-Human Pair Programming vs. Human-AI pAIr Programming." *AIED 2023 workshop*

Human-Human challenges to opportunities

Moderating Factors	Human-Human Challenges	Human-AI Opportunities
Task Types & Complexity: pair work better if the task is not too simple and good for collaboration [7, 22]	Hard to design suitable tasks of appropriate complexity level	AI may be used to generate collaboration tasks and adjust tasks complexity
Compatibility: pairs with similar skill levels and compatible working styles work better [14, 22]	Hard to find a similarly skilled or compatible partner	AI partner should adjust to human skill level and adapt to be compatible with different people
Communication: pairs work better with productive conversations [24], and critiques lead to more debugging success [25]	Hard to teach effective communication and constructive criticism	AI partner should support productive conversations and provide critiques
Collaboration: pairs work better with positive interdependence [27] and clear and balanced responsibilities [18]	Hard to teach collaboration and prevent free riders	AI should support positive social interactions and collaboration and avoid over-assist that eliminates human's need to engage
Logistics: pair programming is costly to implement because of management challenges [21, 28]	Hard to schedule and assess individual contributions in a pair	Scheduling is no longer a problem, but humans should be accountable and responsible when using AI-generated code

How do you make a code gen model useful?

A: Design and iterate on the interface

Communication via UI: Inline explanations

```
def create_window():  
    window = tk.Tk()  
    window.title("COVID-19 Visualization")  
    window.geometry("1000x600")  
    window.resizable(False, False)  
    canvas = tk.Canvas(window, background='white',  
                       width=800, height=600,  
                       scrollregion=(0, 0, 800, 600), orient='horizontal',  
                       cursor='x')  
    scrollbar = tk.Scrollbar(window, command=canvas.yview, width=10)  
    scrollbar.pack(side='right', fill='x')  
    canvas.pack(side='top', fill='both')  
    scrollbar.pack(side='right', fill='y')
```

1

2

3

The method controls whether the main application window can be resized by the user.

The first and second arguments determine if the window's width and height are resizable. Now, the window is non-resizable in both directions.

1. Create a window with the title "COVID-19 Visualization" and size 1000x600.
2. Create a canvas with background color white and size 800x600.
3. Create a scrollbar for the canvas.

Figure 1: Ivie augments the interactive programming assistant with instant explanations that help programmers examine generated code. When a programming assistant suggests code (*italic text above*, ①), Ivie annotates it with brief, informative explanations. Explanations appear at the level of blocks of code (in the right margin, ②) and expressions (anchored beneath the line the programmer hovers over, ③). For single-line suggestions, expression explanations appear automatically. Ivie's explanations help programmers break up complex or unfamiliar suggestions into pieces that can be more readily understood.

Communication via UI: Inline explanations

```
df_all.merge(df_Apr, on='City', how='left', suffixes=('_all', '_apr'))
```

Combines df_all with another DataFrame df_Apr.

Only rows with matching 'City' values in both DataFrames will be merged.

Retains all rows from df_all, with matching rows from df_Apr. If no match is found,

In case of column name conflicts, df_all columns end with _all and df_Apr with

Figure 4: Explaining expressions. After a program consists of a single line, IVIE reveals explanations of calls. The purpose of these explanations is to make programmer seeking to understand the precise behavior

```
def visualize_data(df, max_temp_city, max_rain_city):  
    fig, ax = plt.subplots(2, 1, figsize=(14, 10))  
  
    df[df['City'] == max_temp_city]['Temperature'].plot(ax=ax[0])  
    ax[0].set_title(f'Yearly Average Temperature for {max_temp_city}')  
    ax[0].set_xlabel('Year')  
    ax[0].set_ylabel('Temperature (°C)')  
    ax[0].yaxis.set_major_formatter(ticker.FormatStrFormatter('%0.1f'))  
  
    df[df['City'] == max_rain_city]['Rainfall'].plot(ax=ax[1], color='green')  
    ax[1].set_title(f'Yearly Average Rainfall for {max_rain_city}')  
    ax[1].set_xlabel('Year')  
    ax[1].set_ylabel('Rainfall (mm)')  
    ax[1].yaxis.set_major_locator(ticker.MaxNLocator(nbins=6, integer=True))  
  
    plt.tight_layout()  
    plt.show()
```

1. Creates a figure with two vertical subplots.

2. Plot the temperature data for the city with the highest temperature in the top subplot.

3. Plot the rainfall data for the city with the highest rainfall in the bottom subplot.

4. Format the figure and display it.

Figure 5: Explaining multi-line suggestions. When a programming assistant suggests multiple lines of code, IVIE splits up and explains that code. Its explanations appear in the right margin of the editor. The explanations are meant to help a programmer get a high-level understanding of the code. In the pictured scenario, these explanations might help the programmer understand that the two longest sections of the code suggestion configure each of two subplots, each with a different slice of the data.

Communication via UI: Inline explanations

Usually no need to be super sophisticated methods but just clean communication!

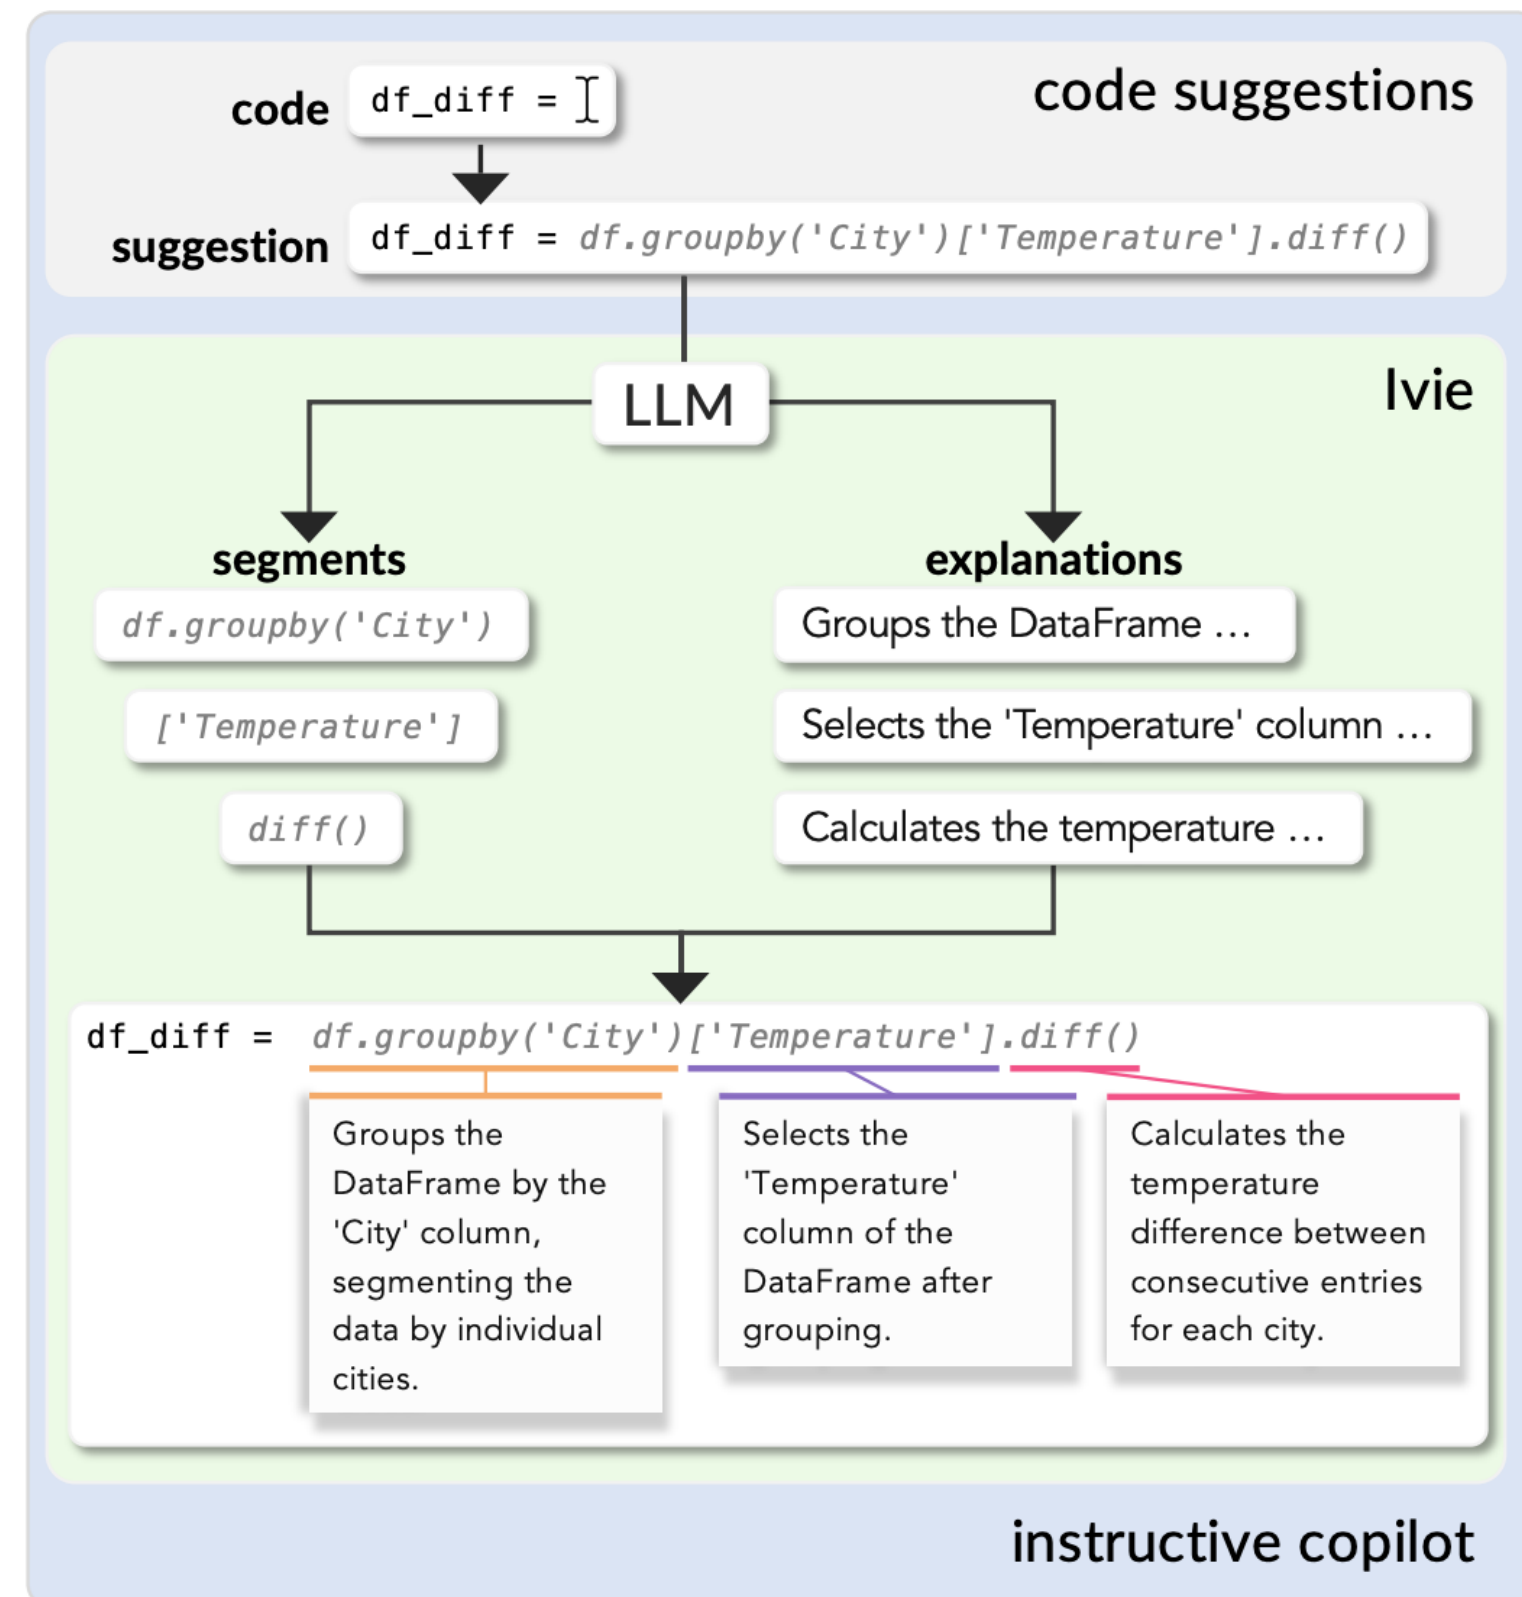


Figure 2: The implementation of an instructive copilot for programming. IVIE creates interactive overlays that explain suggestions made by a programming assistant. When the programming assistant (e.g., Copilot) displays the suggestion, IVIE submits that suggestion in a prompt to an LLM, requesting that the code be segmented and explained. IVIE then integrates the explanations into the editor as overlays beneath the expressions they explain.

Please dissect the following line of code, and explain the unfamiliar vocabulary and structures with less than 15 words each. Include ranges for parameter values and describe how changes in these parameters will affect the output.

Prompt:

```
fig, ax = plt.subplots(2, 1, figsize=(14, 10))
```

Output:

```
plt.subplots $$$ Create a figure and set of subplots.  
2, 1 $$$ 2 rows, 1 column of subplots.  
figsize=(14, 10) $$$ Width and height of entire figure.  
fig $$$ The whole window/figure containing subplots.  
ax $$$ Array of individual subplot axes.
```

Prompt:

```
df_all.merge(df_Apr, on='City', how='left',  
suffixes=('_all', '_apr'))
```

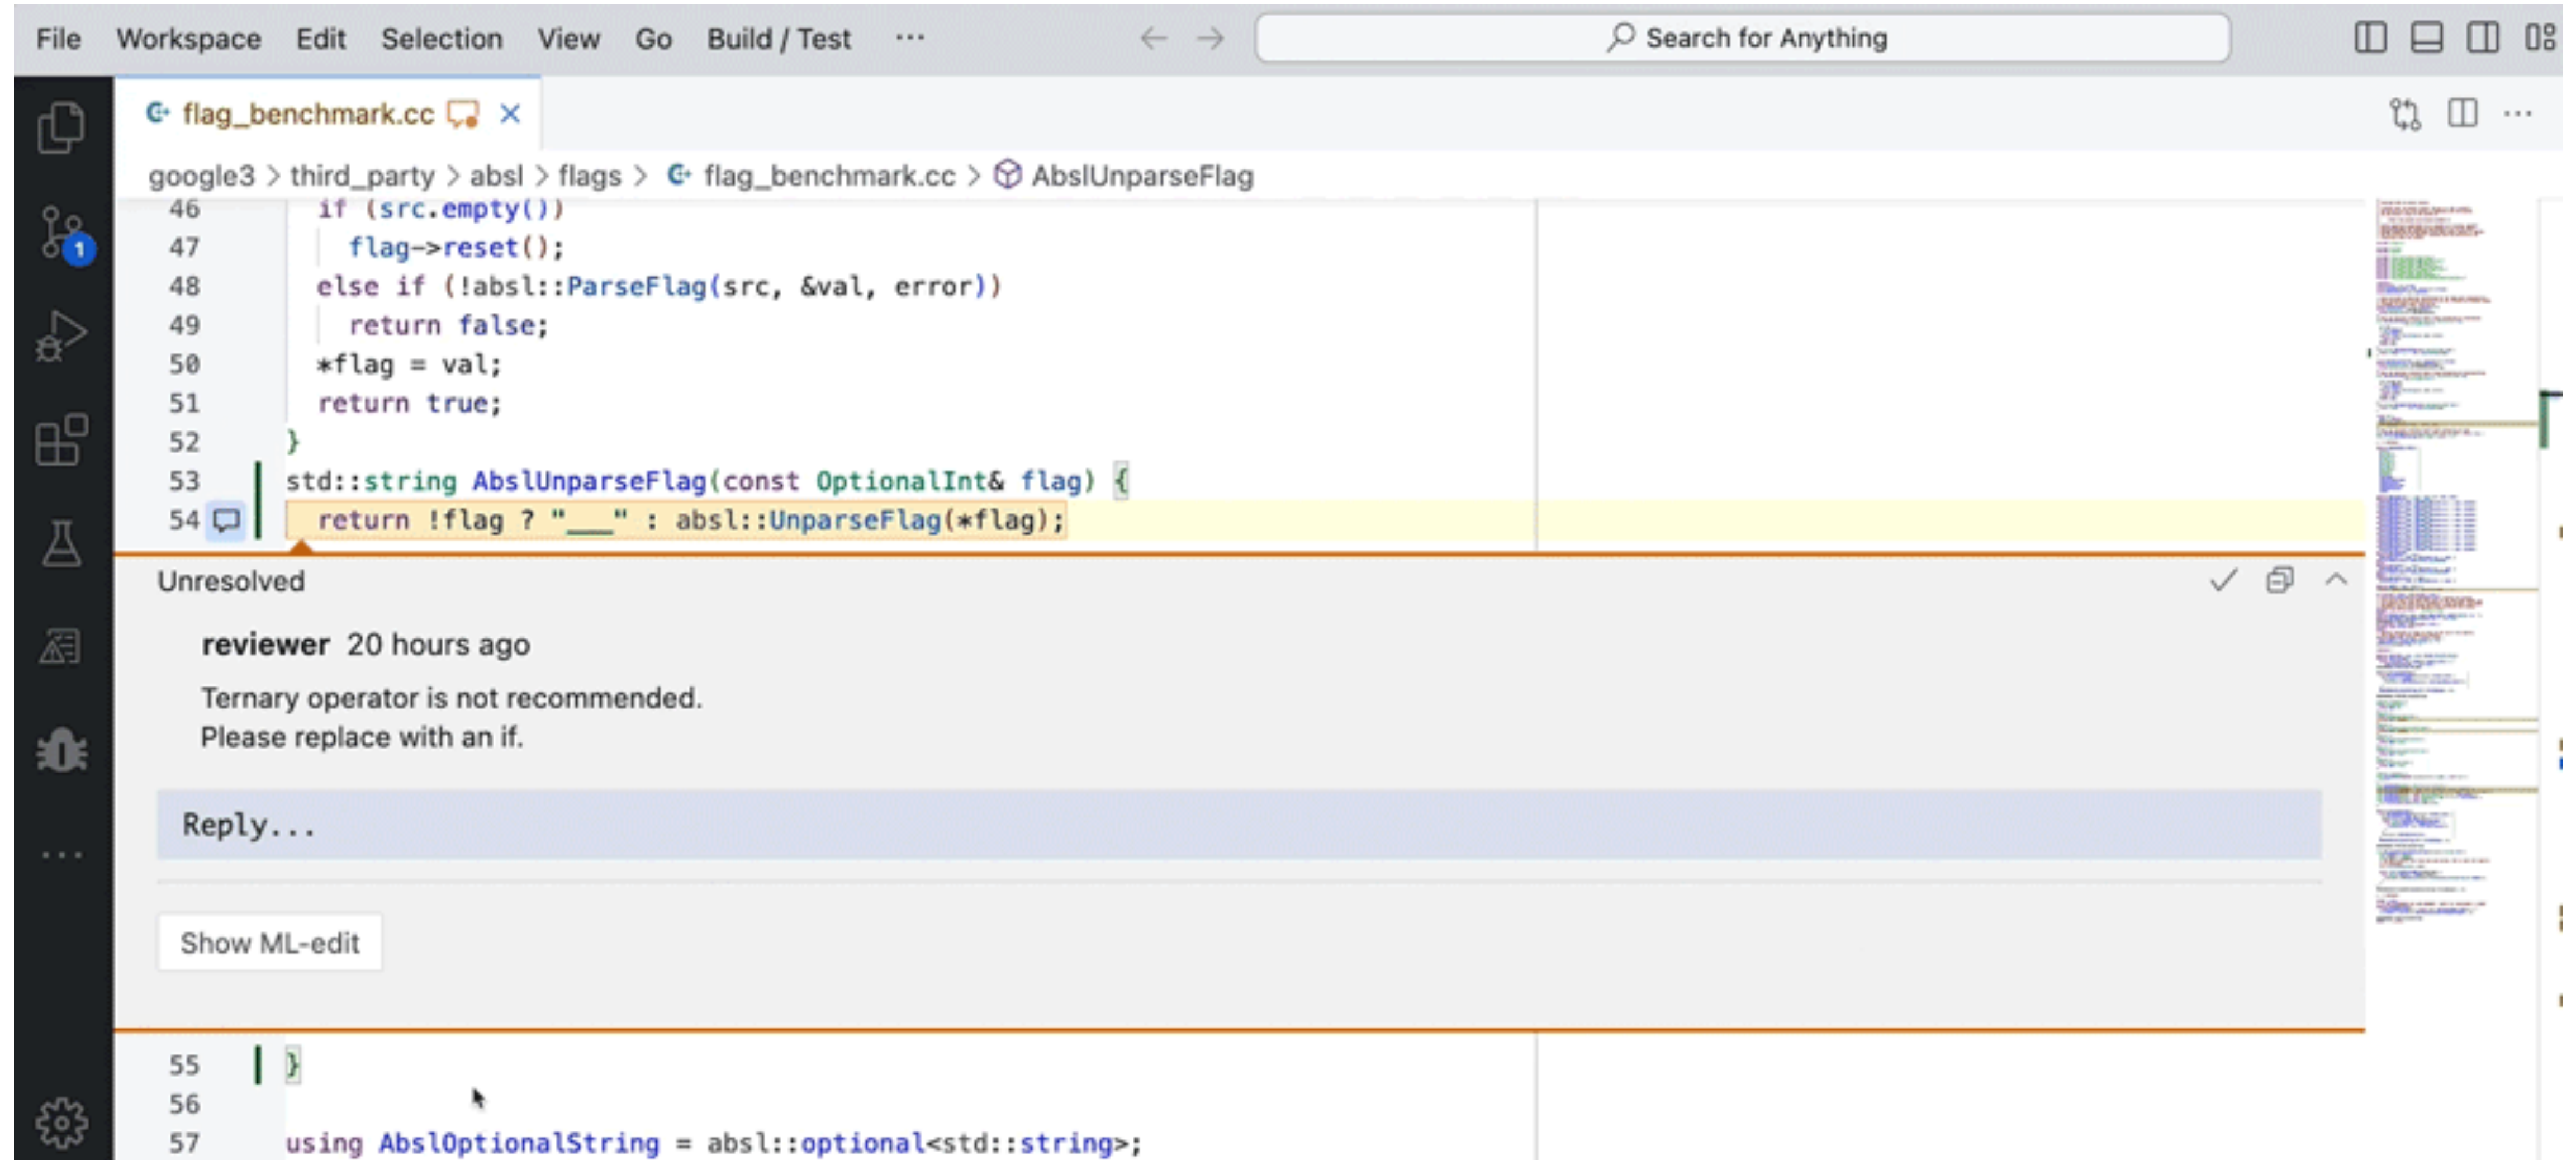
Output:

Figure 3: A prompt for requesting expression-level explanations of generated code. This prompt requests explanations of suggested code. It provides a single example of how it would like code suggestions to be split into expressions with accompanying brief explanations of those expressions.

How do you make a code gen model useful?

A: Design and iterate for specific use cases (case studies!)

Design for Use Case: Addressing Code Review



Design for Use Case: Addressing Code Review

An example of a review comment in Critique. The reviewer asked for a defensive coding practice. The author addressed the comment by updating their changelist with a new review snapshot. The update is shown via colors: green for added text and red for removed. The author responded to the comment with "Done." and marked it "Resolved".

The screenshot displays a code review interface for a change list (CL) 12345 by 'author', which is currently 'Pending'. The file path is `/v8/v11_0/src/objects/contexts.cc`. The interface shows a diff view with line numbers 206 to 217. A comment from a reviewer is shown on line 211, asking to check if `is_sloppy_function_name` is a null pointer before assigning it. The author responded with 'Done.' and the comment is marked as 'Resolved'. The code diff shows that the reviewer's comment was addressed by adding a null check (lines 211-213) and removing the original assignment (line 210).

```
206 206 bool has_seen_debug_evaluate_context = false;
207 207 *index = kNotFound;
208 208 *attributes = ABSENT;
209 209 *init_flag = kCreatedInitialized;
210 210 *variable_mode = VariableMode::kVar;
211 211 *is_sloppy_function_name = false;
212 212 if (is_sloppy_function_name != nullptr) {
213 213     *is_sloppy_function_name = false;
214 214 }
215 215 if (v8_flags.trace_contexts) {
216 216     Printf("Context::Lookup(");
217 217     name->ShortPrint();
```

Design for Use Case: Addressing Code Review

“We started by training a model that predicts code edits needed to address reviewer comments. The model is pre-trained on various coding tasks and related developer activities (e.g., renaming a variable, repairing a broken build, editing a file). It’s then fine-tuned for this specific task with reviewed code changes, the reviewer comments, and the edits the author performed to address those comments.”

```
... +8 common lines ...
import datetime

... +67 common lines ...

def _temporal_split_from_datetime(
    example_datetime: datetime.datetime) -> Optional[str]:
    """Finds the split name using the example datetime."""
    end_train_datetime = datetime.datetime.strptime('2022-01-01', '%Y-%m-%d')
    end_valid_datetime = datetime.datetime.strptime('2022-05-01', '%Y-%m-%d')
    end_test_datetime = datetime.datetime.strptime('2022-06-01', '%Y-%m-%d')

    if example_datetime < end_train_datetime:
        return 'train'
    el if example_datetime < end_valid_datetime:
        return 'validation'
    el if example_datetime < end_test_datetime:
        return 'test'
    else:
        return None
```

Model quality: In-product measurements

Offline evaluation, by computing the recall@X metric described above over a held-out test dataset

Online evaluation / user feedback, by measuring the number of code-review comments produced during day-to-day business, the number of predictions the model made, the number of those predictions that were previewed, and of those how many were applied, or received thumbs up/thumbs down.

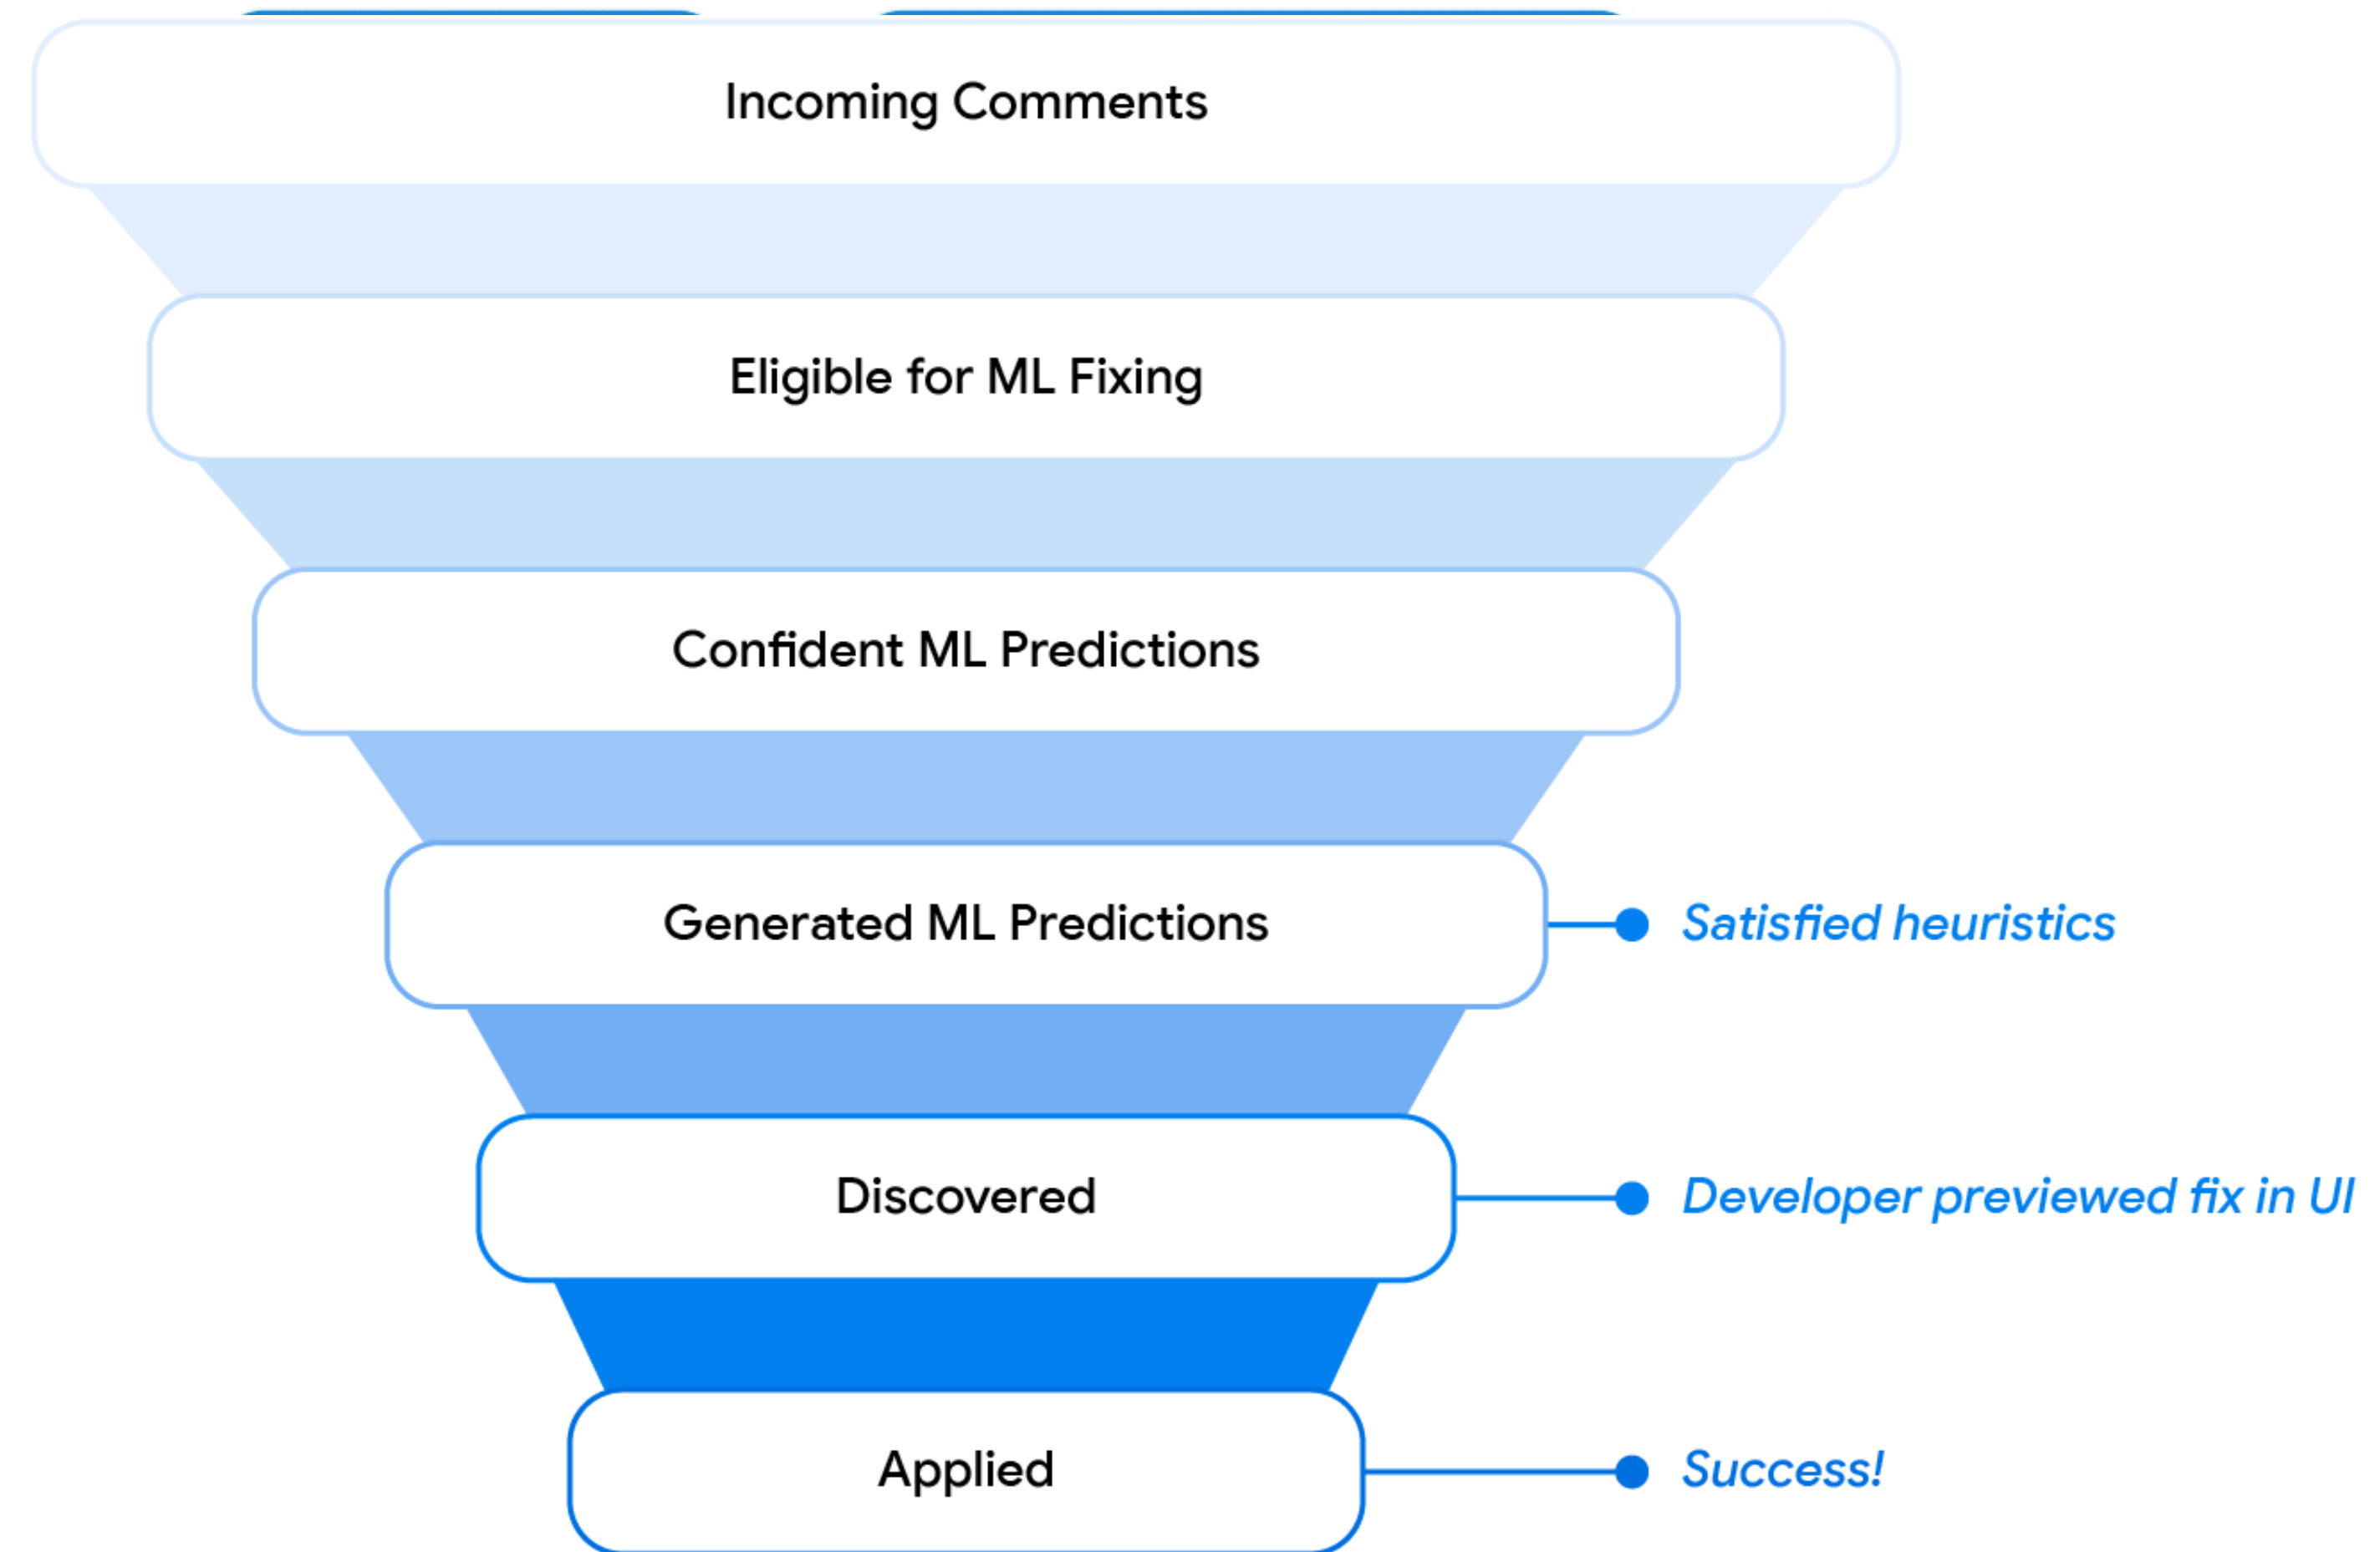
All types of such evaluation are meant to detect an increase in developer productivity, but act as easier-to-measure proxies of that measure.

Acceptance rate: the fraction of comments resolved by an accepted ML suggestion

Discoverability: the fraction of surfaced suggestions that were previewed by system users.

Filters between model and usage

“For every new reviewer comment, we generate the model input in the same format that is used for training, query the model, and generate the suggested code edit. **If the model is confident in the prediction and a few additional heuristics are satisfied**, we send the suggested edit to downstream systems.”



The impact of UIs

```
127 std::vector<AnnotatedCodeRange> ranges_;
```

Auto-generated suggested edit ^

15:25 (0,220) Suggested edit addressing comment:
'Up to you, but I think we might benefit from ...'

Actionable Show fix Was this helpful?

reviewer ^

15:25 Up to you, but I think we might benefit from a brief description of this field.

Unresolved Reply Reply with quote Done Ack

```
211 *is_sloppy_function_name = false;
```

reviewer ^

4:16 PM check if is_sloppy_function_name is nullptr before assigning it

Unresolved Show ML-edit Reply Done Ack

Original: A separate, asynchronous analyzer queried the model and produced the suggested edit as an independent code finding, in a *separate annotation*.

Pitfall: decoupled comment and suggested edit.

-> Duplication of information, wasted precious UI real-estate, and confused the prevailing visual language of review comments.

Revision: combine the two sources of information, by placing a "Show ML edit" in the same box where the reviewer comment appears

Result: improved discoverability considerably

The impact of UIs

```
211 *is_sloppy_function_name = false;
```

reviewer check if is_sloppy_function_name is nullptr before assigning it 4:16 PM

Unresolved Show ML-edit Reply Done Ack

```
211 *is_sloppy_function_name = false;
```

reviewer check if is_sloppy_function_name is nullptr before assigning it 10:32 AM

ML-suggested edit Reviewed

		+208 common lines	+10	+Block
--	--	-------------------	-----	--------

```
209 209 *init_flag = kCreatedInitialized;
210 210 *variable_mode = VariableMode::kVar;
211 211 *is_sloppy_function_name = false;
211 211 if (is_sloppy_function_name != nullptr) {
212 212     *is_sloppy_function_name = false;
213 213 }
212 214
213 215 if (v8_flags.trace_contexts) {
```

		+429 common lines	+10	+Block
--	--	-------------------	-----	--------

Unresolved Show edit Reply Done Ack

Pitfall: click to view. Since code shepherding (i.e., editing the changelist in light of the reviewer comments) takes a significant fraction of developers' time—one study at Google measured the median to be around 60 minutes [7]—efficiency in addressing comments is important

Revision: Showing the suggested edit immediately next to the reviewer comment, rather than requiring a click of the "Show ML edit" button.

Result: ML-suggested edit discoverability for the changelist author improved.

The impact of UIs

The screenshot shows a code review interface. At the top, there's a header with a star icon, 'CL', a dropdown menu, 'by', another dropdown menu, 'Unresolved', 'Request Review', and 'Reply' buttons. Below this is a breadcrumb trail: 'v8/v11_0/src/objects/contexts.cc'. There are also icons for search and 'Snaps'. The main content area shows a code diff. Line 212 has a comment: '*is_sloppy_function_name = false;'. Below the comment is a text box containing the text 'Check if this is null bef'. A red arrow points from the text box to the code line. Below the text box are three checkboxes: 'No action required' (unchecked), 'Attach ML-suggested edit (1) ?' (checked), and 'Markdown' (checked). Below these are three buttons: '+209 common lines', '+10', and '+Block'. The code diff shows lines 210-216. Line 212 is highlighted in pink. Lines 212-214 are highlighted in green. Below the code diff are three buttons: '+429 common lines', '+10', and '+Block'. At the bottom, there are three buttons: 'Unresolved', 'Save draft', and 'Cancel'.

The location of the comment and the mention of "check" and "null" were sufficient to trigger the assistant to suggest the intended edit.

Original: Just show suggested revision to the code author but not the reviewer.

Pitfall: decoupled reviewer from ML assistant.

Reviewers who were uncomfortable having an ML model "interpret" their comment into a suggested edit, and would prefer to preview the suggestion before providing it to the code author. "the pedagogical function of code review – It is often how new engineers are trained on local conventions and programming discipline."

Revision: The reviewer is shown the ML-suggested edit as they type their comment. The reviewer can decide to reject the suggested edit (in which case the author will only see the reviewer's comment).

Result: Many incorrect suggestions are pre-filtered out; Can use a less lower auto-filter because human filter is in the loop!

The impact of UIs

Original: Reviewers are typically pressed for time, and may move on quickly from comment to comment. In an attempt to reduce back-end prediction load, and to avoid showing reviewers suggested edits before they have typed enough of their comment, we set the triggering delay between when the reviewer starts typing a comment and when a prediction is requested to 1500ms.

Pitfall: slow-to-predict edits.; Between this triggering delay, and the additive prediction latency of the model, many predictions “missed” the reviewer, who had already moved on.

Revision: Further reduced the triggering delay to 500ms. and improved the prediction latency through considerable engineering effort.

Result: number of suggested edits previewed by reviewers increased by 12%, and the acceptance rate of ML-suggested edits by authors improved by 18%.

The screenshot shows a code review interface. At the top, there's a header with a star icon, a dropdown menu showing 'CL [redacted]', 'by [redacted]', and 'Unresolved' status. There are buttons for 'Request Review' and 'Reply'. Below the header, there's a breadcrumb trail: 'v8/v11_0/src/objects/contexts.cc'. A 'Snaps' icon is visible on the right. The main content area shows a code diff. Line 212 is highlighted in pink, and line 213 is highlighted in green. A comment box is open over line 212, containing the text 'Check if this is null bef' with a red arrow pointing to the end of the text. Below the comment box, there are checkboxes for 'No action required', 'Attach ML-suggested edit (1)', and 'Markdown'. There are also buttons for '+209 common lines', '+10', and '+Block'. The code diff shows the following changes:

```
210 210 *variable_mode = VariableMode::kVar;
211 211
212 212 *is_sloppy_function_name = false;
212 212 if (is_sloppy_function_name) {
213 213     *is_sloppy_function_name = false;
214 214 }
213 215
214 216 if (v8_flags.trace_contexts) {
```

At the bottom, there are buttons for 'Unresolved', 'Save draft', and 'Cancel'.

The impact of UIs

Original: Our original design of the UI assumed that the changelist author and reviewers operate in lock step: one stops when the other starts working on the changelist.

Pitfall: code review is serialized.; This is not how code review operates in practice. Sometimes the changelist is edited by the author as the reviewer is reviewing, or perhaps the reviewer thinks of a new comment after they have passed the bulk of their review to the author, and sometimes the reviewer attaches a comment to an older review snapshot of the changelist. All in all, this means that sometimes even an MLsuggested edit that the author wishes to accept is incompatible with the current state of the code.

Revision: detect those cases, and opening a three-way merge window (Figure 10) for the author to resolve any merge conflicts.

Result: the number of accepted suggested edits increased.

```
File Workspace Edit Selection View Go Build/Test ... Search for Anything
flag_benchmark.cc flag_benchmark.cc (ML-suggested edit)
Accept Merge Add ML-suggested Edits Add Workspace Edits Was this suggestion helpful?
ML-suggested Edit Merge Results Workspace
43 bool AbslParseFlag(absl::string_view src, 0 43 bool AbslParseFlag(absl::string_view src, 0 43 bool AbslParseFlag(absl::string_view src, 0
44 | | | | | std::string* error) { 44 | | | | | std::string* error) { 44 | | | | | std::string* error) {
45 int val; 45 int val; 45 int val;
46 if (src.empty()) 46 if (src.empty()) 46 if (src.empty())
47 | flag->reset(); 47 | flag->reset(); 47 | flag->reset();
48 else if (!absl::ParseFlag(src, &val, erro 48 else if (!absl::ParseFlag(src, &val, erro 48 else if (!absl::ParseFlag(src, &val, erro
49 | return false; 49 | return false; 49 | return false;
50 *flag = val; 50 *flag = val; 50 *flag = val;
51 return true; 51 return true; 51 return true;
52 } 52 } 52 }
53 std::string AbslUnparseFlag(const Optiona 53 std::string AbslUnparseFlag(const Optiona 53 std::string AbslUnparseFlag(const Optiona
54 | if (!flag) return ""; 54 | return !flag ? "" : absl::UnparseFlag(*fl 54 | return !flag ? "" : absl::UnparseFlag(
55 | return absl::UnparseFlag(*flag); 55 | 55 |
56 } 56 } 56 }
57 57 57
58 using AbslOptionalString = absl::optiona 58 using AbslOptionalString = absl::optiona 58 using AbslOptionalString = absl::optiona
59 struct OptionalString : AbslOptionalString 59 struct OptionalString : AbslOptionalString 59 struct OptionalString : AbslOptionalString
60 | using AbslOptionalString::AbslOptiona 60 | 60 |
61 }; 61 }; 61 };
62 // Next two functions represent Abseil Flag 62 // Next two functions represent Abseil Flag 62 // Next two functions represent Abseil Flag
63 bool AbslParseFlag(absl::string_view src, 0 63 bool AbslParseFlag(absl::string_view src, C 63 bool AbslParseFlag(absl::string_view src, 0
64 | | | | | std::string* error) { 64 | | | | | std::string* error) { 64 | | | | | std::string* error) {
65 | std::string val; 65 | std::string val; 65 | std::string val;
```

Some takeaways

When a model is in a specific use case it usually means **blending into existing workflows**

Test-in-product is not the most ideal but usually quite **useful**

There will be **metrics** not relevant to model, but just relevant to **usability** (e.g. discoverability)


Little things like latency in suggestion can easily change usability

UI iteration is a BIG aspect

Need to consider **all users touching** the system (reviewers, and authors).

Also consider the **original objective** of the task (a bit of education and training going on!)

Case Study: LLM for CS Education



When AI writes code

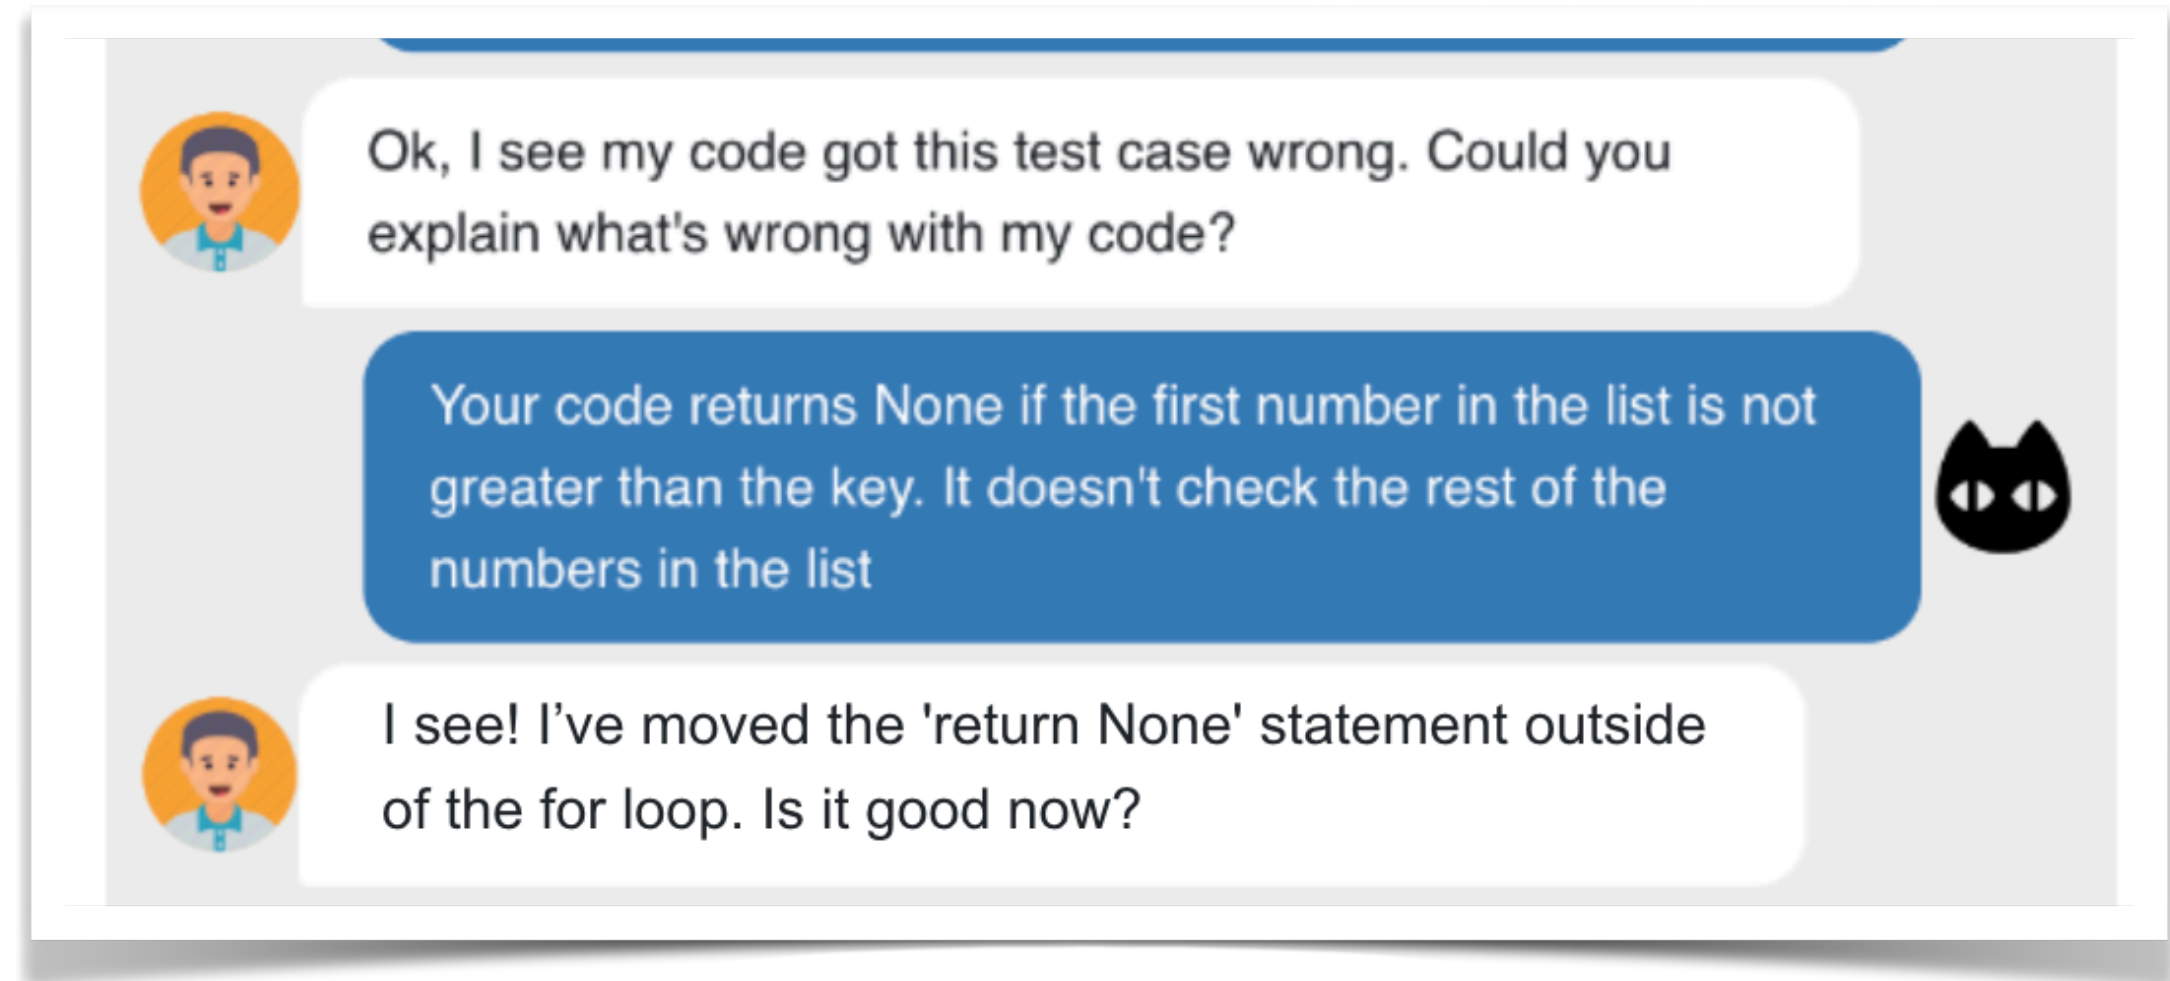
More than autocomplete

GitHub Copilot is powered by Codex, the new AI system created by OpenAI. GitHub Copilot understands significantly more context than most code assistants. So, whether it's in a docstring, comment, function name, or the code itself, GitHub Copilot uses the context you've provided and synthesizes code to match. Together with OpenAI, we're designing GitHub Copilot to get smarter at producing safe and effective code as developers use it.

```
1 package main
2
3 type Run struct {
4     Time int // in milliseconds
5     Results string
6     Failed bool
7 }
8
9 // Get average runtime of successful runs in seconds
10 func averageRuntimeInSeconds(runs []Run) float64 {
11     var totalTime int
12     var failedRuns int
13     for _, run := range runs {
14         if run.Failed {
15             failedRuns++
16         } else {
17             totalTime += run.Time
18         }
19     }
20
21     averageRuntime := float64(totalTime) / float64(len(runs) - failedRuns) / 1000
22     return averageRuntime
23 }
```



Humans might do more debugging!



Qianou Ma, et al. "How to Teach Programming in the AI Era? Using LLMs as a Teachable Agent for Debugging." *ICSE 2024*

Problem: first_num_greater_than

Write a Python function `first_num_greater_than(numbers_list, key)` that takes a list of integers (`numbers_list`) and an integer key (`key`), and returns the first number in the list that is greater than the key. If there is no number greater than the key, then you should return `None`.

Now you are chatting with a student. Please explain to them why their code is wrong by selecting the right explanation from the list. If you are right, the student will fix their code accordingly! Otherwise, they may get frustrated and leave.

Student's Current Code

```

1 def first_num_greater_than(numbers_list, key):
2     for i in range(len(numbers_list)):
3         if numbers_list[i] > key:
4             return numbers_list[i]
5     return None
    
```

View Code Differences

<pre> for i in range(len(numbers_list)): if numbers_list[i] > key: return numbers_list[i] else: return None </pre>	<pre> for i in range(len(numbers_list)): if numbers_list[i] > key: return numbers_list[i] return None </pre>
---	---

Test Suite Development

Add Test Case:

Input: Expected output:

Add Test Group:

Enter test group name:

Your Test Cases Passed? Actual Output

Default Group		
<code>assert(first_num_greater_than([1, 2, 3], 2) == 3)</code>	✓	3 delete
<code>assert(first_num_greater_than([1, 2, 3], 1) == 2)</code>	✓	2 delete

1. No number in list greater than key Delete Group		
<code>assert(first_num_greater_than([3, 2, 1], 3) == None)</code>	✓	None delete

2. Key in middle of list Delete Group		
<code>assert(first_num_greater_than([3, 2, 1], 2) == 3)</code>	✓	3 delete

Office Hour!

AI: Your function returns None as soon as it encounters a number that is smaller than the key, it only continue to check the loop if num is equal to key. But there might be numbers greater than the key later after a number smaller than key.

Student: But if that's the case, shouldn't `first_num_greater_than([1, 2, 3], 1)` give me 2? Currently my code would output None.

AI: Consider this test case: given input [1, 2, 3], 2, it is supposed to outputs 3, however, the actual behavior is unexpected as the program outputs None

Student: Ok, I see my code got this test case wrong. Could you explain what's wrong with my code?

AI: Your code returns None if the first number in the list is not greater than the key. It doesn't check the rest of the numbers in the list

Student: Now I've updated the code with your fix. Is it good now?

AI: Is the student correct?

Your code returns None if the first number in the list is not...

All of your test cases passed.

Problem: first_num_greater_than

Write a Python function `first_num_greater_than(numbers_list, key)` that takes a list of integers (`numbers_list`) and an integer key (`key`), and returns the first number in the list that is greater than the key. If there is no number greater than the key, then you should return `None`.

Student's Current Code

```
1 def first_num_greater_than(numbers_list, key):
2     for i in range(len(numbers_list)):
3         if numbers_list[i] > key:
4             return numbers_list[i]
5     return None
```

View Code Differences

```
for i in range(len(numbers_list)):
    if numbers_list[i] > key:
        return numbers_list[i]
else:
    return None
```

Your Test Cases

Default Group

```
assert(first_num_greater_than
([1, 2, 3], 2) == 3)
assert(first_num_greater_than
([1, 2, 3], 1) == 2)
```

1. No number in list greater than key

```
assert(first_num_greater_than
([3, 2, 1], 3) == None)
```

2. Key in middle of list

```
assert(first_num_greater_than
([3, 2, 1], 2) == 3)
```

Office Hour Queue

There are several students waiting for your help, please click to start chatting with them.



Bob: Finished!



Chelsea: Start helping



Dave: Waiting...

Problem: first_num_greater_than

Write a Python function first_num_greater_than(numbers_list, key) that takes a list of integers (numbers_list) and an integer key (key), and returns the first number in the list that is greater than the key. If there is no number greater than the key, then you should return None.

Test Suite Development

Add Test Case:

Input Expected output

+

Add Test Group:

Enter test group name +

Evaluate Test Suite

Passed? Actual

greater_than ✓ 3 delete

greater_than ✓ 2 delete

Student's Current Code

```
1 def first_num_greater_than(numbers_list, key):
2     for i in range(len(numbers_list)):
3         if numbers_list[i] > key:
4             return numbers_list[i]
5     return None
```

View Code Differences

```
for i in range(len(numbers_list)):
    if numbers_list[i] > key:
        return numbers_list[i]
else:
    return None

for i in range(len(numbers_list)):
    if numbers_list[i] > key:
        return numbers_list[i]
return None
```


1. No number in list greater than key

assert(first_num_greater_than([3, 2, 1], 3) == None) ✓ None delete

2. Key in middle of list

assert(first_num_greater_than([3, 2, 1], 2) == 3) ✓ 3 delete

Office Hour!



Hi Christina! Here is my code and I think I have some problem with it. Can you walk me through that?

Evaluate student's code against your test suite on the left!

Send

Problem: first_num_greater_than

Write a Python function first_num_greater_than(numbers_list, key) that takes list of integers (numbers_list) and an integer key (key), and returns the first number in the list that is greater than the key. If there is no number greater than the key, then you should return None.

Now you are chatting with a student. Please explain to them why their code is wrong by selecting the right explanation from the list. If you are right, the student will fix their code accordingly! Otherwise, they may get frustrated and leave.

Student's Current Code

```
1 def first_num_greater_than(numbers_list, key):
2     for i in range(len(numbers_list)):
3         if numbers_list[i] > key:
4             return numbers_list[i]
5     return None
```

View Code Differences

```
for i in range(len(numbers_list)):
    if numbers_list[i] > key:
        return numbers_list[i]
else:
    return None

for i in range(len(numbers_list)):
    if numbers_list[i] > key:
        return numbers_list[i]
    return None
```

Test Suite Development

Add Test Case:

Input Expected output

+

Add Test Group:

Enter test group name +

Submit Test Suite & Start Helping Students

Your Test Cases Passed? Actual Output

Your Test Cases	Passed?	Actual Output
Default Group		
1. No number in list greater than key		Delete Group
2. Key in middle of list		Delete Group
3. All numbers in list greater than key		Delete Group

Hour!

Your function returns None as soon as it encounters a number that is smaller than the key, it only continue to check the loop if num is equal to key. But there might be numbers greater than the key later after a number smaller than key.

But if that's the case, shouldn't first_num_greater_than([1, 2, 3], 1) give me 2? Currently my code would output None.

Consider this test case: given input [1, 2, 3], 2, it is supposed to outputs 3, however, the actual behavior is unexpected as the program outputs None

Ok, I see my code got this test case wrong. Could you explain what's wrong with my code?

Your code returns None if the first number in the list is not greater than the key. It doesn't check the rest of the numbers in the list

Now I've updated the code with your fix. Is it good now?

Is the student correct?
Yes No I don't know

code returns None if the first number in the list is not...

your test cases passed.

Send

Problem: first_num_greater_than

Write a Python function first_num_greater_than(numbers_list, key) that takes a list of integers (numbers_list) and an integer key (key), and returns the first number in the list that is greater than the key. If there is no number greater than the key, then you should return None.

Test Suite Development

Add Test Case:

Input Expected output



Add Test Group:

Enter test group name +

Evaluate Test Suite

Passed? Actual

greater_than	✓	3	delete
greater_than	✓	2	delete

Student's Current Code

```

1 def first_num_greater_than(numbers_list, key):
2     for i in range(len(numbers_list)):
3         if numbers_list[i] > key:
4             return numbers_list[i]
5     return None

```

View Code Differences

```

for i in range(len(numbers_list)):
    if numbers_list[i] > key:
        return numbers_list[i]
else:
    return None

```

1. No number in list greater than key

```
assert(first_num_greater_than([3, 2, 1], 3) == None) ✓ None delete
```

2. Key in middle of list

```
assert(first_num_greater_than([3, 2, 1], 2) == 3) ✓ 3 delete
```


Office Hour!



Hi Christina! Here is my code and I think I have some problem with it. Can you walk me through that?

Select... | v

Test suite evaluated! Select a test case to start explain.



Problem: first_num_greater_than

Write a Python function `first_num_greater_than(numbers_list, key)` that takes a list of integers (`numbers_list`) and an integer key (`key`), and returns the first number in the list that is greater than the key. If there is no number greater than the key, then you should return `None`.

Now you are chatting with a student. Please explain to them why their code is wrong by selecting the right explanation from the list. If you are right, the student will fix their code accordingly! Otherwise, they may get frustrated and leave.

Student's Current Code

```
1 def first_num_greater_than(numbers_list, key):
2     for i in range(len(numbers_list)):
3         if numbers_list[i] > key:
4             return numbers_list[i]
```

Test Suite Development

Add Test Case:

Input Expected output

+

Add Test Group:

Enter test group name +

Evaluate Test Suite

Your Test Cases

Passed? A

Default Group

```
assert(first_num_greater_than
([1, 2, 3], 2) == 3)
```

✓ 3

```
assert(first_num_greater_than
```

✓ 2

View Code Differences

```
for i in range(len(numbers_list)):
    if numbers_list[i] > key:
        return numbers_list[i]
else:
    return None
```


```
for i in range(len(numbers_list)):
    if numbers_list[i] > key:
        return numbers_list[i]
return None
```

Office Hour!

Your function returns `None` as soon as it encounters a number that is smaller than the key, it only continue to check the loop if `num` is equal to `key`. But there might be numbers greater than the key later after a number smaller than `key`.




But if that's the case, shouldn't `first_num_greater_than([1, 2, 3], 1)` give me 2? Currently my code would output `None`.



Ok, I see my code got this test case wrong. Could you explain what's wrong with my code?

Your code returns `None` if the first number in the list is not greater than the key. It doesn't check the rest of the numbers in the list



Now I've updated the code with your fix. Is it good now?

Is the student correct?

Yes

No

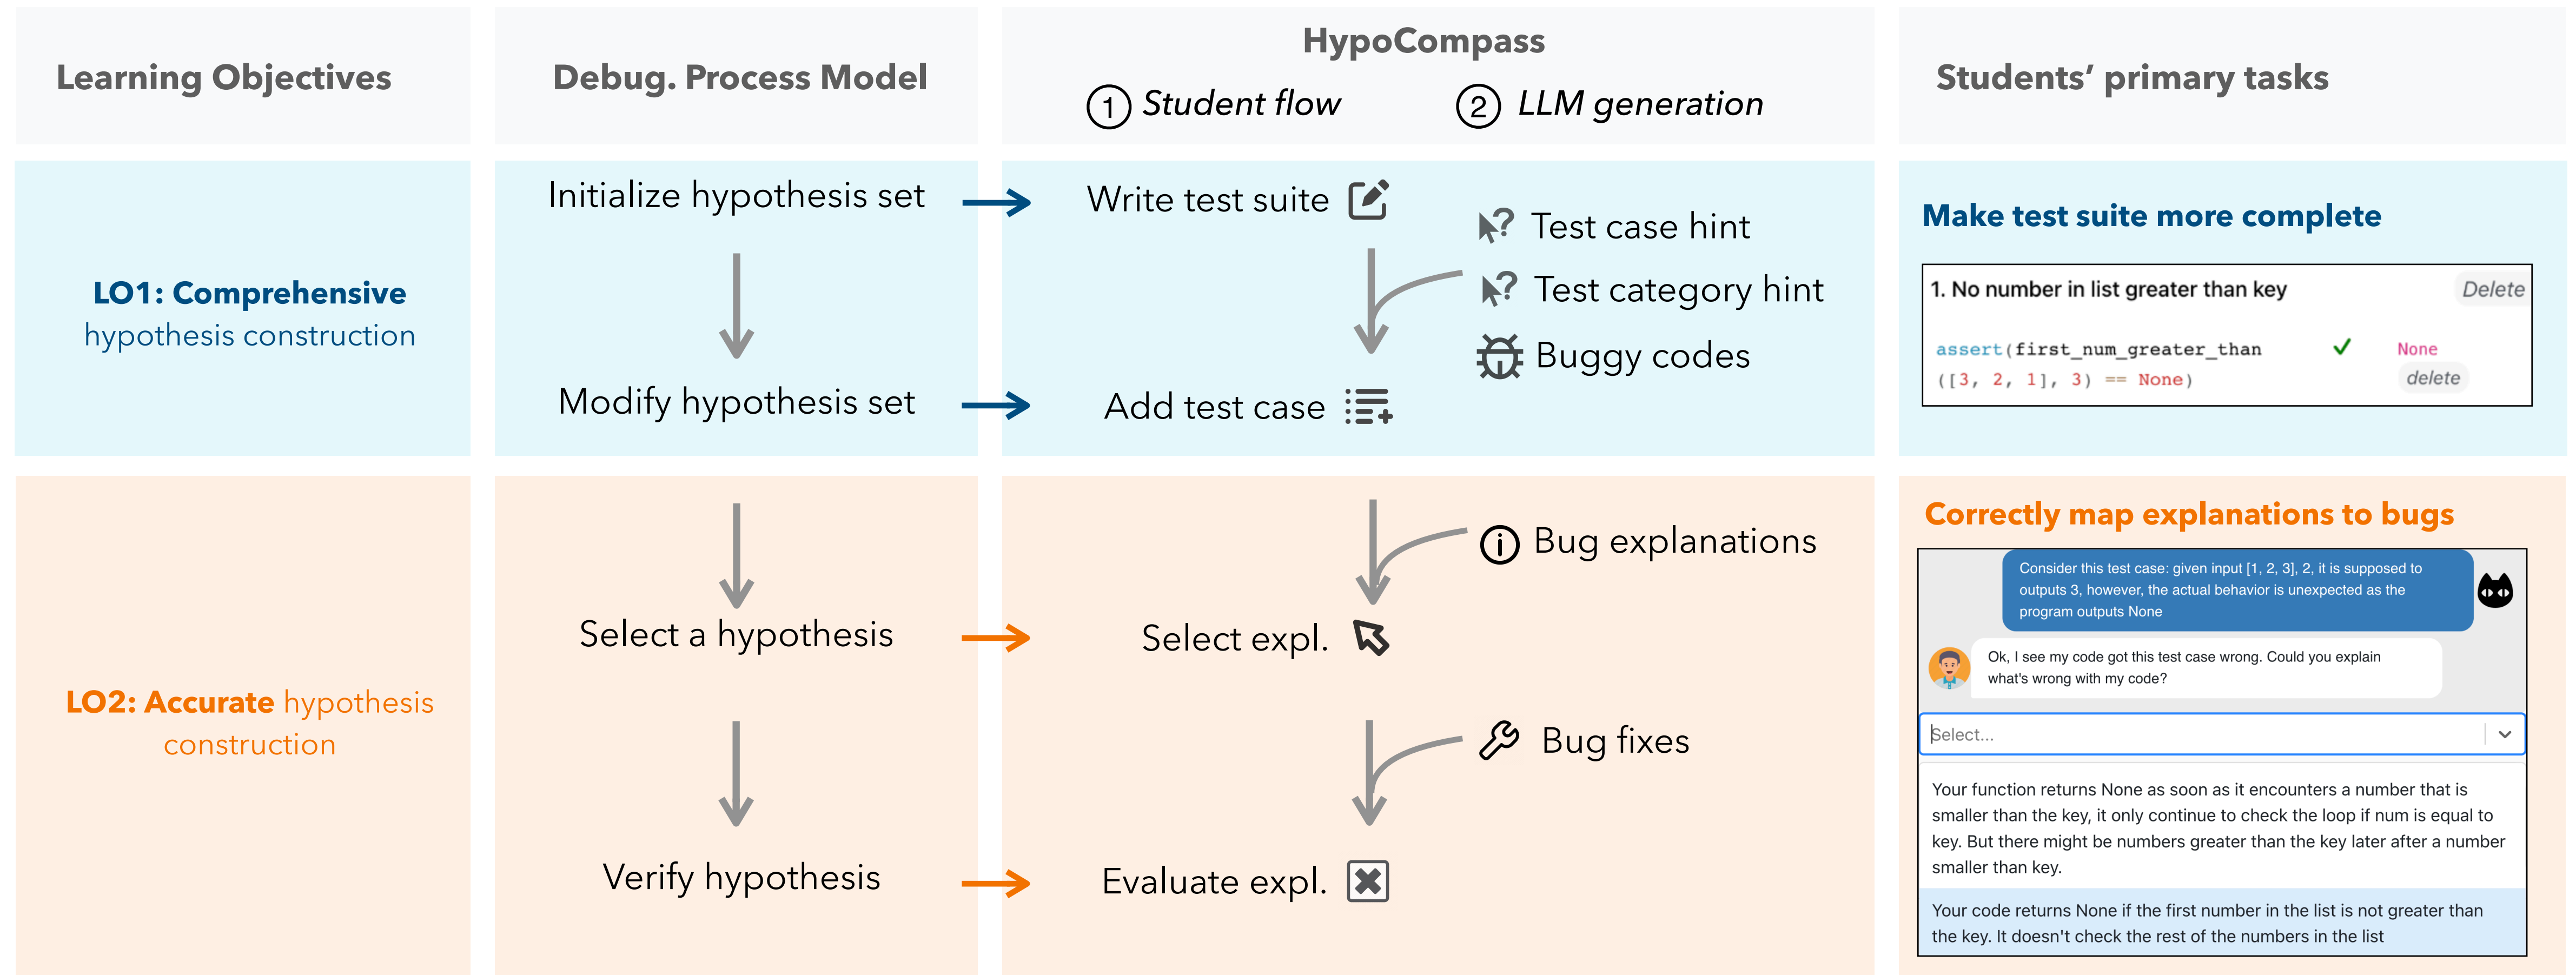
I don't know



All of your test cases passed.

Send

HypoCompass: Learning Theory Inspired Design



Explicit & deliberate training just on debugging, by off-loading other necessary sub-tasks to LLMs (e.g., writing the bug, correcting the bug, etc.)

HypoCompass: Effectiveness

HypoCompass consistently generates high-quality training materials: 90% success rate + only take 15 minutes to label and edit (reduce instructor effort!)

HypoCompass brings learning gain: In a pre-to-post test setup, 10 novices improved their performances by **17%**, with a reduced completion time of **13%**.

It is possible to eventually train students to be better at debugging!

Question 2 Select all that apply.

Imagine a solution to this problem that fails on all of these following test case(s) because of the same bug. What could the bug be?

Python

```
assert(num_smaller([10, 10, 10, 20, 30], 10) == 0)
assert(num_smaller([10, 20, 20, 30, 30], 20) == 1)
assert(num_smaller([10, 10, 20, 30, 30], 30) == 3)
```

Select all that apply.

- A. The buggy codes may have overlooked duplicated elements in seq
- B. The buggy codes may have overlooked x when it is the smallest element in seq
- C. The buggy codes may only handle the case where x is not in seq
- D. The buggy codes may only handle the case where x is already in seq

Answer:

Question 3.1 Select one answer.

Given the addition which test suite works

Question 7 Select one answer.

Test case 2: `assert(remove_extras_code2([1, 1, 2, 3]) == [1, 2, 3])`
Actual behavior: `'TypeError' 'int' object is not iterable.`

Python
asse
asse
asse

- A.
- B.

Python

```
1 def remove_extras_code2(lst):
2     new_lst = []
3     for i in lst:
4         if i == lst[i+1]:
5             continue
6         else:
7             new_lst += i
8     return new_lst
```

What's the bug exposed by this test case?

- A. The bug occurs because the loop variable `i` is mistakenly used as both the element and index of the list. This leads to incorrect comparisons and triggers a `TypeError` in `lst[i+1]` because `i` is an element of the list, not an index.
- B. The bug is caused by not initializing the `new_lst` properly. The code fails to explicitly assign an empty list to `new_lst`, so when concatenating elements to `new_lst` using the `+=` operator, a `TypeError` occurs because `new_lst` is not iterable.
- C. The bug is due to an incorrect conditional statement. The code incorrectly compares `i` with `lst[i+1]` instead of comparing adjacent elements of the list, which triggers `TypeError` when trying to compare an integer `i` with a list element.
- D. The bug occurs because the code incorrectly assumes that `i` is iterable when concatenating it to `new_lst` with the `+=` operator. In this case, `i` is an integer, which is not iterable, and it causes a `TypeError`.

HypoCompass: Learning Theory Inspired Design

Learning Objectives

Explicit & deliberate training just on debugging, by **off-loading** other necessary sub-tasks to LLMs (e.g., writing the bug, correcting the bug, etc.)

Debug. Process Model

HypoCompass

① Student flow

② LLM generation

Test suite



Test case hint

Test category hint

Buggy codes

Test case



① Bug explanations

Select expl.



Bug fixes

Evaluate expl.



Students' primary tasks

Make test suite more complete

1. No number in list greater than key

Delete

```
assert(first_num_greater_than  
([3, 2, 1], 3) == None)
```

✓ None

delete

Correctly map explanations to bugs

Consider this test case: given input [1, 2, 3], 2, it is supposed to output 3, however, the actual behavior is unexpected as the program outputs None

Ok, I see my code got this test case wrong. Could you explain what's wrong with my code?

Select...

Your function returns None as soon as it encounters a number that is smaller than the key, it only continue to check the loop if num is equal to key. But there might be numbers greater than the key later after a number smaller than key.

Your code returns None if the first number in the list is not greater than the key. It doesn't check the rest of the numbers in the list

There will be skills that can be offloaded to LLMs. What skills to offload and, in turn, what skills to train humans on, become an interesting HCI question.

Explicit & deliberate training just on debugging, by off-loading other necessary sub-tasks to LLMs (e.g., writing the bug, correcting the bug, etc.)

Task formation in HypoCompass (bug fixing)

LLM task: To edit the buggy code according to the fix instruction without over- or under- fix

```
def first_num_greater_than (numbers_list, key):  
    for i in range(len(numbers_list)):  
        if numbers_list[i] <= key:  
            return num  
    else:  
        return None
```

"Change the comparison line to be larger than key."


You fix bugs in Python code closely following the instructions.
Original code: {buggy_code};
Code modification: {explanation}
Modified code:

Over-fixing, because LLM wants to continue to generate correct code!

```
def first_num_greater_than (numbers_list, key):  
    for i in range(len(numbers_list)):  
        if numbers_list[i] > key:  
            return num  
    return None
```

Task formation in HypoCompass (bug fixing)

LLM task: To edit the buggy code according to the fix instruction without over- or under- fix



```
def first_num_greater_than (numbers_list, key):  
    for i in range(len(numbers_list)):  
        if numbers_list[i] <= key:  
            return num  
        else:  
            return None
```

"Change the comparison line to be larger than key."

You fix bugs in Python code closely following the instructions.

Original code: {buggy_code};

Code modification: {explanation}

Translate the statement into actual, minimal code change in this format:

{original code snippet: "copy the lines of code that need editing"}

-> edited code snippet: "write the edited code snippet"}

Task formation helps avoid competing tasks of code editing and code completion!

`numbers_list[i] <= key` → `numbers_list[i] > key`

That's all for today!!