Human-aware AI for Procedural Content Generation

Zhiyu Lin

UC SANTA CRUZ

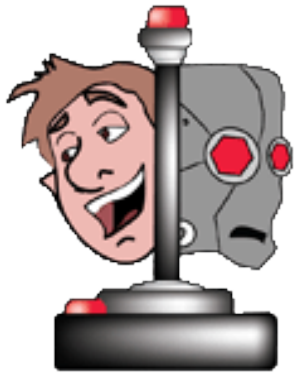

Georgia Tech

Who am I?

- Postdoctoral Fellow at University of California, Santa Cruz
  - Working at Game User Interaction and Intelligence Lab in Computational Media
- PhD in Computer Science from Georgia Institute of Technology
- Research focus: Especially in computational creative domains, how human and computers can collaborate and excel beyond either human or AI alone
  - Mixed-Initiative and Co-creative systems
  - Generative AI, Reinforcement Learning
  - Human's role in human-AI collaboration, AI and education w/ Gamification…

In a nutshell: My stories on how to…

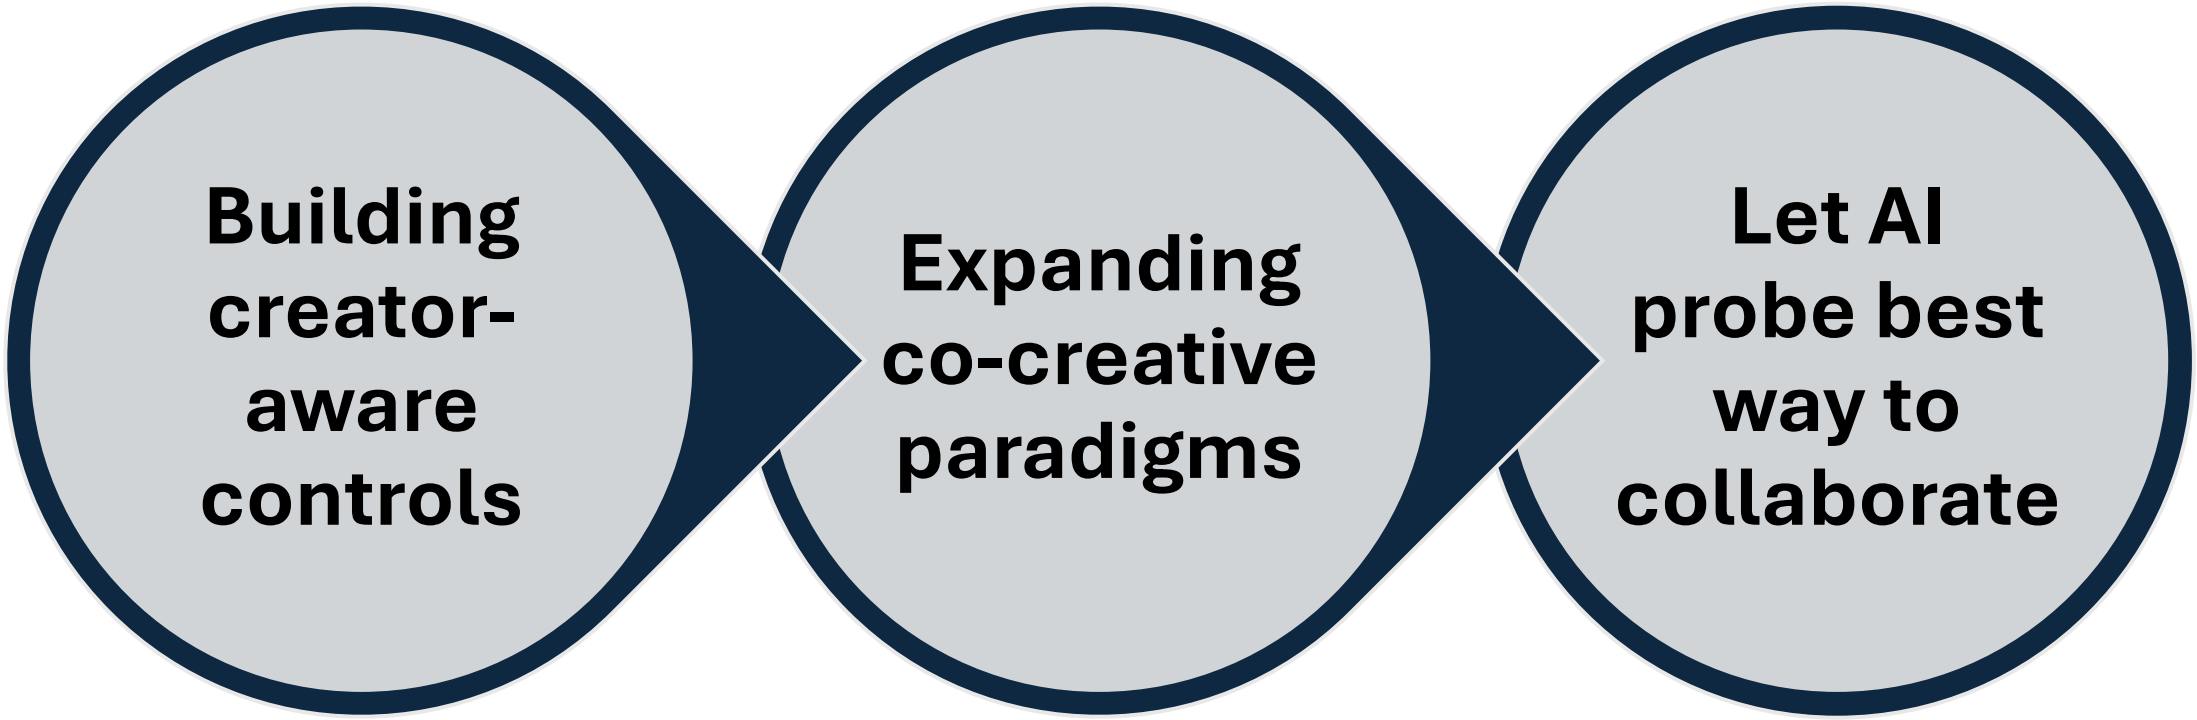

Let’s make this interactive!

AI went through a long way

So do computational creativity

A Chinese professor used AI to write a science fiction novel. Then it was a winner in a national competition

- Journalism professor Shen Yang plans to detail his creation process so anyone can 'create good fiction with AI'
- But artificial intelligence poses threats to writers and irreversible damage to literary language, a publisher says

These models are good at generating contents *like what they learned from*.

Is this the *end of the story*?

Challenge: Procedural Content Generation (PCG)

ALGORITHMIC CREATION OF **GAME CONTENT** WITH ***LIMITED OR INDIRECT USER INPUT*** (SHAKER ET AL. [2016])

Super Mario (1985)

Goal of game

Music

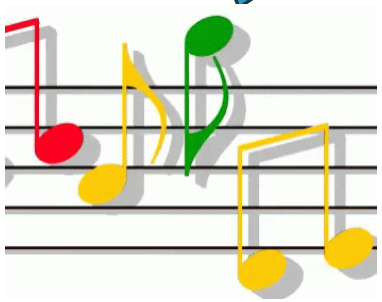

Graphics

Difficulty progression

Stage / Objects

Character

Interactive Mechanism

Narrative

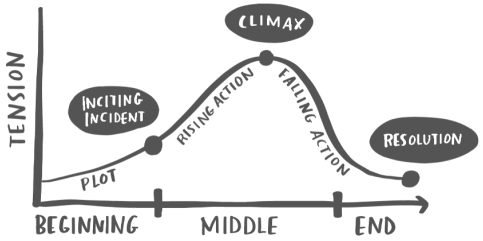

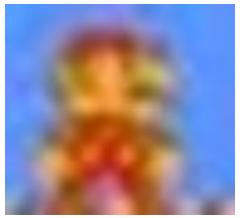

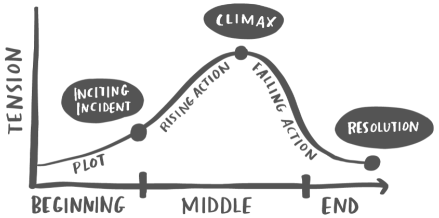

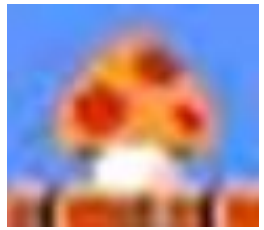

Based on “Mario saving the mushroom kingdom”

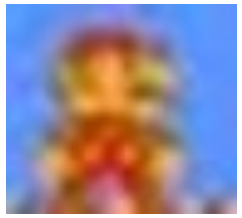

Objects used to teach new mechanisms

Elements correlates and constrain each other

A model generating arbitrary contents are not enough for PCG

Controllability

We can make better models, filter training data, add postprocessing steps, prompt engineer…

Controllability… **From the perspective of the model**

We can make better models, filter training data, add postprocessing steps, prompt engineer…

Revisit: ML-based AI

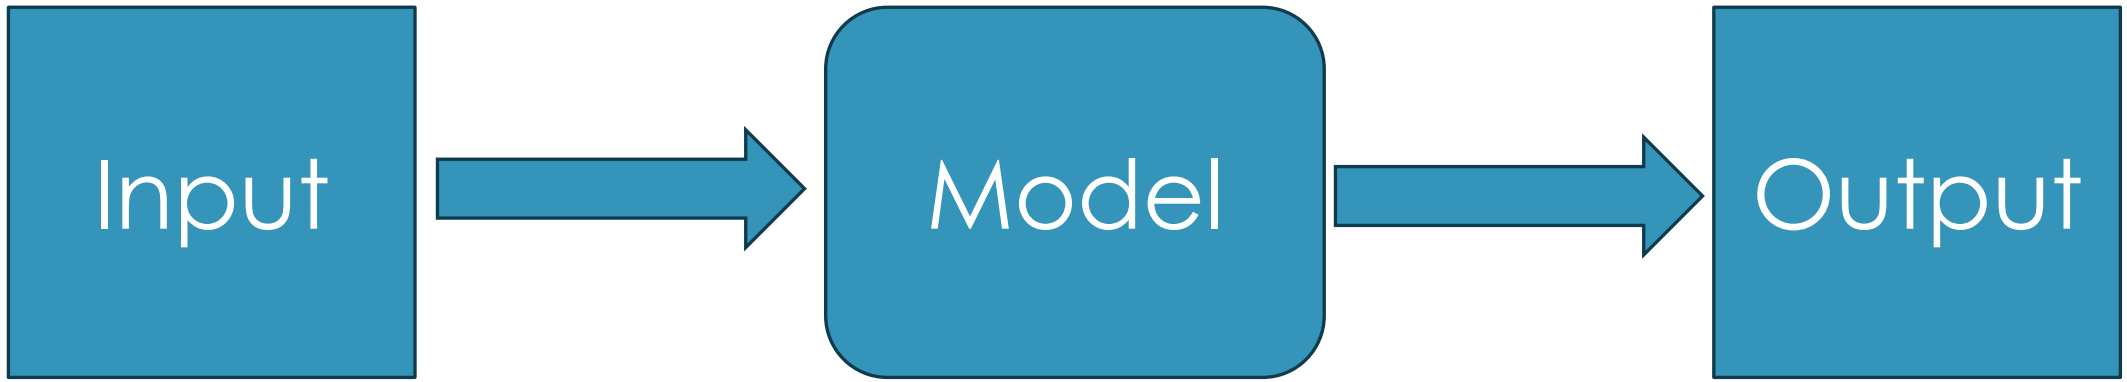

Bad input may lead to bad output!

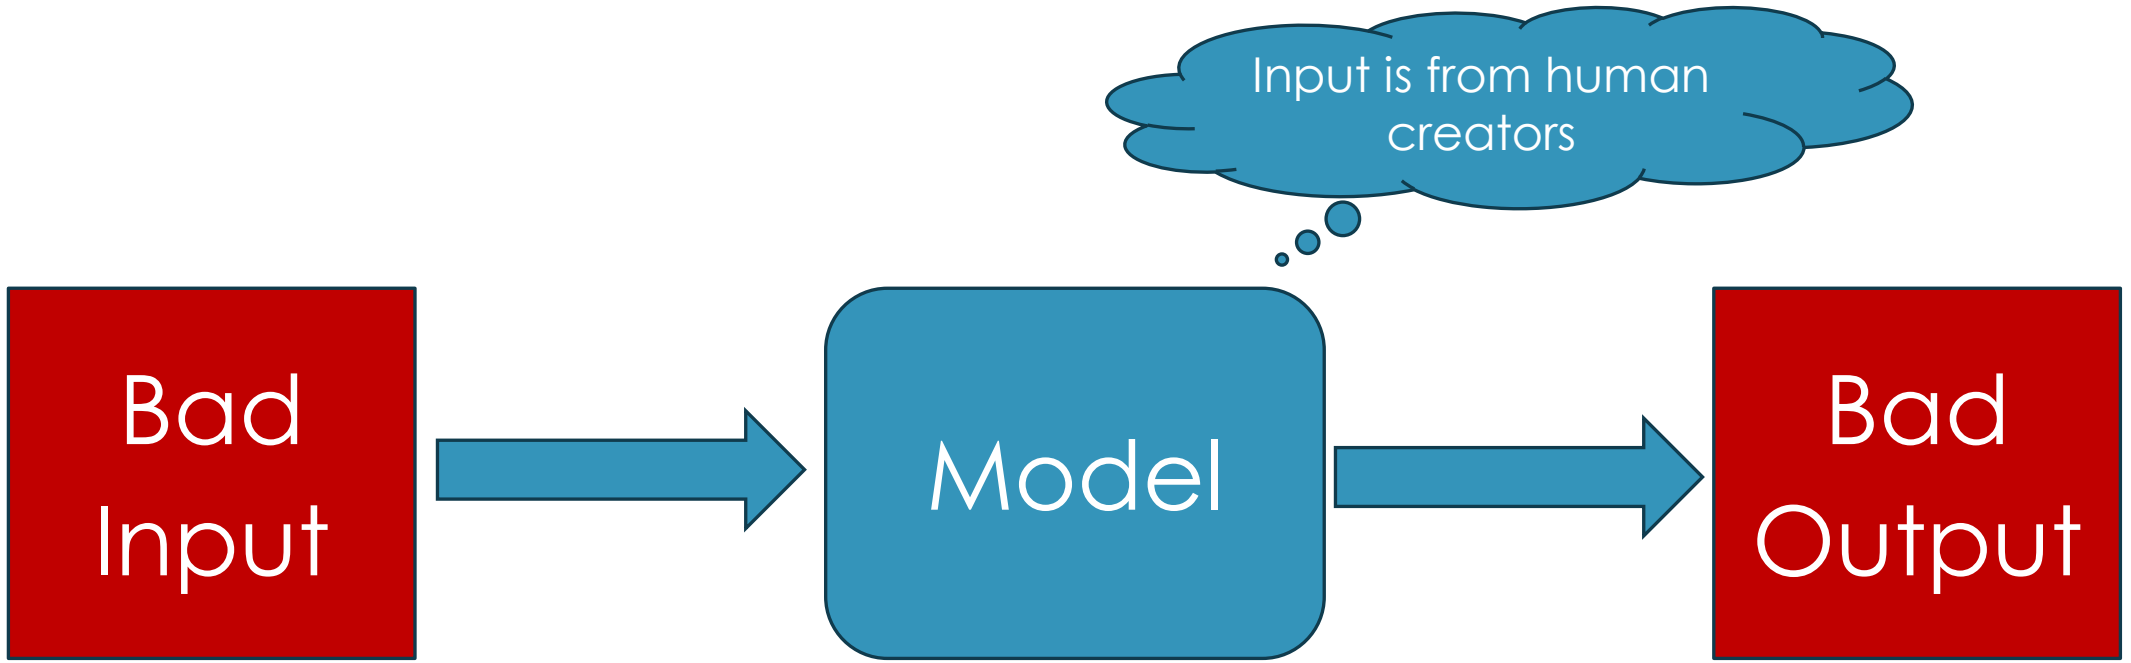

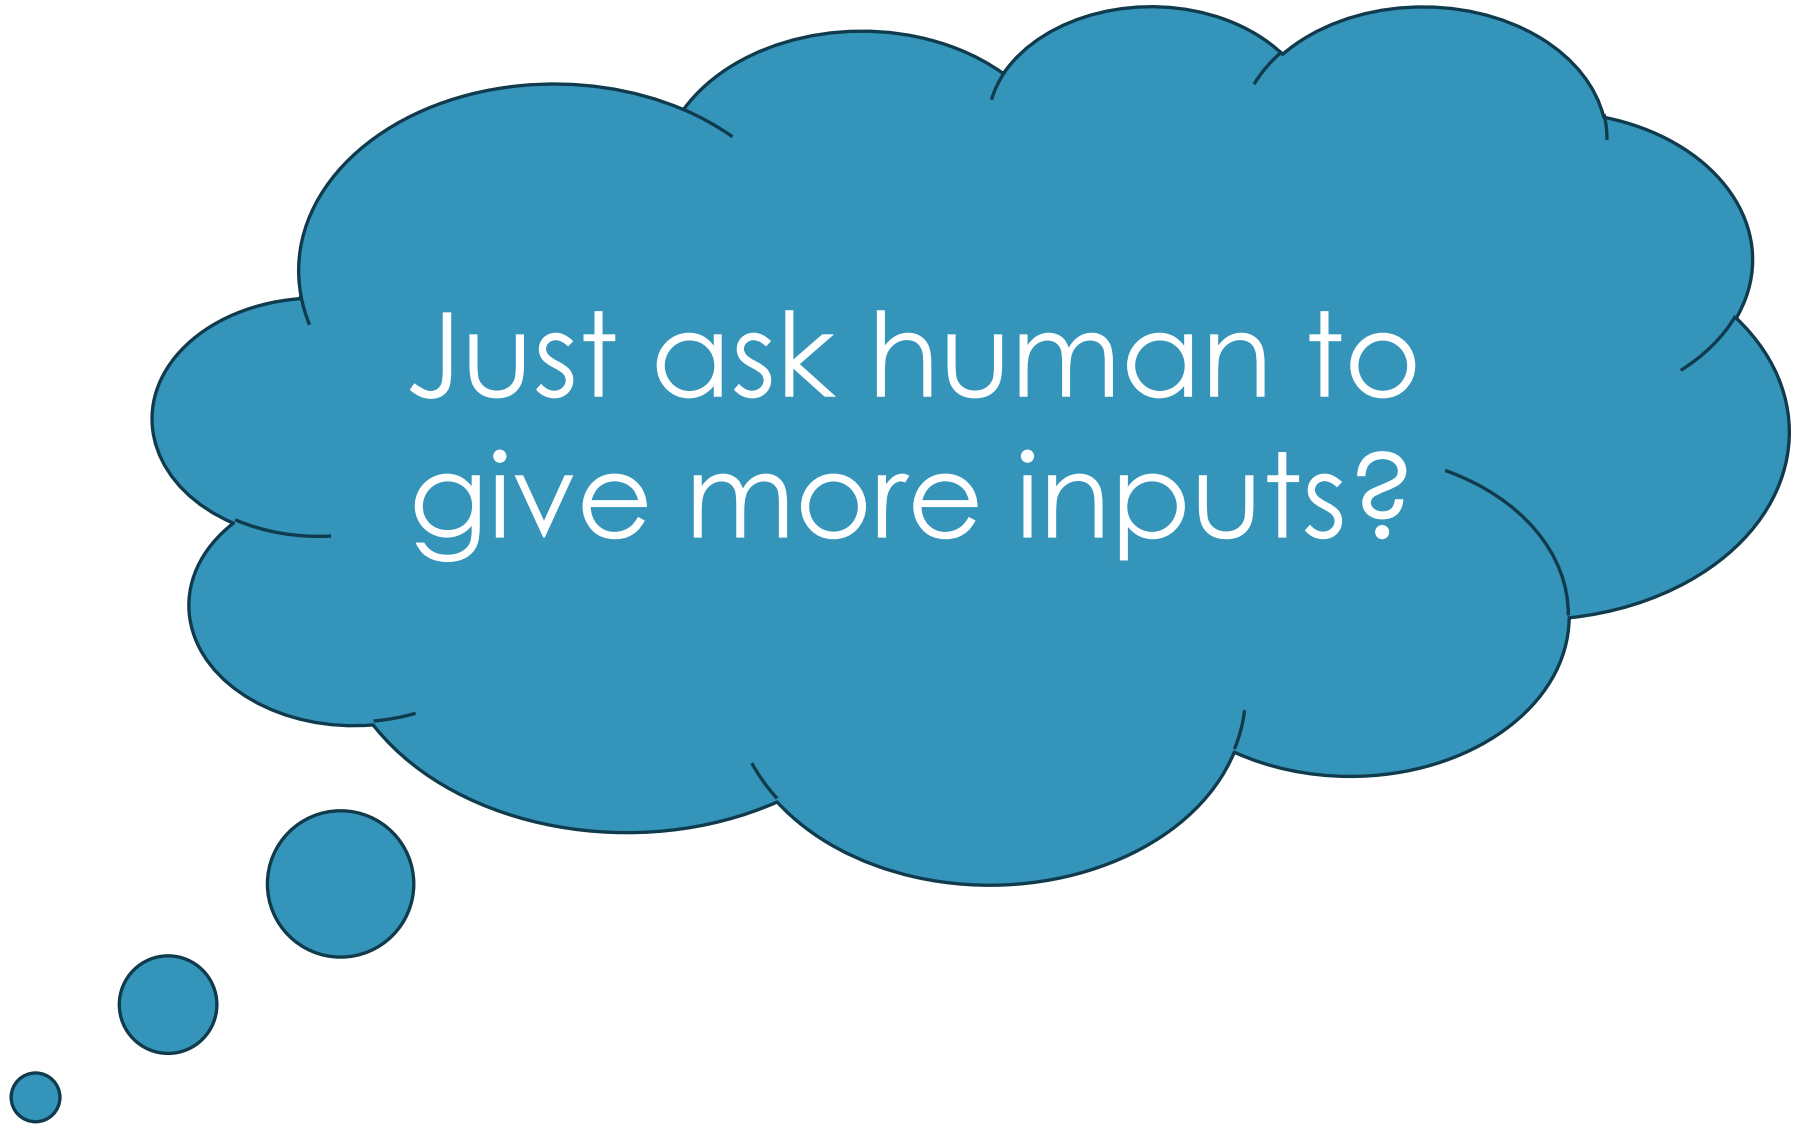
Just ask human to
give more inputs?

“More inputs” not as straightforward!

Creators may have trouble communicating	• We are asking the creators for a different skillset than creating contents!
Creators may not know what the final product looks like	• We are asking a concise prompt which may never be possible!
AI may not understand the creator	• AI are interacting with the user without enough information!
There may never be a “best action”!	• There are more than one way to improve the contents!

Controllability from the perspective of human

"This control was clearly not designed with me in mind!"

"I want to control them in diverse ways that work for me."

"There are so many kinds of control, which one should I use?"

INCREASED COGNITIVE LOAD

FRUSTRATION AND FAILURES

PREVENTING EQUITABLE USE

Issues

Spotlight: Mixed-Initiative Co-Creativity

Mixed-Initiative	Human initiative and a computational initiative cooperate towards a shared goal
Co-Creative	System having the capability to modify the contents as if it's the human counterpart

My goal

Improve the experience of the human creators using AI Procedural Content Generation systems

My Path

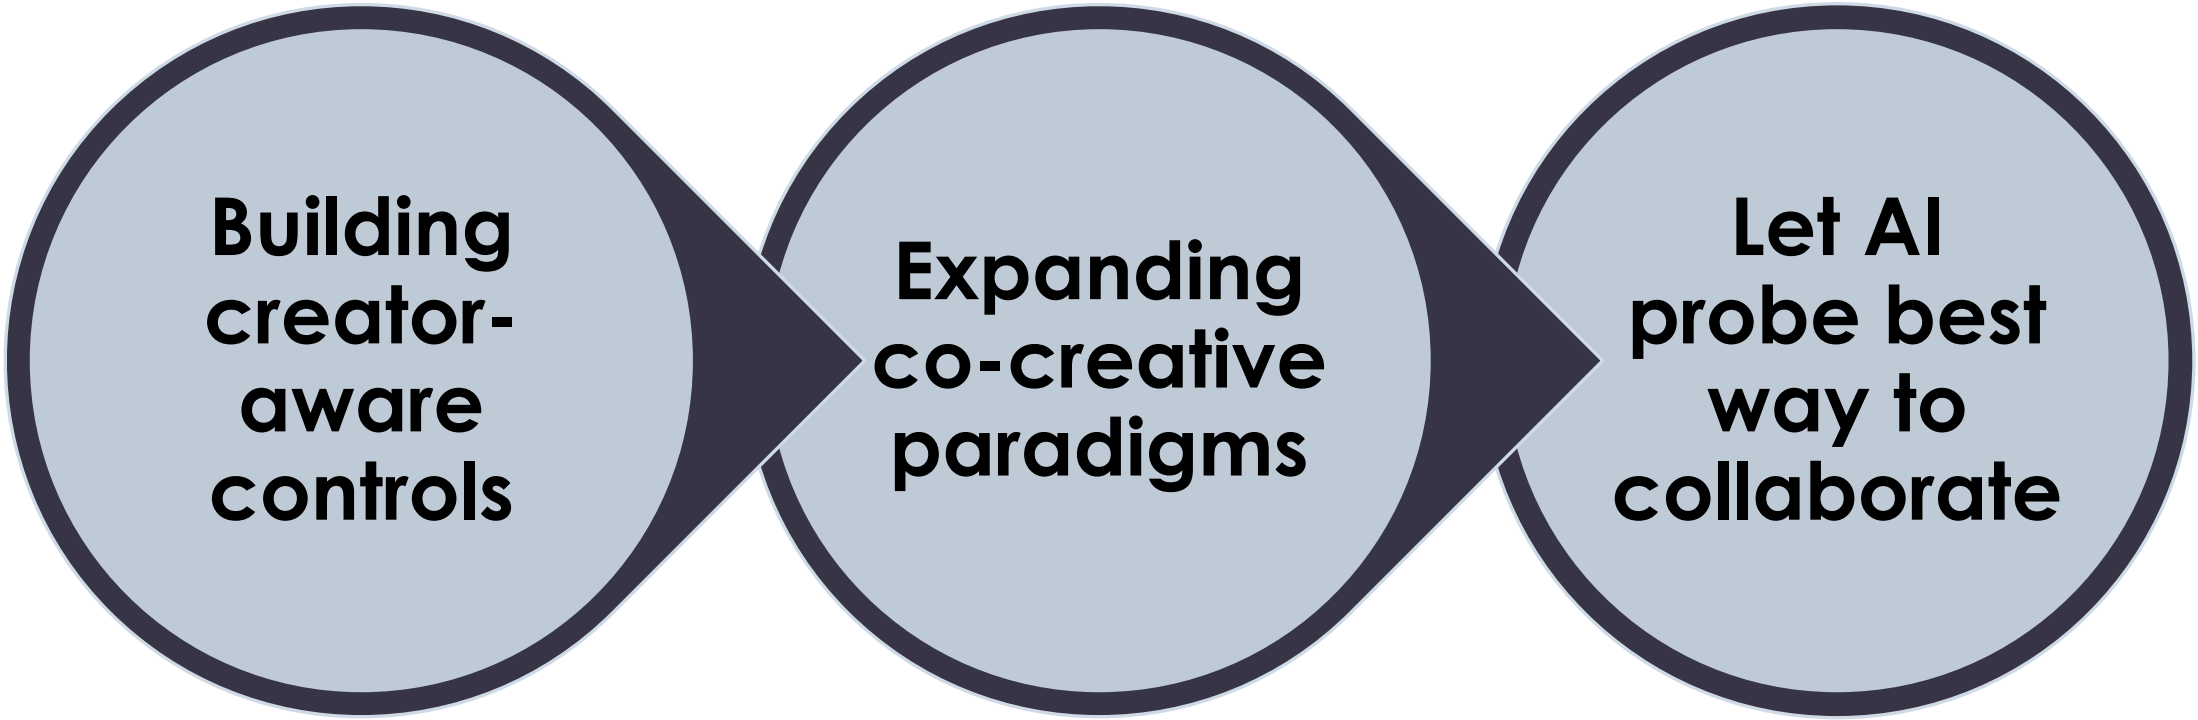
Building creator-aware controls
Expanding co-creative paradigms
Let AI probe best way to collaborate

Building creator-aware controls

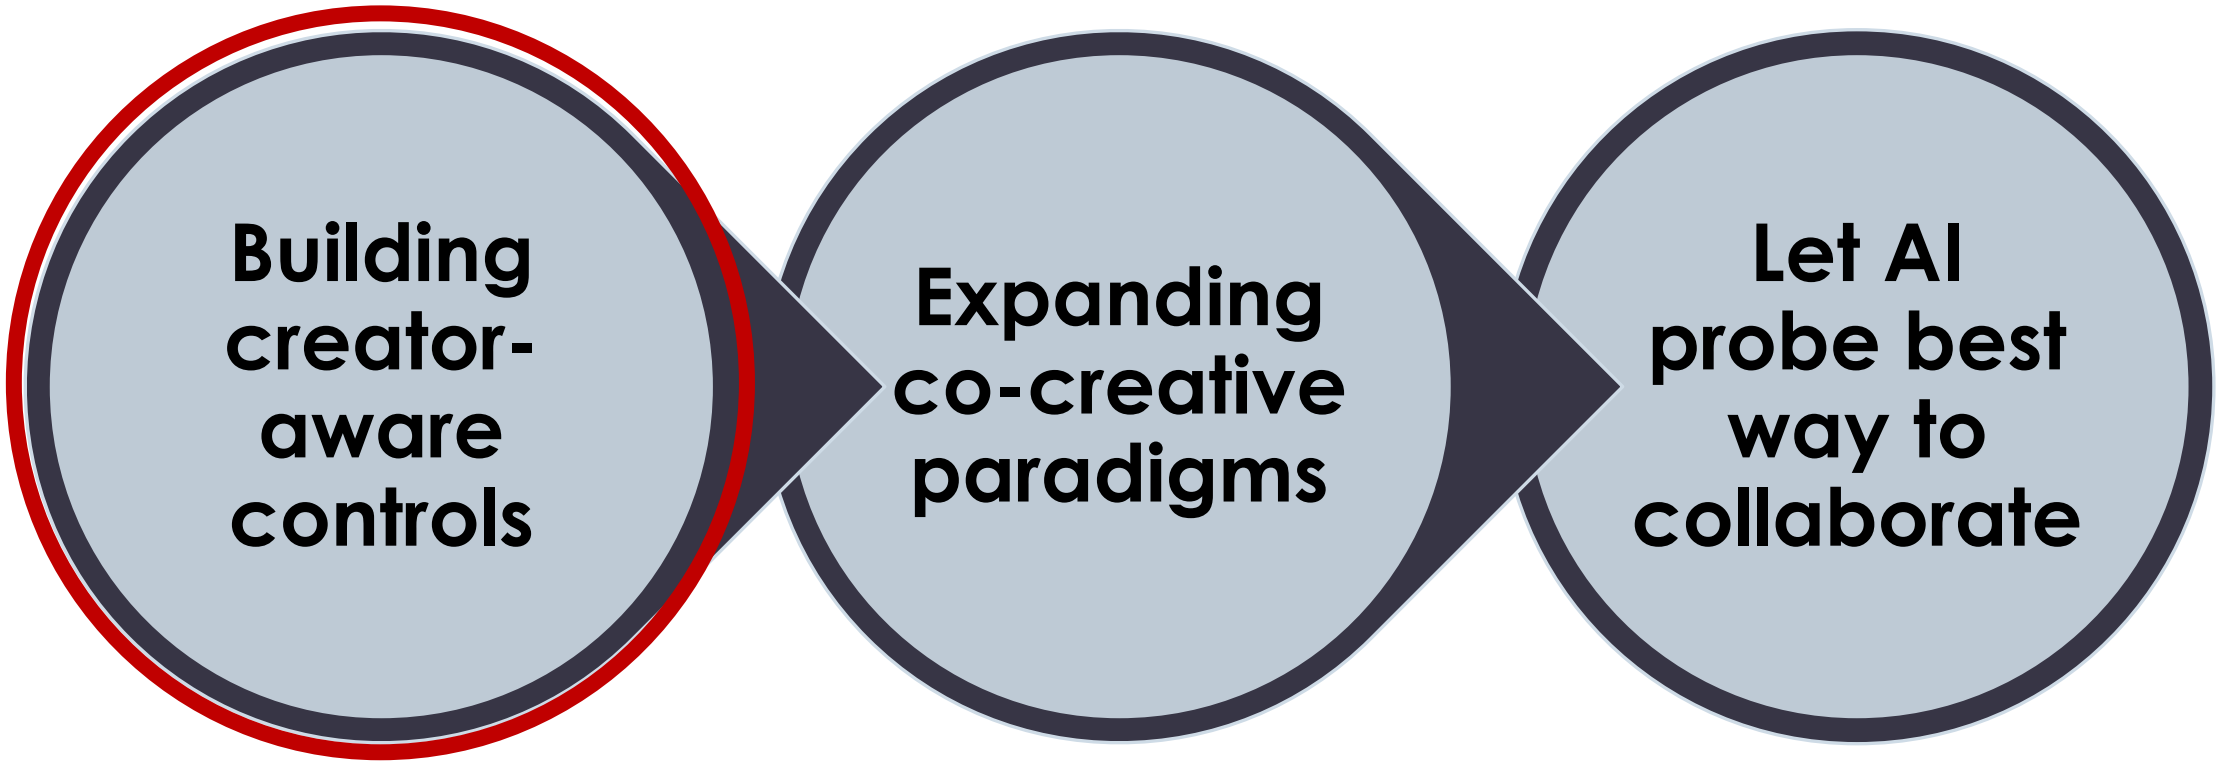

Are controls provided by AI **creator-aware** and designed for them?

Instead of imposing paradigms based on underlying models and algorithms?

Lin, Zhiyu, Kyle Xiao, and Mark Riedl. "Generationmania: Learning to semantically choreograph." AIIDE 2019.

Domain 1: Rhythm Action Games

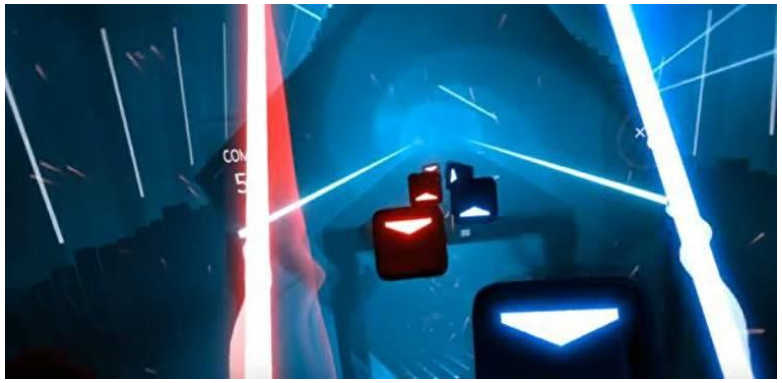

Focus: Beatmania

Recreate music as DJ

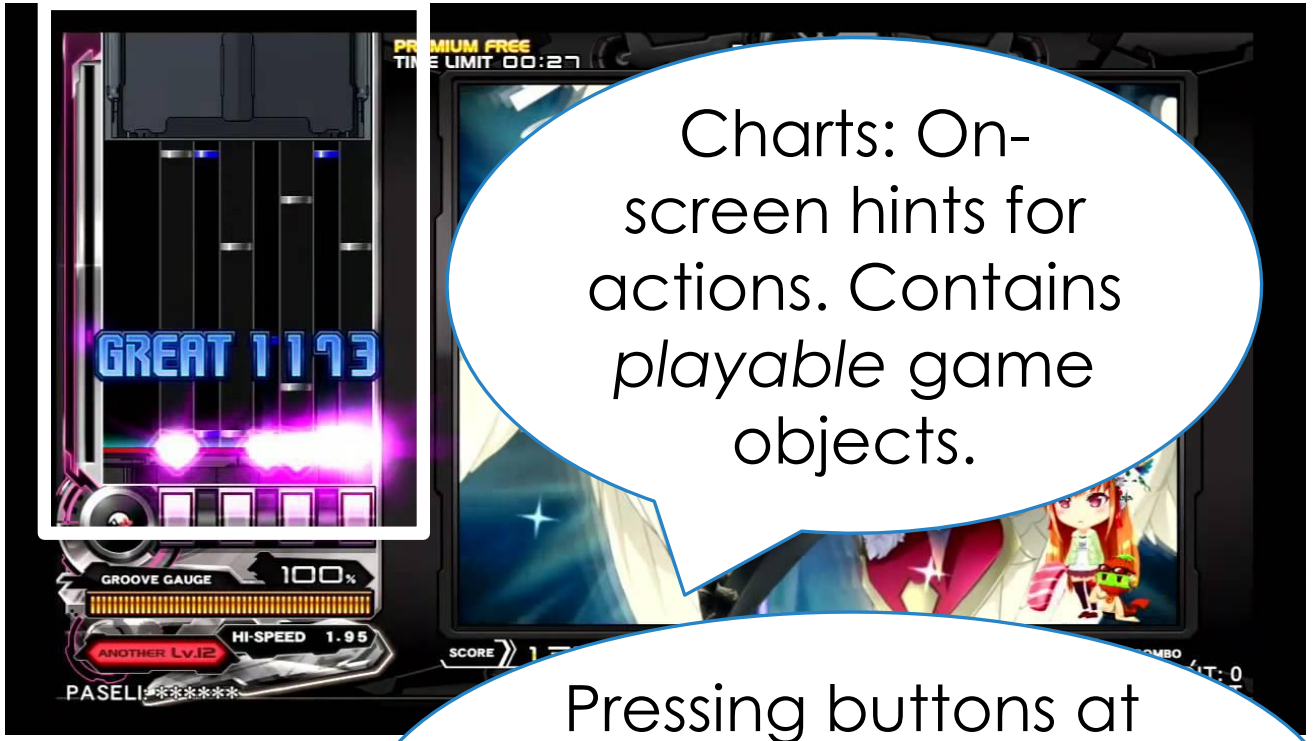

Charts: On-screen hints for actions. Contains *playable* game objects.

Pressing buttons at the correct time plays audio samples to recreate music

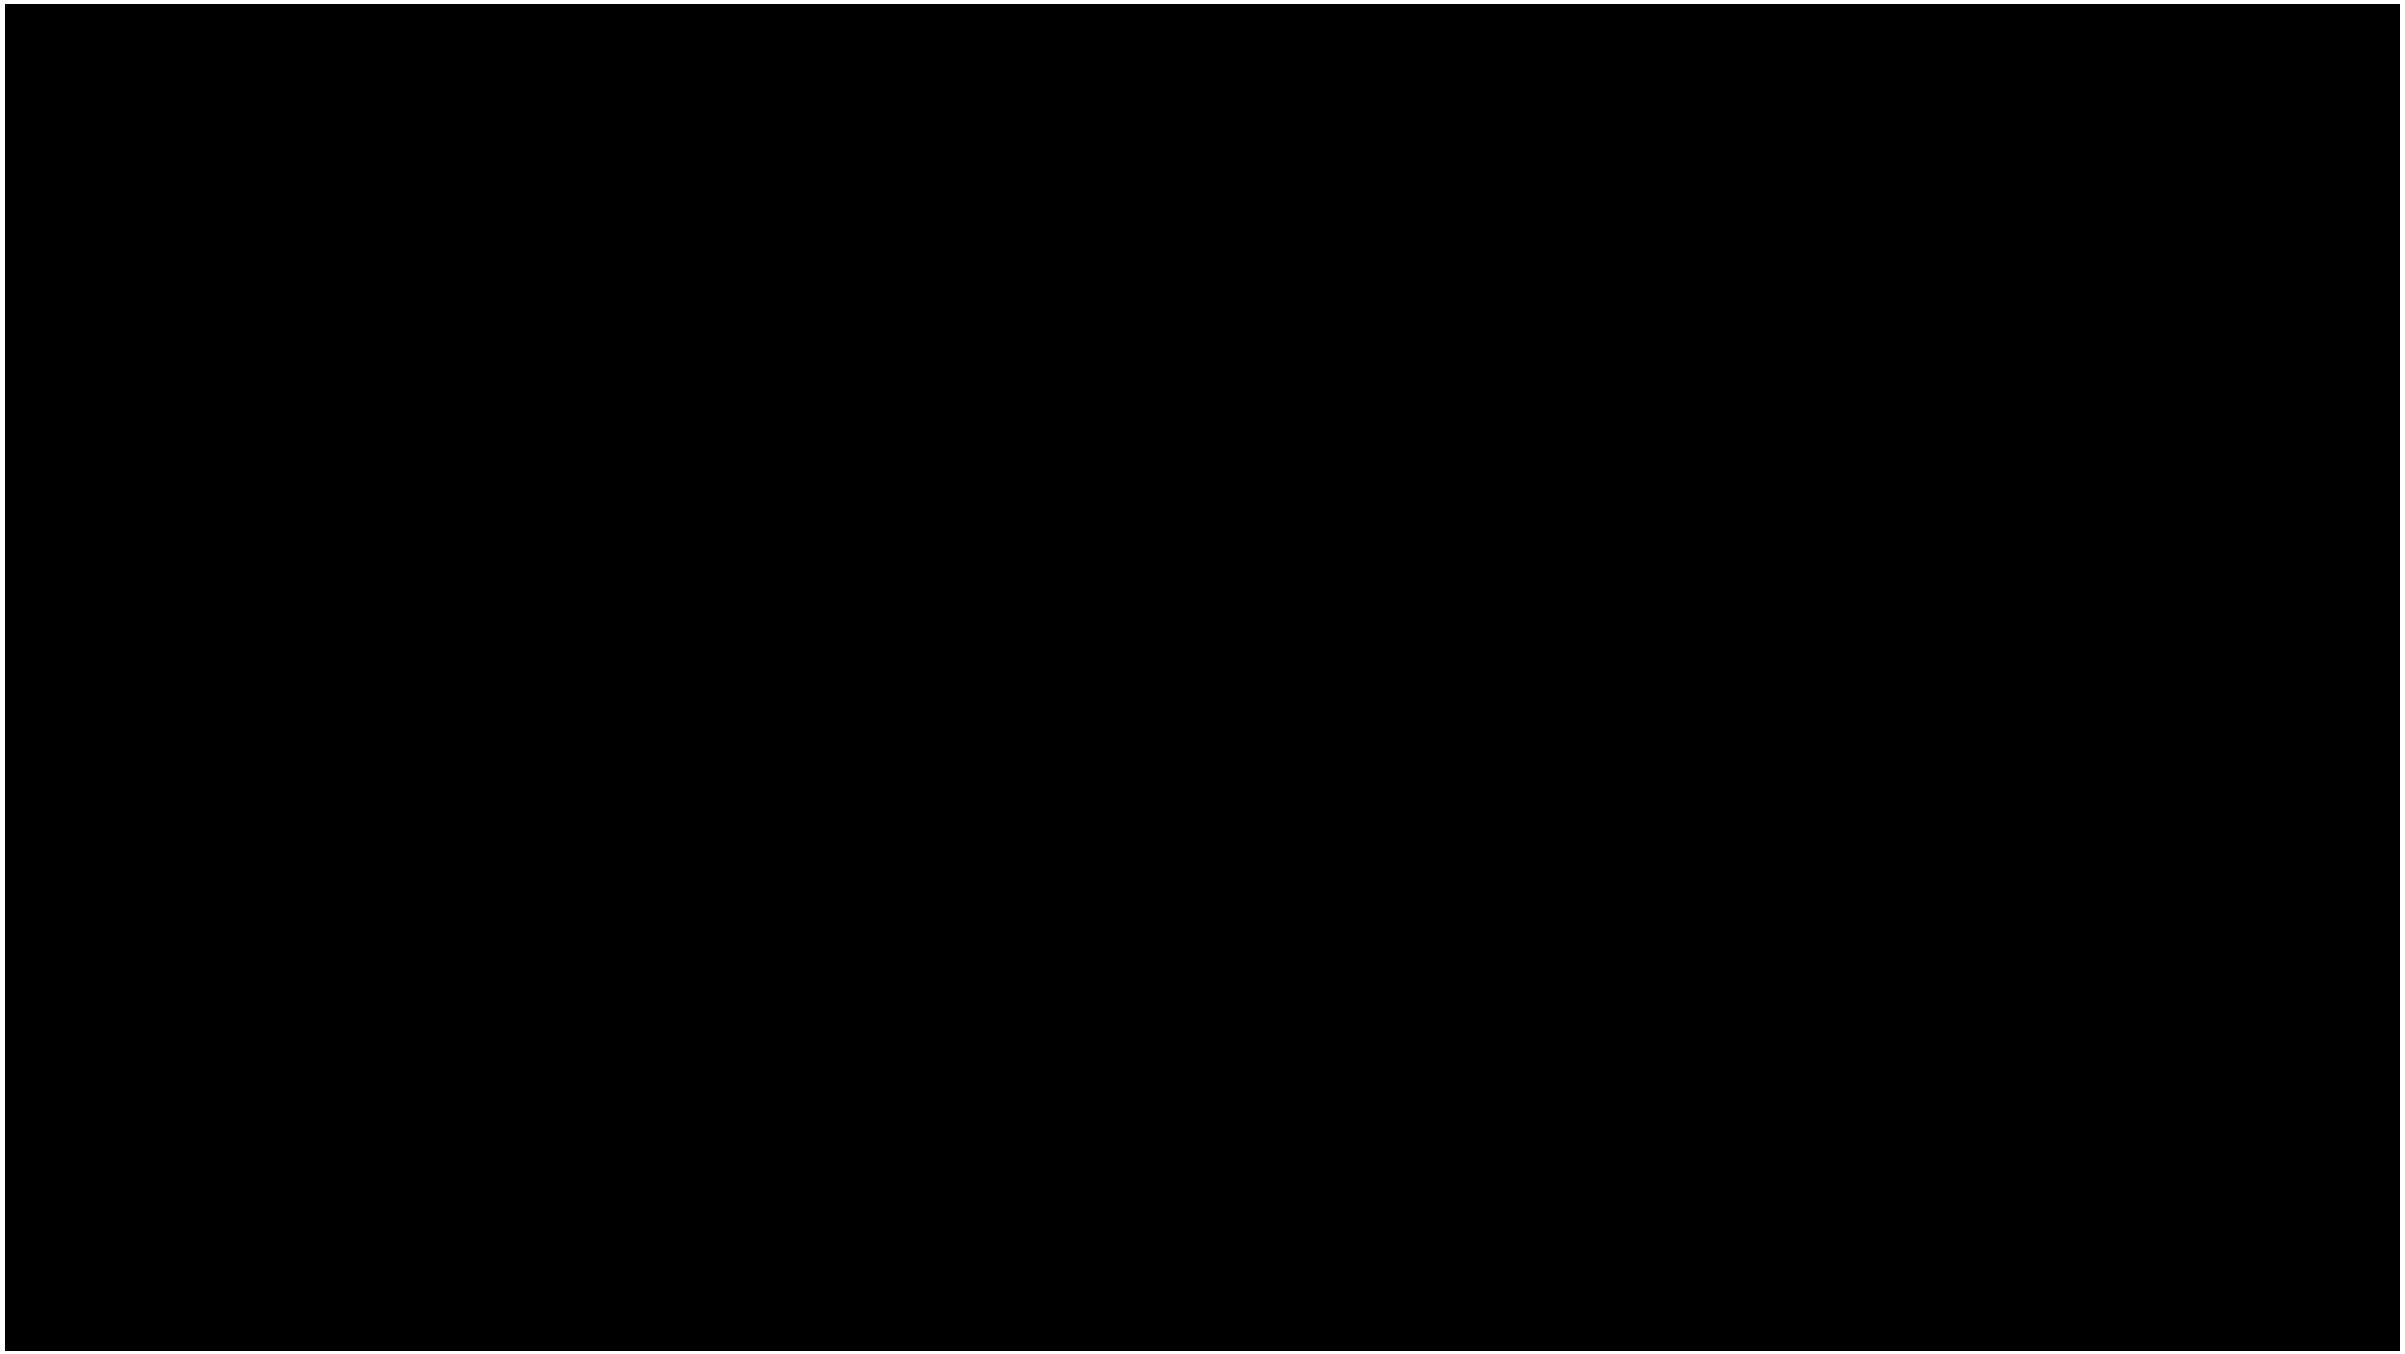

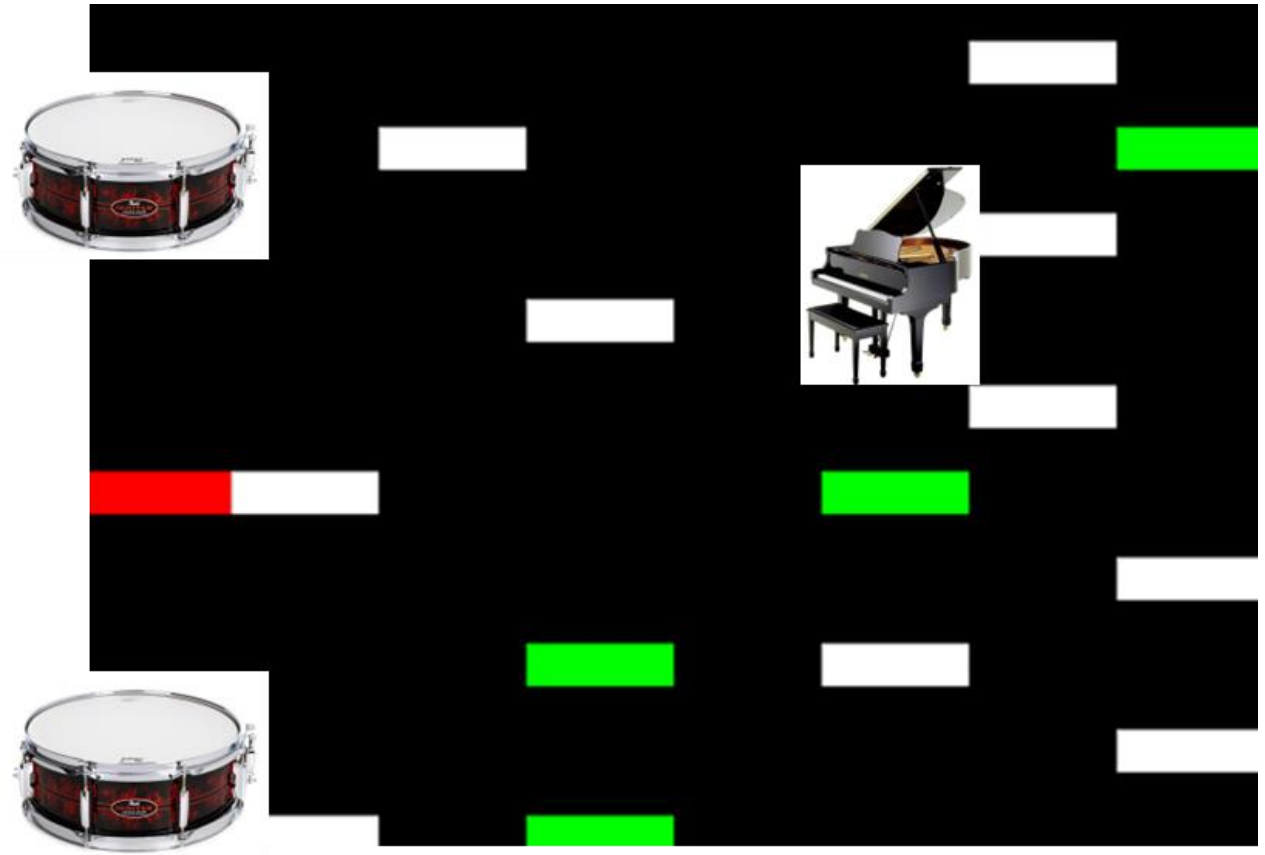

Considerations for creator-awareness

- Music to follow
- Intended challenge level
- Artistic intents

Our Solution: *GenerationMania*

- Generate music-grounded charts
- Enable creator-aware controls

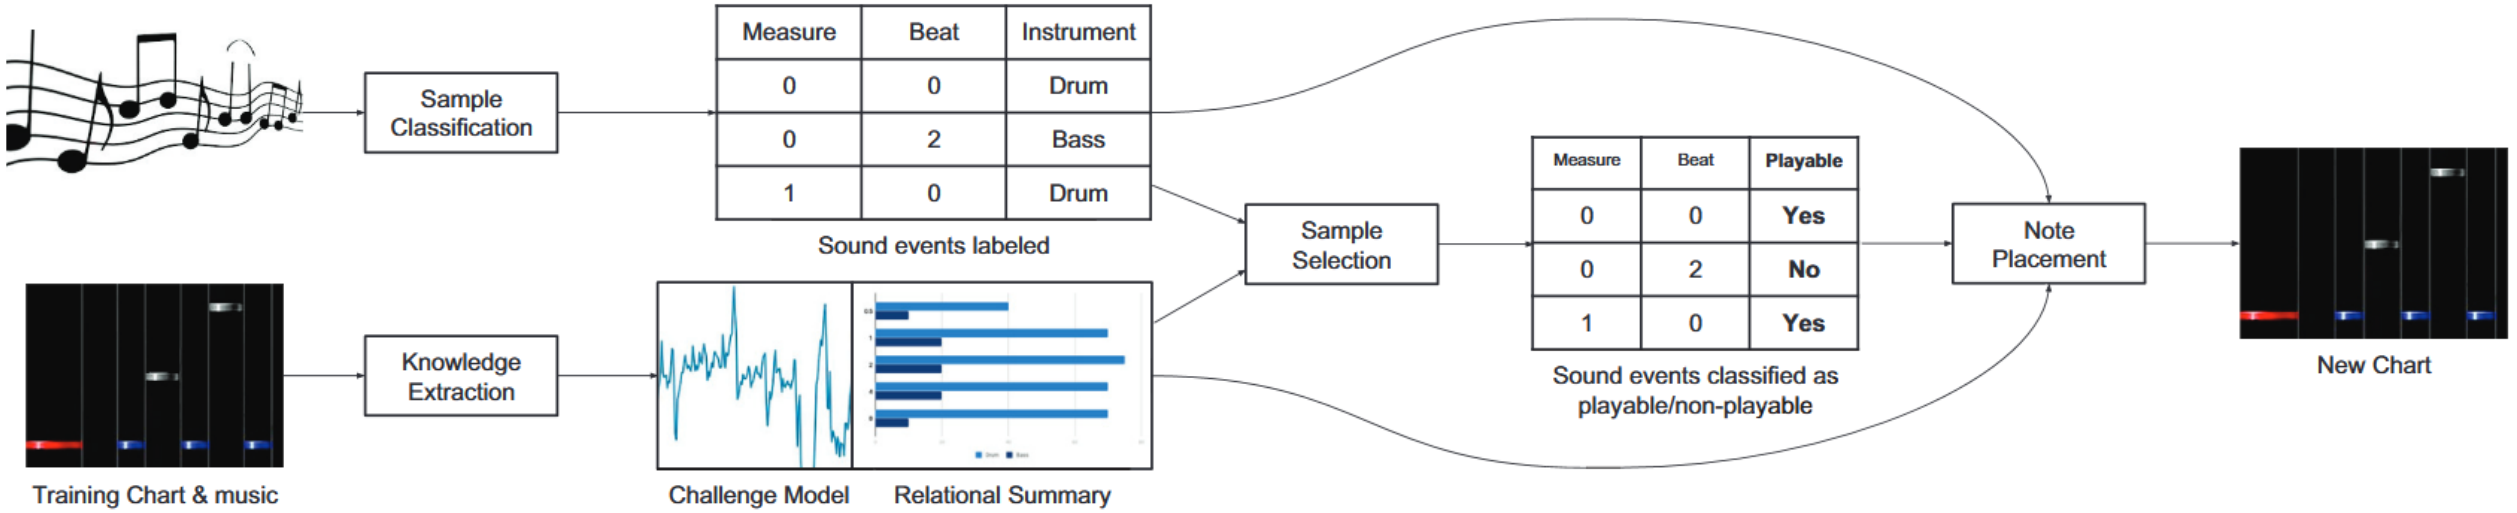

Highlighted control #1: Intended Challenge Level

I want to warm up the player with an easy beginning

But challenge the player halfway into the stage

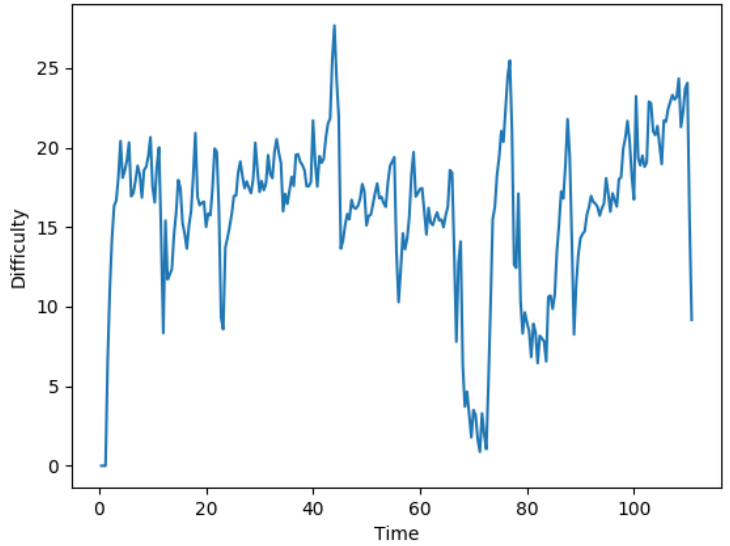

Note frequency
- Interval between consequent notes

Relationship between notes
- Simultaneous key presses
- Key press when holding down other keys

Highlighted control #2: Artistic Intents

I think including this piano solo will be interesting…

Since the genre of the music is Drum & Bass, I want to highlight drum and bass samples…

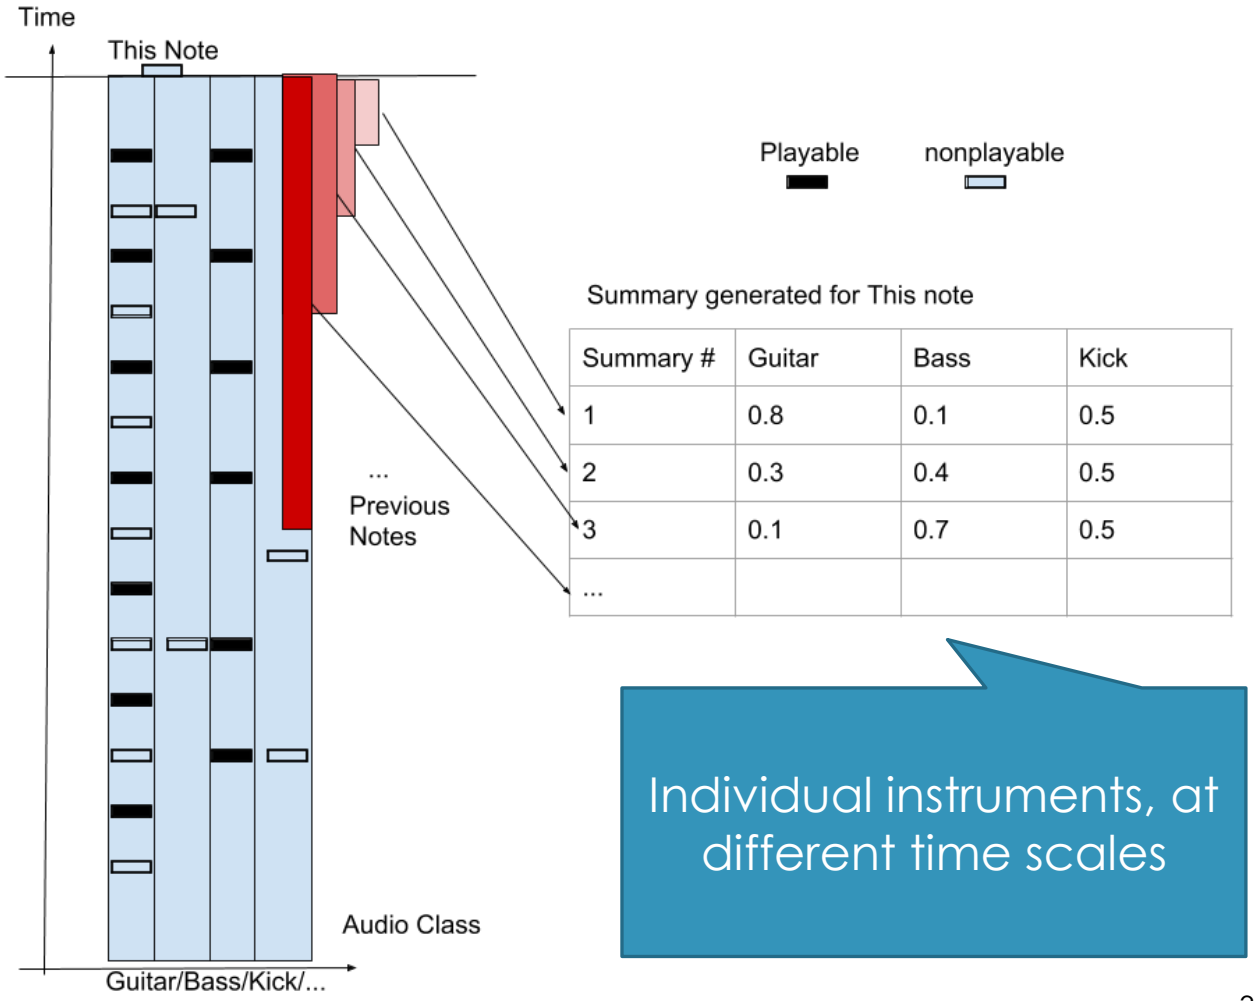

We built GenerationMania that learns how to use these intent-based controls

Model	F_1-score	Precision	Recall
Reference Baselines			
Random	0.291 ± 0.089	0.335 ± 0.200	0.299 ± 0.020
All Playable	0.472 ± 0.207	0.335 ± 0.199	1.000 ± 0.000
LSTM Baselines Models			
LSTM + Audio Features + BP + CM	0.424 ± 0.154	0.767 ± 0.176	0.353 ± 0.248
LSTM + Audio Features + BP + CM + RS	0.564 ± 0.149	$\mathbf{0.776 \pm 0.117}$	0.475 ± 0.194
Feed-Forward Models			
FF + Audio Features + BP + CM	0.253 ± 0.143	0.523 ± 0.266	0.179 ± 0.113
FF + Audio Features + BP + CM (Self Summary)	0.368 ± 0.198	0.422 ± 0.213	0.392 ± 0.258
FF + Audio Features + BP + RS	0.621 ± 0.206	0.760 ± 0.110	0.568 ± 0.254
FF + Audio Features + BP + CM + RS	$\mathbf{0.700 \pm 0.158}$	0.762 ± 0.114	$\mathbf{0.662 \pm 0.193}$

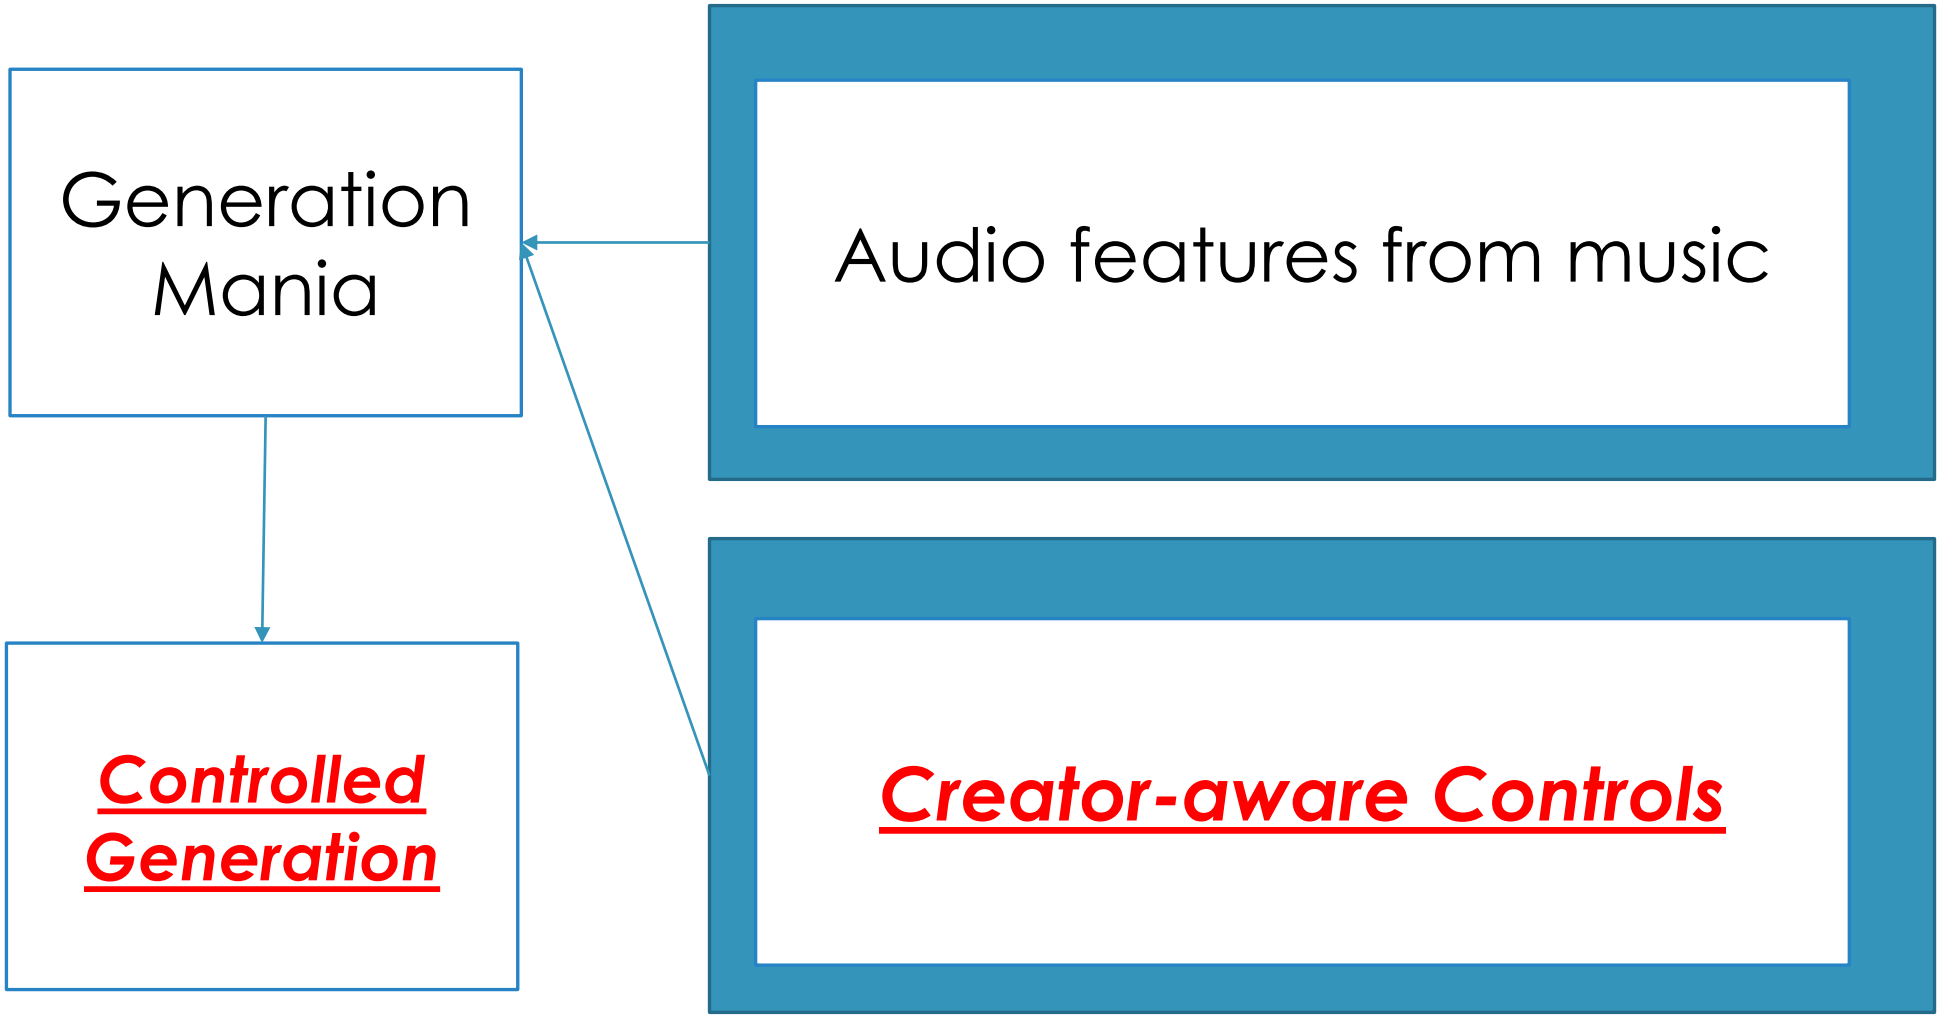
Generation Mania
Audio features from music
Controlled Generation
Creator-aware Controls

What did we learn?

- We can build a model that generates a stage from **high-level artistic intents**
- This enables **injection of intent from the creators**, while still **adhering to inherent constraints.**

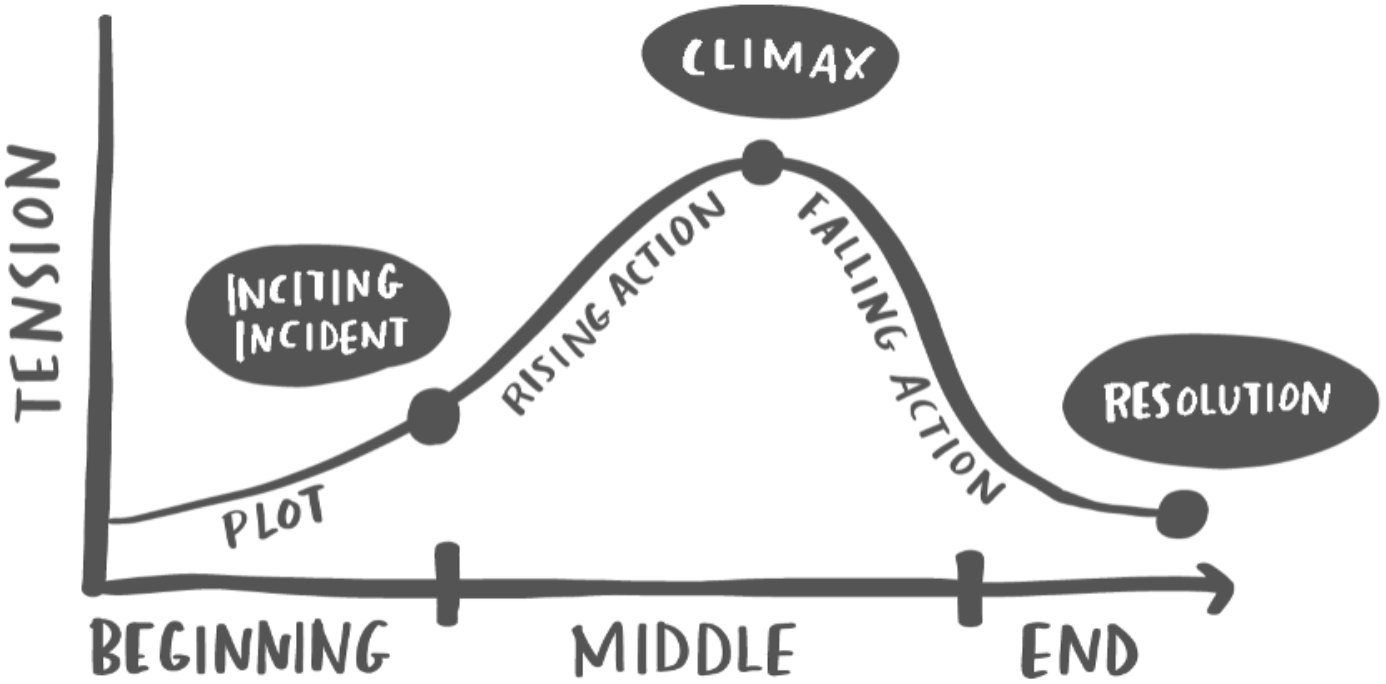

Let's apply this to

STORY GENERATION

Lin, Z., & Riedl, M. O. (2021). Plug-and-Blend: A Framework for Plug-and-Play Controllable Story Generation with Sketches. AIIDE 2021.

Domain 2: Story Generation

Suppose we want to turn this short story…

It is a sunny day.	We will see clear skies today.	Let's go to the park because the weather is good.

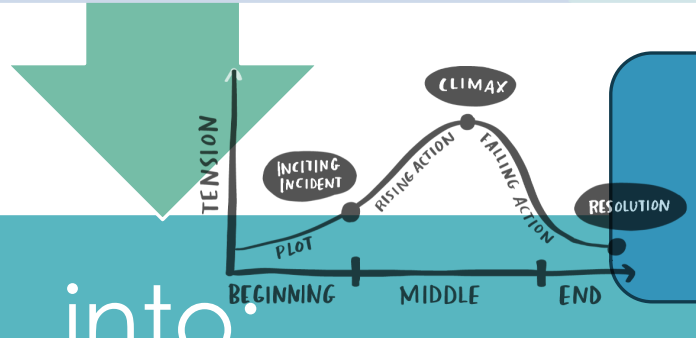

"I want to introduce some Science gradually…"

into:

It is a sunny day.	**I went to the classroom.**	**I learned that nothing travel faster than light.**

Can we...

Generate stories with these creator-aware controls?

While adapting to ever changing internal of the models?

Solution: Plug-and-Blend

Works with **the logit outputs of any continuation language model** *even without instruction following capabilities*

Facilitates fine control of generated sentences by allowing **continuous steering towards specific control codes**

Enables **high level area-of-control with a planner** that allows multiple control codes representing style and topic constraints on overlapping contexts.

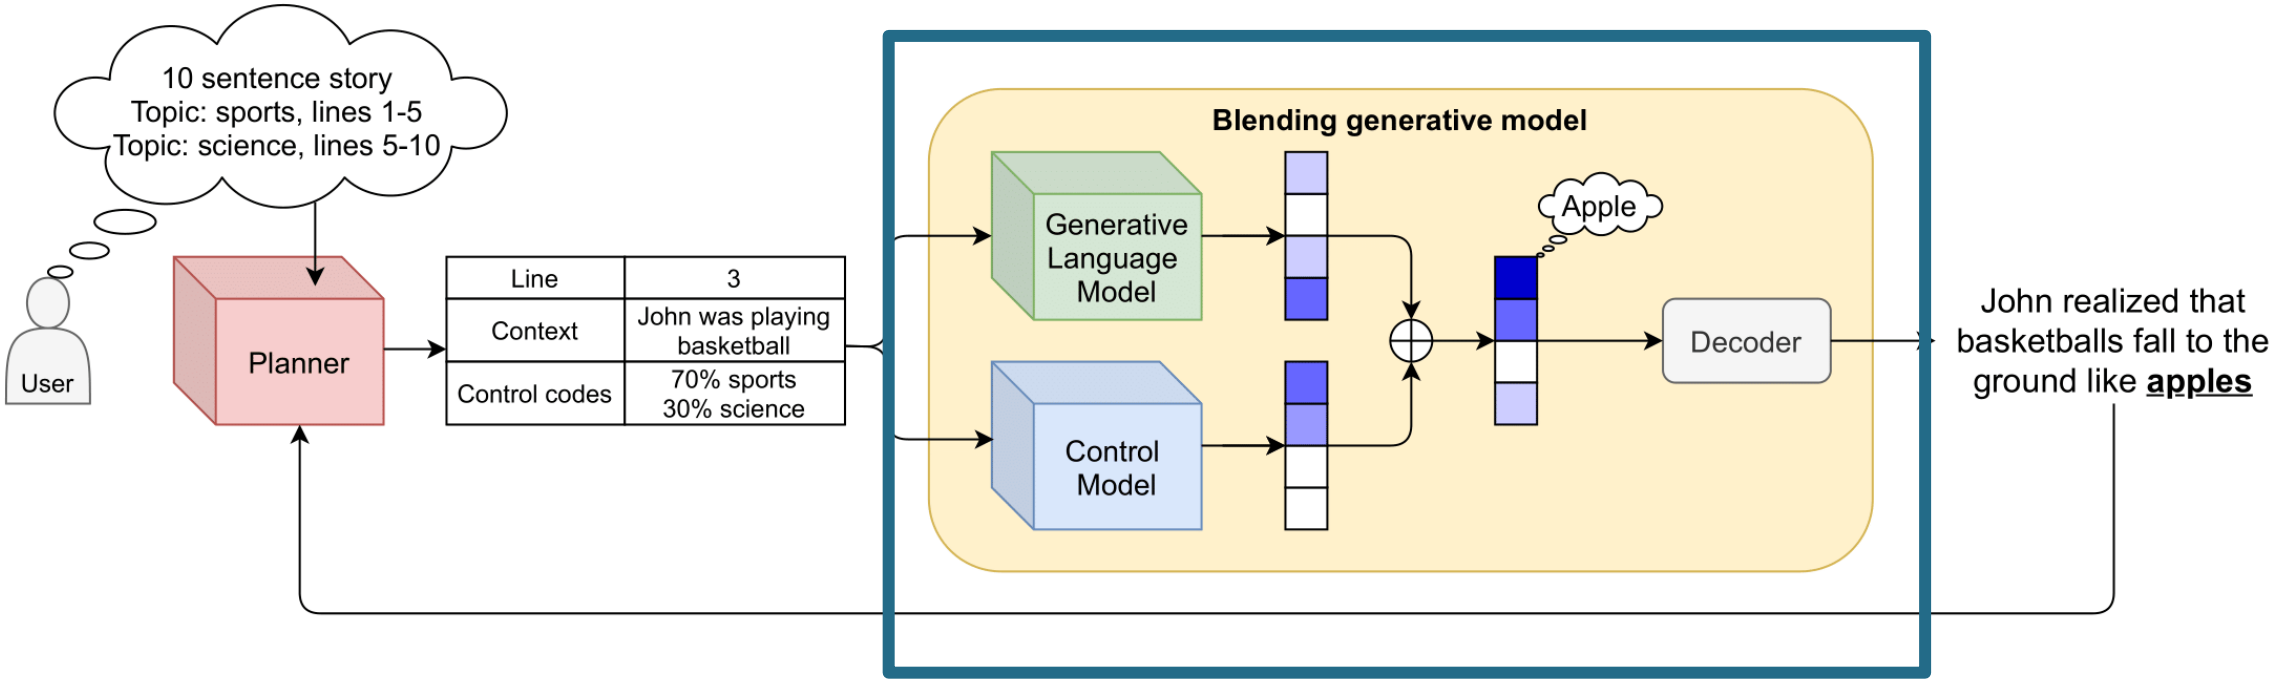

Fine-grained topic control

Language Model Preliminaries

Language models complete a sentence by picking the most likely token that finishes it.

$$P_\theta\left(x_{1:T}\right) = \prod_{t=1}^{T} P_\theta\left(x_t \mid x_{<t}\right)$$

What if we want to *control*?

We can control the generation by conditioning the process on a specific control code (attribute of topic, sentiment, etc...)

Adding new control codes naively requires a fine-tuning.

Hard for bigger, ever-changing and closed source models!

$$P_\theta\left(x_{1:T} \mid c\right) = \prod_{t=1}^{T} P_\theta\left(x_t \mid x_{<t}, c\right)$$

Control signal

The Blending Generation Model

- We can train an additional classifier (Krause et al., 2020) on a specific control attribute

Does the sentence fit control c?

$$P_\theta\left(c \mid x_{1:t}\right) = \alpha P(c) \prod_{j=1}^{t} P_\theta\left(x_j \mid x_{<j}, c\right)$$

- By contrasting it against an anti-attribute (not-c):

$$\alpha = 1/\left(\sum_{c' \in \{c, \bar{c}\}} \prod_{j=1}^{t} P\left(c'\right) P_\theta\left(x_j \mid x_{<j}, c'\right)\right)$$

Blending in

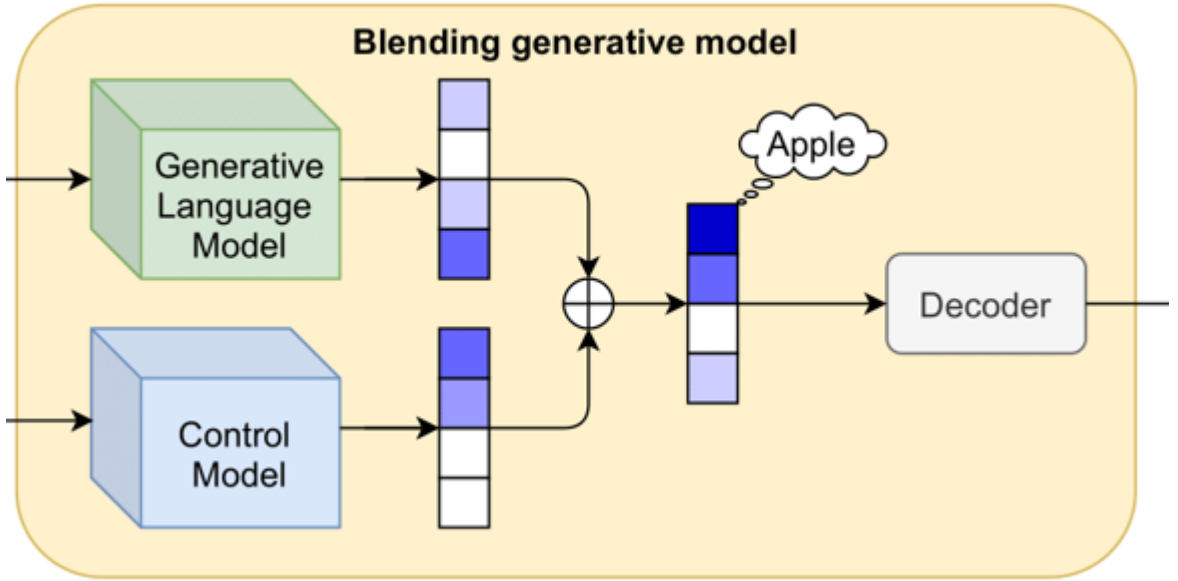

Control strength hyperparameter

$$P\left(x_t \mid x_{<t}, c\right) \propto P_{LM}\left(x_t \mid x_{<t}\right) P_\theta\left(c \mid x_t, x_{<t}\right)^\omega$$

Extending this to multi-control-code

This can be applied multiple times

$$P\left(x_t \mid x_{<t}, C\right) = P_{LM}\left(x_t \mid x_{<t}\right) \prod_{c^* \in C} P_\theta\left(c^* \mid x_t, x_{<t}\right)^{\omega_c^*}$$

See it in action

A different model for topic classification

Prompt: The people gathered to protest the court's ruling last week.

c_1 = Sport ω_{c_1}	c_2 = Business ω_{c_2}	BGM-Generated Sentence (Overall Strength 2x)	Classifier score c_1	Classifier score c_2
100%	0%	Coach Leeman was in a wheelchair and had been taken to hospital for treatment.	86%	14%
75%	25%	Coach Reebok was one of them.	65%	35%
50%	50%	The players were joined by a few of them.	84%	16%
25%	75%	The company that owns the team was fined $1,000 for violating a rule prohibiting employees from using their own equipment.	37%	63%
0%	100%	Bankruptcy Judge William H. said that the bank had failed to pay its creditors and was in default on $1 billion of loans it owed them.	24%	76%

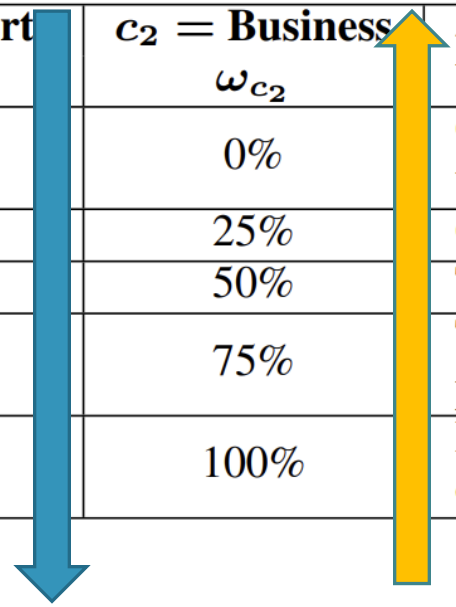

Less Sports, More Business

- Generated sentence becomes more related to a topic when requested with higher weights on that topic.

Local controls to global control

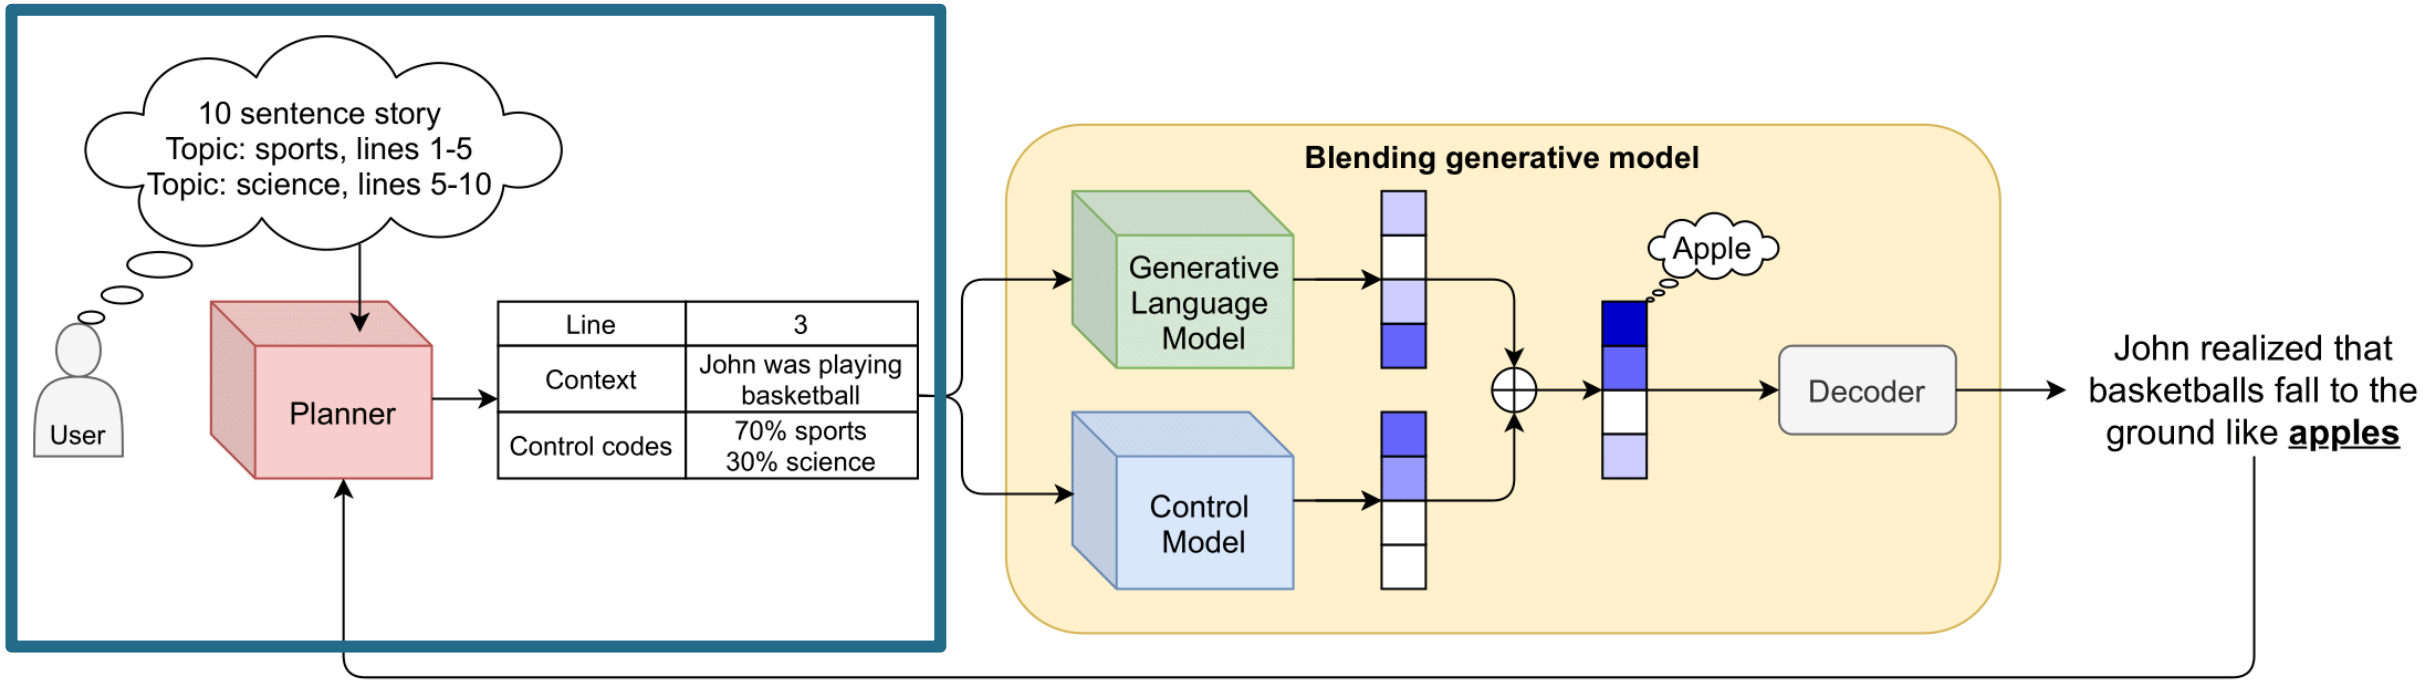

Capturing creator intents with "Sketches"

Inspired by Story Arcs

Translates to a continuous topic control to an area in the story

#	Sketch 1 On Nature	Sketch 2 On Tech
1		
2		
3		

Planner

#	Sketch 1 On Nature	Sketch 2 On Tech
1	0.9	0.1
2	0.5	0.5
3	0.1	0.9

Generated sentences

It is a sunny day.

I went to the classroom.

I learned that nothing travel faster than light.

Control Sketches	Generated Paragraph
Sketch 1 c_1 = Sports, $\{0 \rightarrow 5\}$ c_2 = Science, $\{\mathbf{4} \rightarrow 10\}$	Jackie Robinson was playing in the NBA. He had just been drafted by the Cleveland Cavaliers. He decided to play for a team that was trying to win a championship. "I'm gonna go out there and beat them," he said "They're going to be great." "Yeah, they're good," he said "they'll be great." He added that his wife had been working on the project for about a year. He was happy to see her progress. She was able to get some work done and she started getting more comfortable with it. She had a lot of fun doing it! Sh[illegible]go got her own comput[illegible]which is now in the [illegible]ving ro[illegible]m
Sketch 2 c_1 = Sports, $0 \rightarrow 5$ c_2 = Science, $\mathbf{5} \rightarrow 10$	Jack[illegible]ot[illegible]nso[illegible]as[illegible]l[illegible]n[illegible]i[illegible]t[illegible]BA He had just played basketball for a few years. He was looking forward to his first game of the season. He decided to play with his friends and play against them in the court. He got a lot of feedback from everyone who played against him, including some that were very excited about it! I was really happy when I saw how he played. I also had to admit that my favorite player was the guy who beat me in the finals. The computer game Super Mario 64 is a great game, but it's not perfect. I played it on my laptop and found that I couldn't play it properly because of some bugs. The problem was that the graphics were bad, so I had to use an emulator instead of playing the game.

Earlier Transition following control

Human Evaluations

We invited 45 participants to randomly read 30 pairs of stories and say if they see a transition between topics.

75.1% ($p<0.01$) chose story generated with P&B sketches.

This demonstrates **strong ability of our system injecting perceivable transitions**, blending in two topics when provided sketches that instruct so.

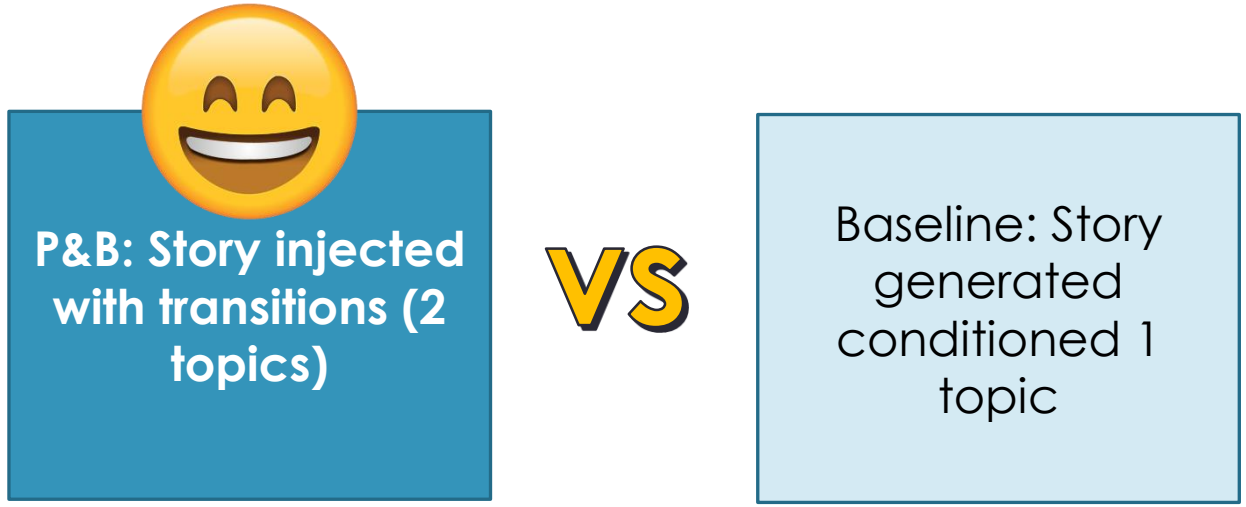

What did we learn?

- Inspired by the Narrative Arc, a tool used by storytellers, we can build a model that understands blending control, while creating a plan to apply control signals globally and locally.

What's next?

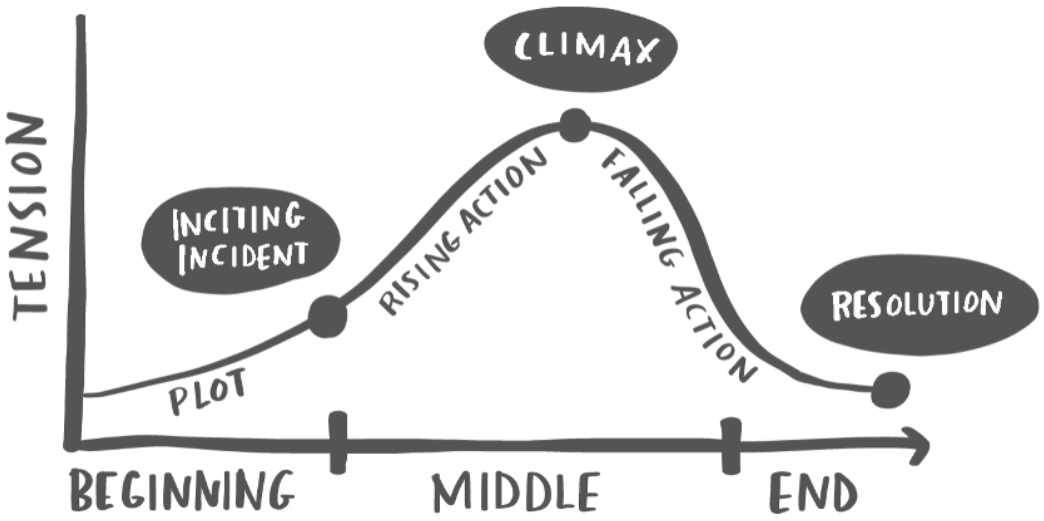

Generate with hierarchy?

Other creator-aware controls?

Player awareness?

"Why did you wrote this?"

How would you build creator-aware controls?

Expanding co-creative paradigms

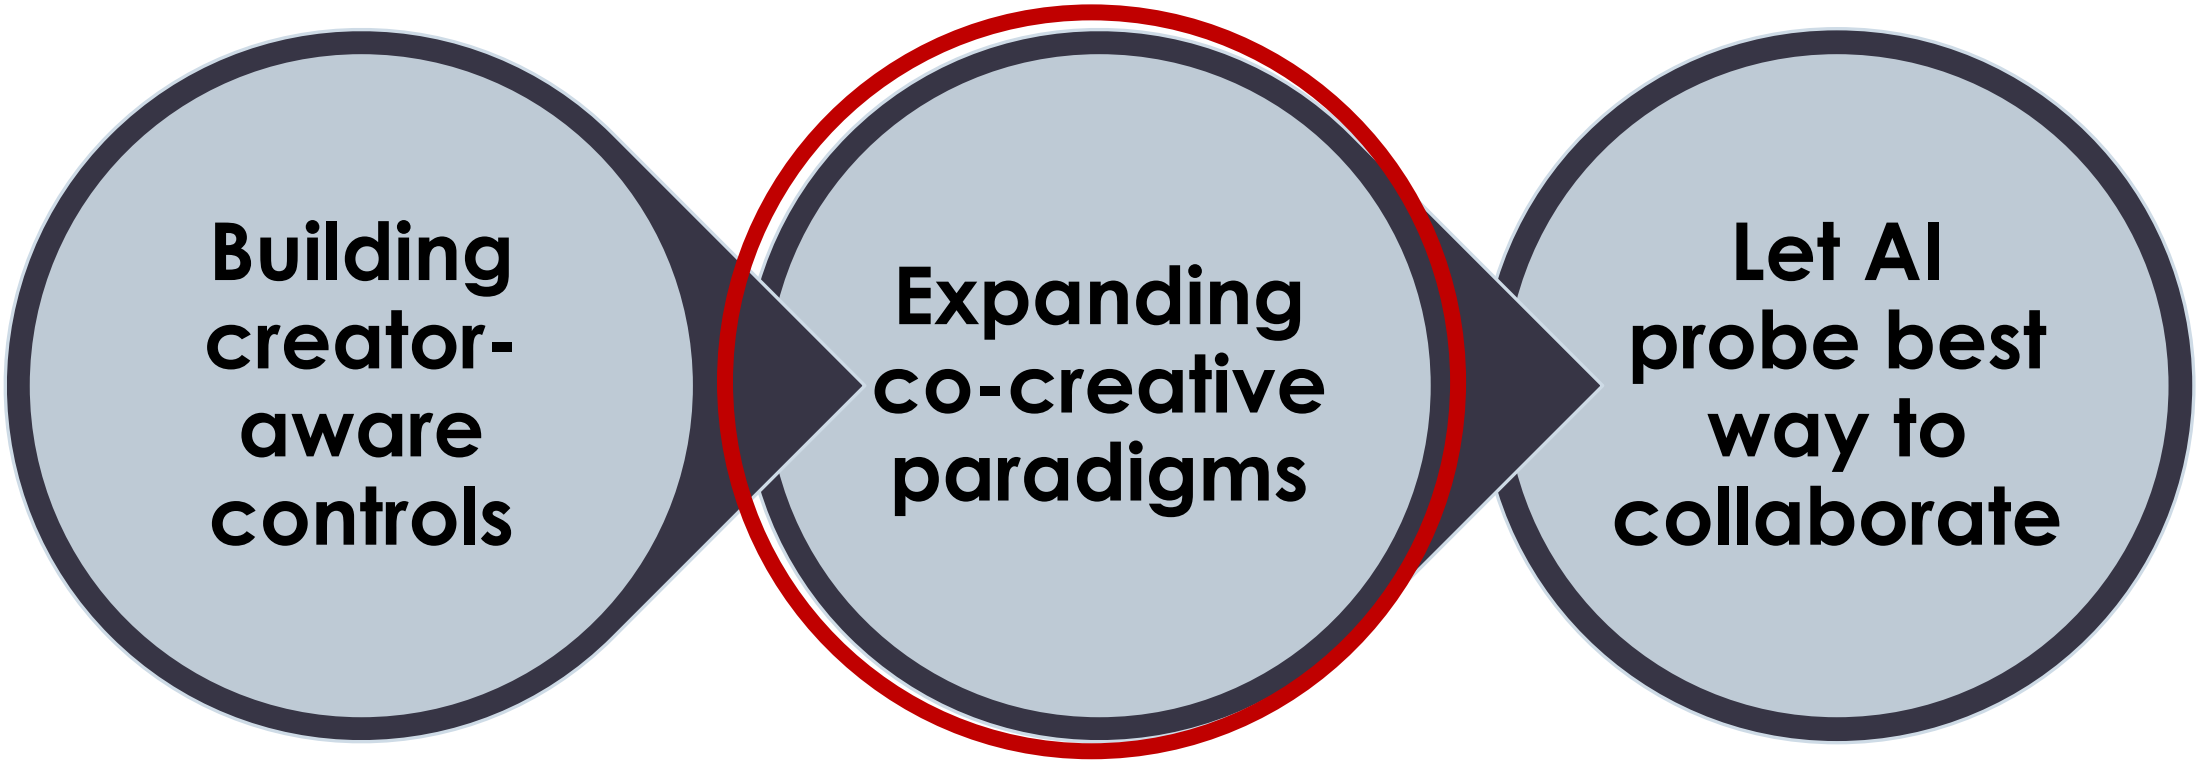

Revisit: What goes beyond "More Inputs"?

Users may not have needed communication expertise	• We are asking the creators for a different skillset than creating contents!
Users may not know what the final product looks like	• We are asking a concise prompt which may never be possible!
AI may not have a working understanding of users	• AI are interacting with the user without enough information!
There may never be a "best action"!	• There are more than one way to improve the contents!

Revisit: Mixed-Initiative Co-Creativity

Mixed-Initiative	Human initiative and a computational initiative cooperate towards a shared goal
Co-Creative	System having the capability to modify the contents as if it's the human counterpart

Let's take a deeper look into Co-creativity

Co-Creative AI having the capability to modify the contents like their human counterpart

- Researchers tag a variety of interactive setups as co-*creative*
- What does co-creativity actually mean for them?
- **Let's explore the space of co-creativity with an ontology!**

Lin Z, Riedl M. An Ontology of Co-Creative AI Systems
NeurIPS Workshop on Machine Learning for Creativity and Design 2023

How should responsibilities be distributed?

How should human creators and creative AI ***communicate*** with each other?

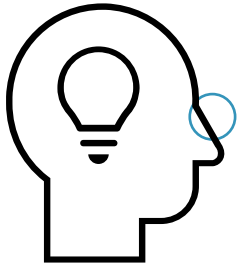

From computer creative support (Non-AI except Computer-as-colleague)…

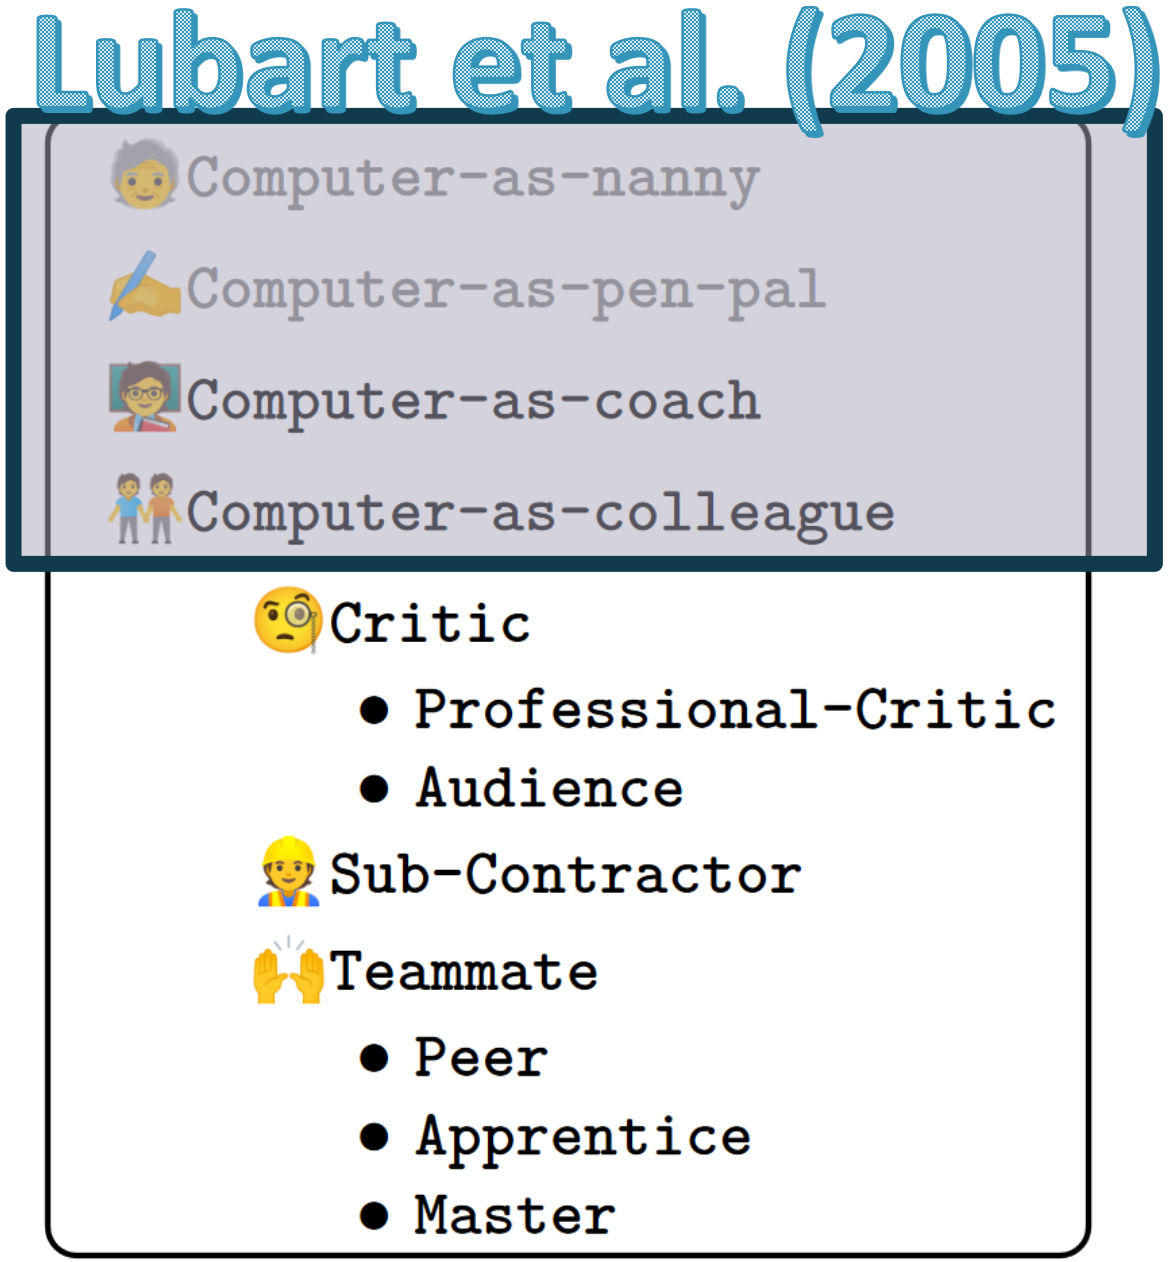

To co-creativity

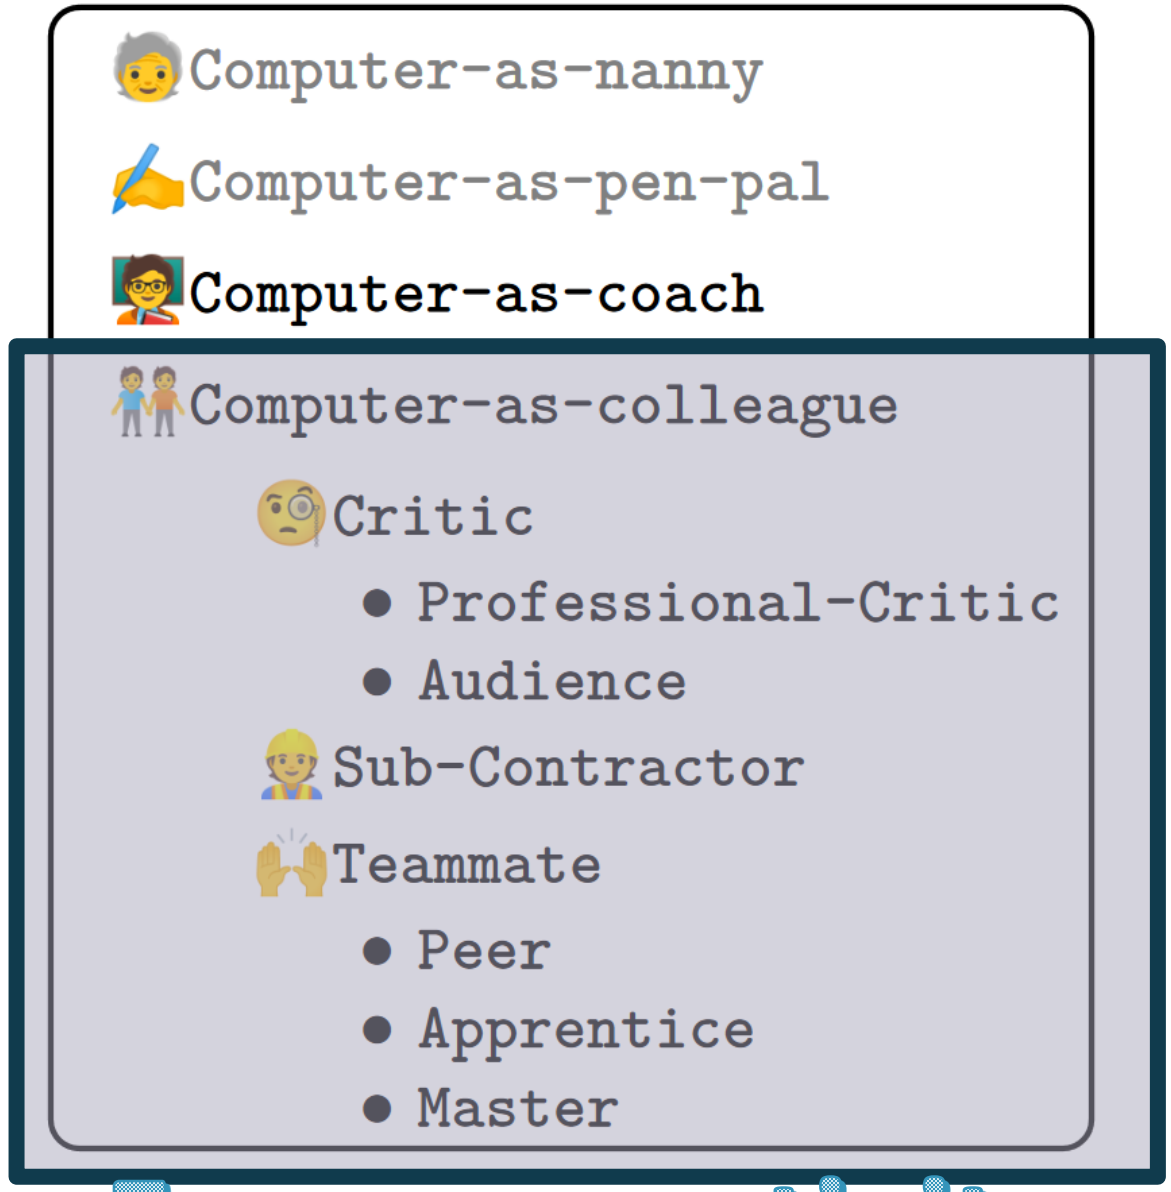

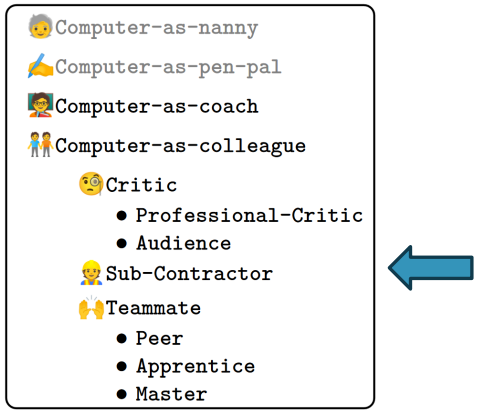

Computer-as subcontractor

- “This one is on me…”
- Human do not interfere with the generative process until the artifact is generated.
- Strict delineation of responsibilities.

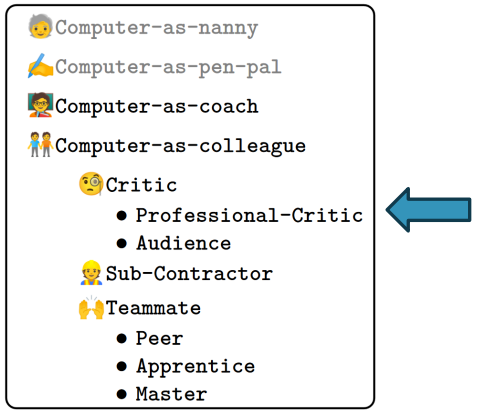

Conputer-as-critic

- “I think this looks…”
- The human ***retains all responsibilities*** for altering the creative artifact in response to advice.
- A **professional** critic gives feedback on established norms and conventions.
- An **audience** critic emulates people who will experience the final creative artifact.

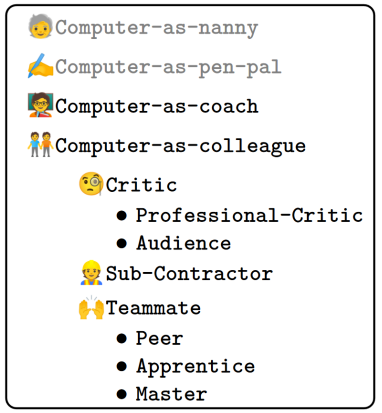

Computer-as-teammate

- “Let’s work this out together…”
- **Mixed-initiative in nature.**
- AI being an **apprentice, peer or master** based on its capability or human preference.
- When should the agent take initiative? How to adapt to, augment and extend the human creator?

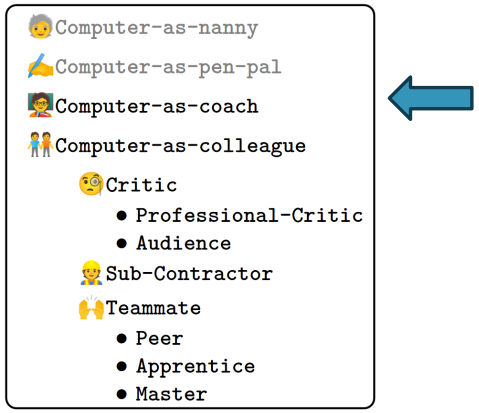

Computer-as-coach

- "Here is a guide…"
- AI teaches the human how to carry out a creative task, without working on the task
- AI can augment human's creative power, or provide actionable feedbacks (like a **master**)
- Although out of the "colleague" category, AI and non-AI opportunities awaits

How should responsibilities be distributed?

How should human creators and creative AI ***communicate*** with each other?

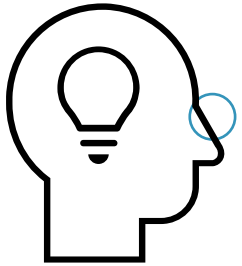

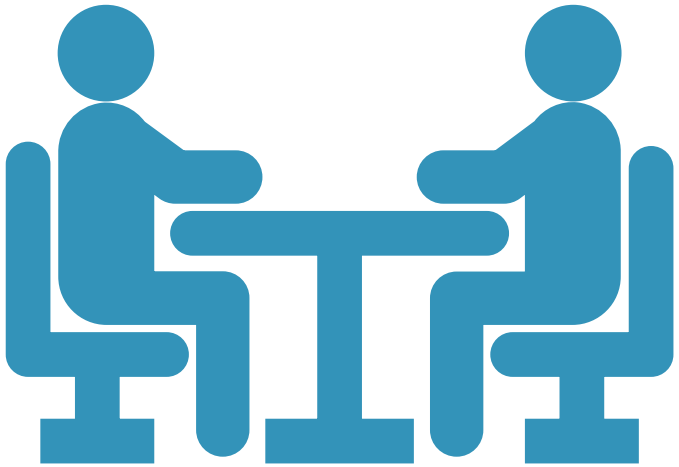

How should we effectively study the design space of co-creative communications between creators and AI?

Lin Z, Agarwal R, Riedl M. Creative wand: a system to study effects of communications in co-creative settings, AIIDE 2022

Creative Wand is a framework to study MI-CC Communications, with

Communications between the designer and the agent as first-class citizen	Modular design and API that helps connect **all kinds of creative domains**	Suggests **protocol and metrics** that helps researchers gain key insights

Components of Creative Wand

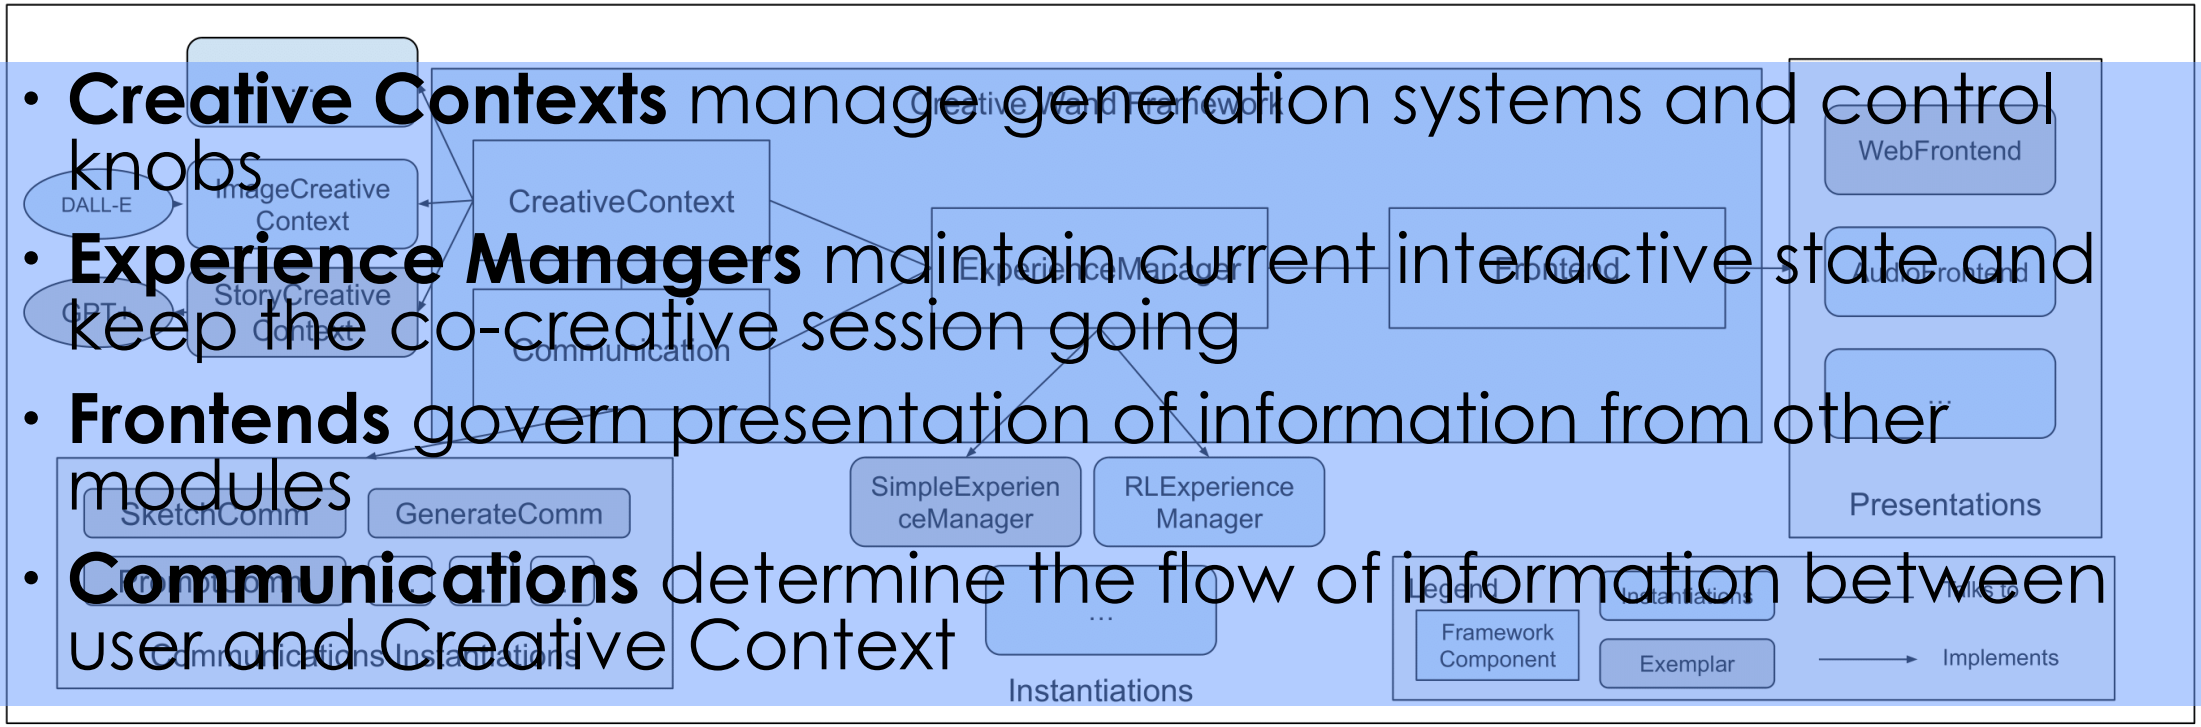

- **Creative Contexts** manage generation systems and control knobs
- **Experience Managers** maintain current interactive state and keep the co-creative session going
- **Frontends** govern presentation of information from other modules
- **Communications** determine the flow of information between user and Creative Context

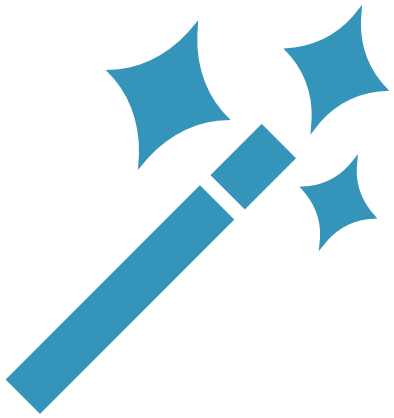

See an Exemplar Experiment

CREATIVE WAND IN ACTION

Exemplar Research Question

How will *Plug & Blend* enhance a MI-CC experience regarding co-creative goals and experience of human creators?

Exemplar Experiment Protocol

1. Configure systems with different MI-CC capabilities
2. Let participants use the systems
3. Collect feedbacks on the experience

Local communication: Substitute a sentence

#	Sentence
[0]	a girl was walking along the street and had to stop for a walk.
[1]	She stopped at a store to pick up some groceries.
[2]	She looked at the store clerk and said, "What's that?"
[3]	She said she was looking for a grocery store.
[4]	She went through the shelves and found a whole bunch of groceries.
[5]	She walked into the store and picked up a box of cereal.
[6]	She then walked out with a pack of carrots.
[7]	She picked them up and ate them.
[8]	She said she had never eaten carrots before.
[9]	She told them that she was going to the grocery store.

User rewrites sentence 5

[0]	a girl was walking along the street and had to stop for a walk.
[1]	She stopped at a store to pick up some groceries.
[2]	She looked at the store clerk and said, "What's that?"
[3]	She said she was looking for a grocery store.
[4]	She went through the shelves and found a whole bunch of groceries.
[5]	She decided to leave the grocery store.
[6]	Her parents said she should buy a few more.
[7]	They told her that she needed to get some food.
[8]	She went to the store and bought a bag of beans.
[9]	She got the beans and started eating them.

Changed

Regenerated continuation

Global communication: Plug and Blend

#	Sentence
[0]	a girl was walking along the street and had to stop for a walk.
[1]	She stopped at a store to pick up some groceries.
[2]	She looked at the store clerk and said, "What's that?"
[3]	She said she was looking for a grocery store.
[4]	She went through the shelves and found a whole bunch of groceries.
[5]	She walked into the store and picked up a box of cereal.
[6]	She then walked out with a pack of carrots.
[7]	She picked them up and ate them.
[8]	She said she had never eaten carrots before.
[9]	She told them that she was going to the grocery store.

Sentence	Sketch Topics
you had to buy a book to read about it.	
You bought it and then you read it again.	
Now you know how to use this book.	*'Business'*
This time, you can use it to teach your kids about the importance of reading.	*'Business'*
You can also teach them to read and write.	*'Business'*
You will need to make sure that you are reading the right way.	*'Business'*
You may want to learn how to use a pen and paper to write your book.	*'Business'*
In addition, you may need a little help with your reading.	*'Business'*
If you have trouble writing a book, check out our bookshelves.	*'Business'*
If you want to know how to write a good story, here are some tips.	

User wants a part of story to focus on Business

Whole story regenerated based on topic requested

Two conditions: Local and Global

Module	Instantiation	Description
CreativeContext	StoryCreativeContext	A backend interfacing with an implementation of Plug and Blend (Lin and Riedl 2021) with GPT-J(Wang and Komatsuzaki 2021) as the base language model, supporting both prompts and "sketch-based" high-level control.
ExperienceManager	SimpleExperienceManager	A turn and rule-based agent that shows all available Communications and allow the user to make a choice, or request for activation of Interrupted Communication when there is one.
Frontend	WebFrontend	A React.js web application with Chatbox interface and a canvas showing the artifact and additional information. (See figure 3)
Communications	Local condition: *UserWorkComm*, *GenerateWithFreezeComm*	Unique to this condition, allow the user to manually edit any line in the story, and "freeze" any line of the story
	Global condition: *UserSketchComm*	Unique to this condition, provides a global communication type allowing the user to set topics for a "sketch"(Lin and Riedl 2021) to influence part of the story. Also see Figure 2 for an example.

Table 1: CREATIVE-WAND module instantiations used in the experiment.

Help

Creative Wand Experiment System

End Session

Instructions and Your goals (click to reopen)

#	Sentence
[0]	
[1]	
[2]	
[3]	
[4]	
[5]	
[6]	
[7]	
[8]	
[9]	

And co-create a story

Sketch Topics

Human communicate with the agent with a chatbox interface…

Creative Wand

Let's work together!

Hello! I'm your Creative Wand.

Select a way to work with the Wand:
[1] Apply a topic to part of the story to guide the generation process.
[2] Let the Wand suggest a new topic to work on.
[3] Let the Wand (re)write the story, fill in the blanks.
[-1] We're done!
Take an action by either describing in natural language or replying with its number in []:

1 2 3

Type a message...

Exemplar experiment details

Your goal: Make a story with *creative wand* that start from talking about business and ends in something related to sports,mentioning soccer.

G#1 G#2 G#3

60 Participants are recruited from Prolific and asked to create a story **meeting the goal above (G#1,2 and 3)**

We ask participants to use both systems, and **collect feedback** for 15 communications

We ask participants whether they **completed the goals**, and level of **satisfaction and frustration**

Your goal: Make a story with *creative wand* that start from talking about business and ends in something related to sports,mentioning soccer.

G#1 G#2 G#3

Results

Reports of sub-goal Completion	#1	#2	#3
Local condition (n=28)	25.0%	17.9%	10.7%
Global condition (n=32)	40.6%	25.0%	15.6%
p-value ($H_0 : p_{global} \leq p_{local}$)	0.095	0.249	0.285

More participants reported goal completion on Global condition

Table 2: Metrics on rate of sub-goal reported as completed.

Of people completing goals, global condition made it faster, substantial for some

Interactions needed for sub-goal	#1	#2	#3
Local condition	8.71	9.40	9.33
Global condition	7.08	6.25	5.80
p-value ($H_0 : t_{global} \geq t_{local}$)	0.140	0.027	0.047

Table 3: Metrics on interactions taken to achieve sub-goals.

Your goal: Make a story with *creative wand* that start from talking about business and ends in something related to sports,mentioning soccer.

G#1 G#2 G#3

Results

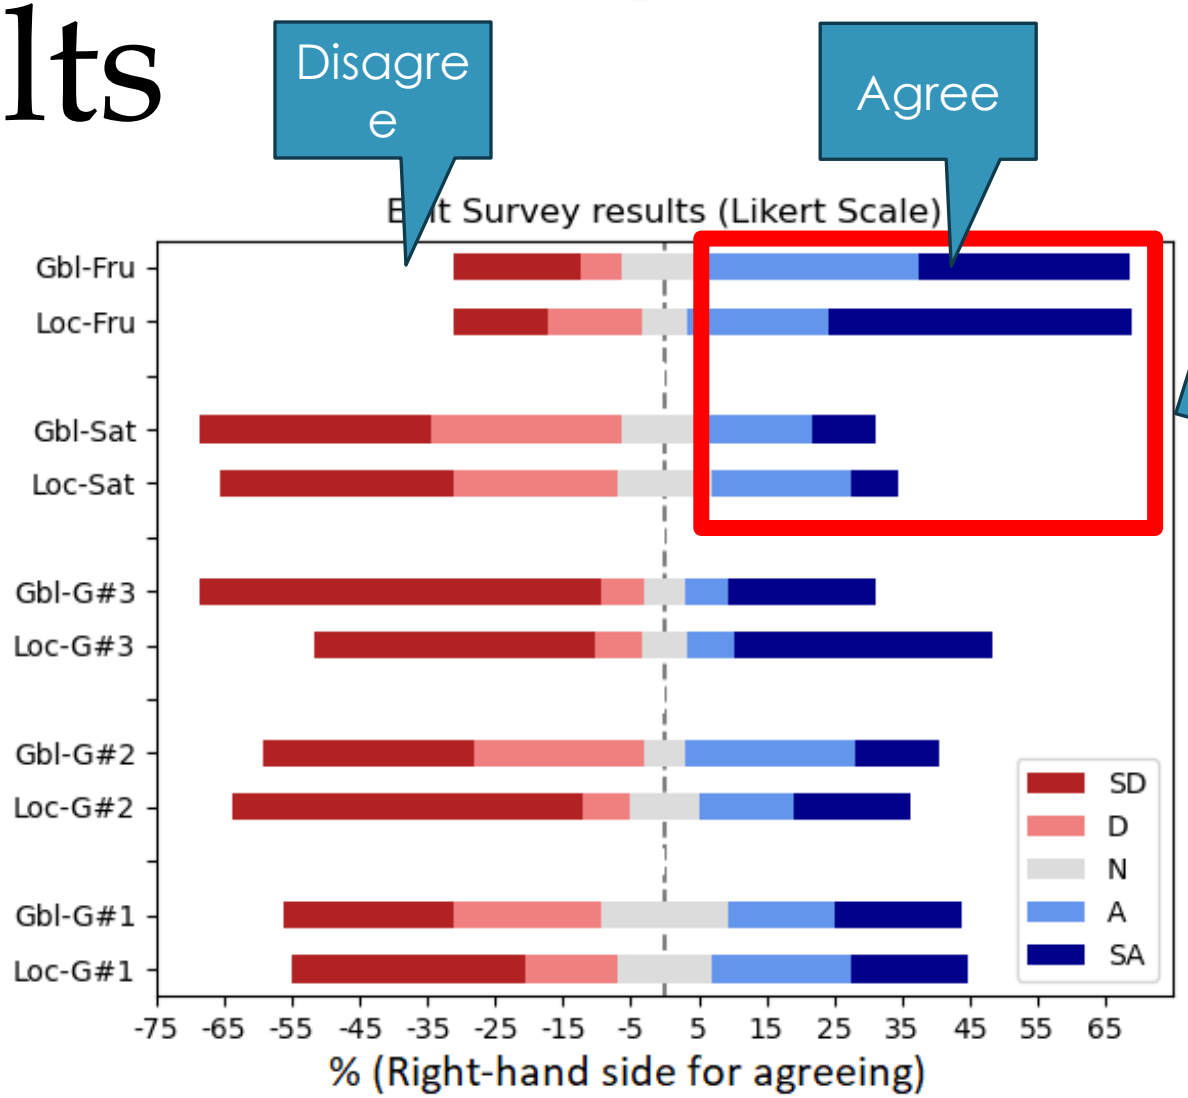

Participants are frustrated with these systems with only one type of communication available, wished that they have both.

A starting point

Although an *artificial comparison*, we can **easily expand the communications provided**.

We get **useful insights for MI-CC systems** from this exemplar experiment.

We've seen **Creative Wand and its flexibility in action.**

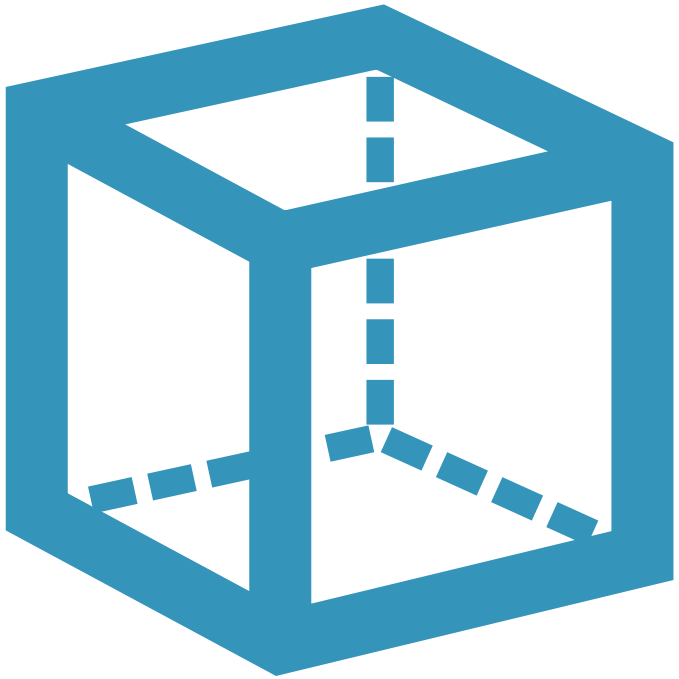

Let's **expand the design space of** communications!

Lin, Z., Ehsan, U., Agarwal, R., Dani, S., Vashishth, V., & Riedl, M. , Beyond Prompts: Exploring the Design Space of Mixed-Initiative Co-Creativity Systems, ICCC 2023

Ontology of Communications

Continuous and non-exhaustive

Human vs. Agent-initiated

- Which of the two parties initiated this communication?

Elaboration vs. Reflection

- Is the communication related to previous contents (reflection) or newly planned action(elaboration)?

Global vs. Local

- What is the scope of the contents that is targeted by the communication?

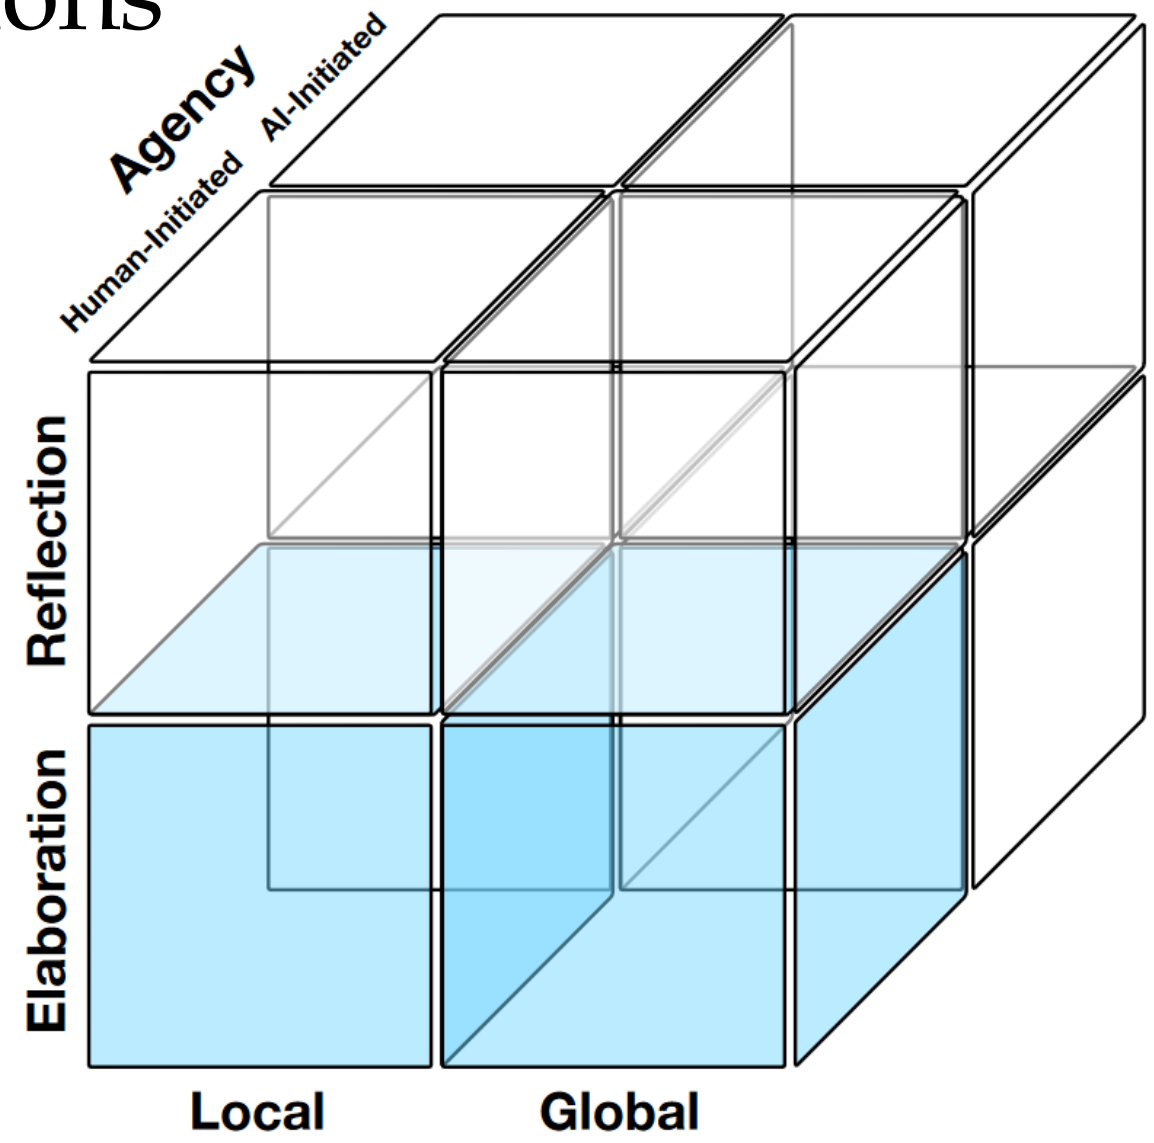

Communication examples fitting our ontology

(Human Initiated)	Elaboration	Reflection
Local Scope	Human write a sentence	H. provide why writing a sentence so that A. can verify
Global Scope	Human apply topic control	H. Ask a question about current story so that A. can answer

(Agent Initiated)	Elaboration	Reflection
Local Scope	Agent suggest a sentence	Agent point out that a sentence is out of topic
Global Scope	Agent suggest a topic	Agent point out that the whole story is a cliché

(Human Initiated)	Elaboration	Reflection
Local Scope	Human write a sentence	H. provide why writing a sentence so that A. can verify
Global Scope	Human apply topic control	H. Ask a question about current story so that A. can answer

We are ready to go beyond the exemplar experiment!

What do we explore?

How does the **availability of different means of communication** affect the **creative experience**?

Specifically, we ask…

How the **presence or absence of different modes of human-AI communications** affect perceptions of creative support?

What are the **individual variables**, such as creative background and familiarity with AI, that influences the above?

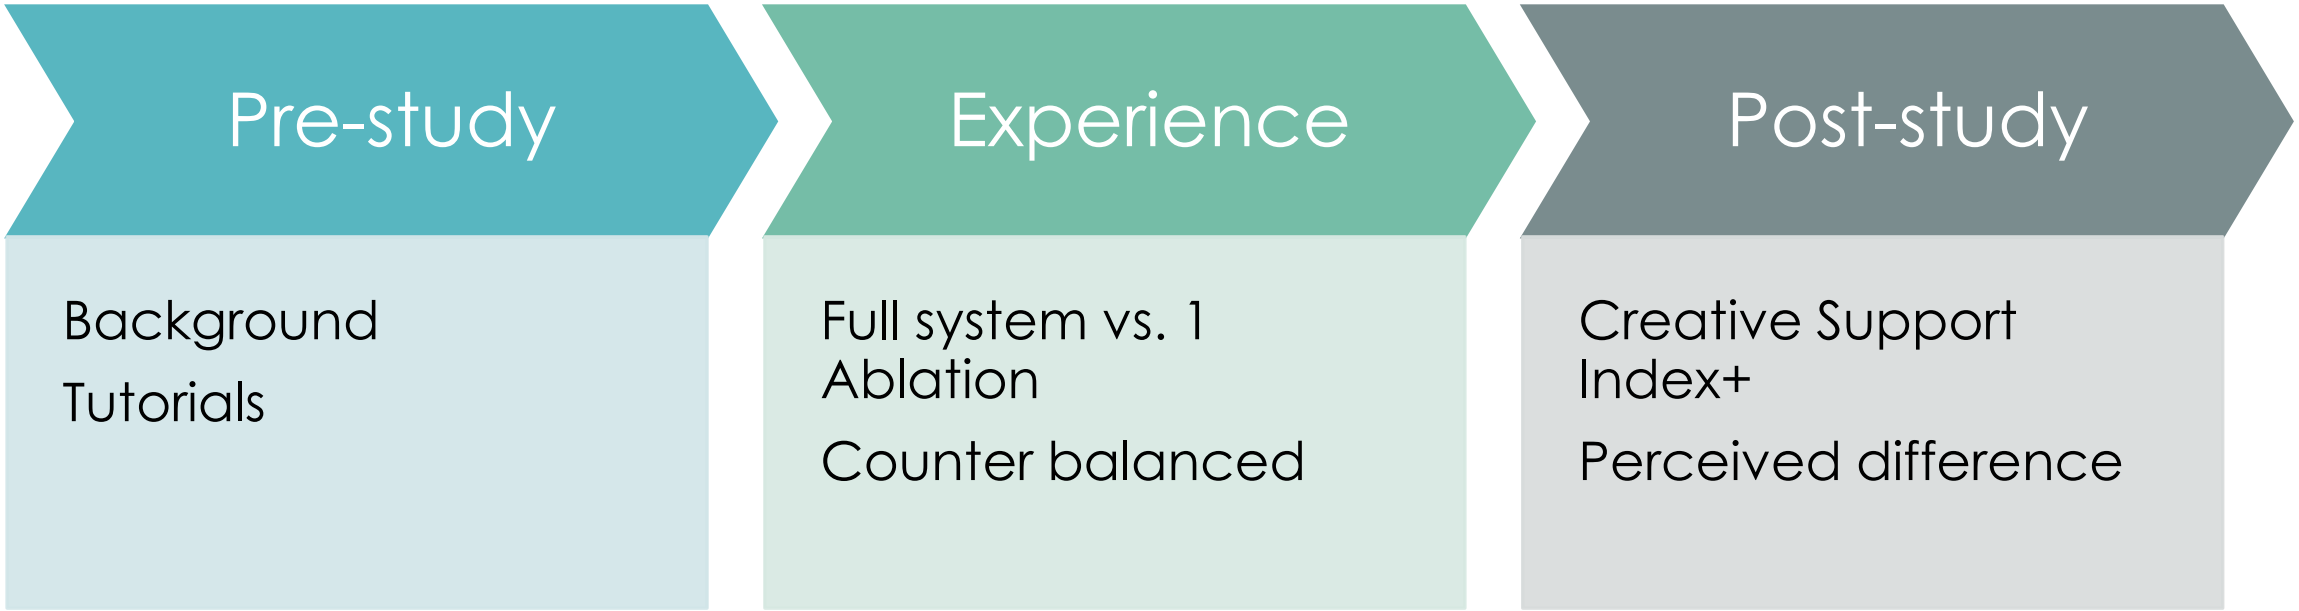
Pre-study
Background
Tutorials
Experience
Full system vs. 1 Ablation
Counter balanced
Post-study
Creative Support Index+
Perceived difference

The study

“Full System” vs. Ablations

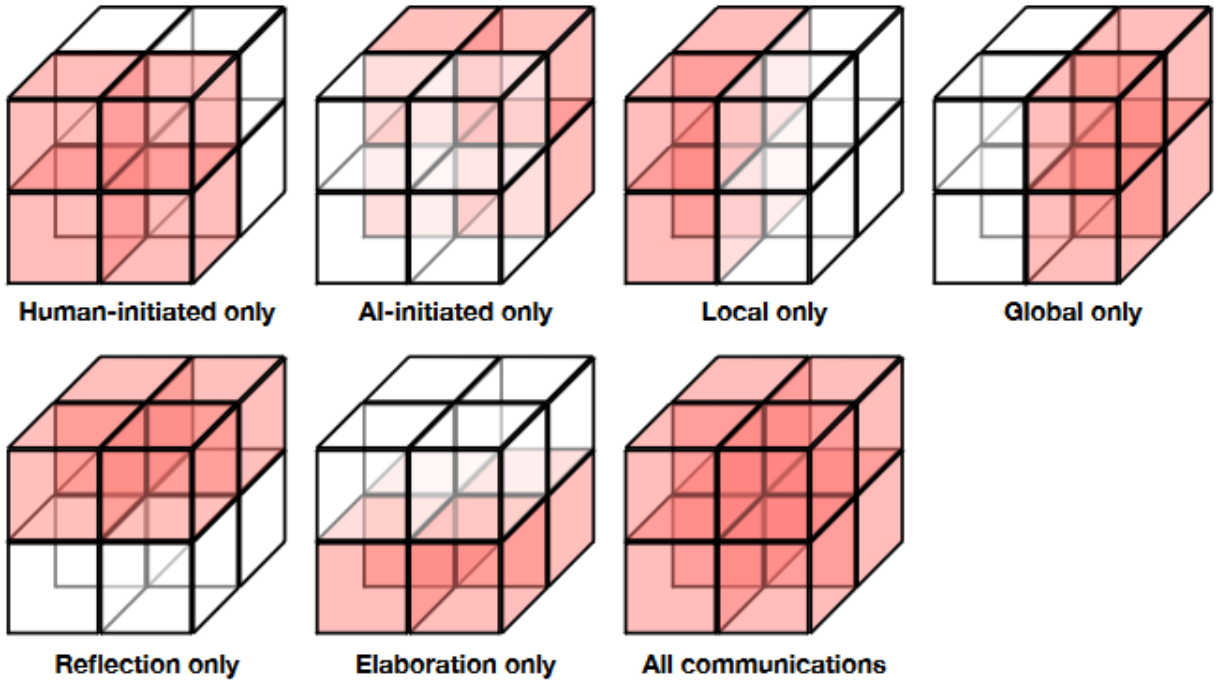

n

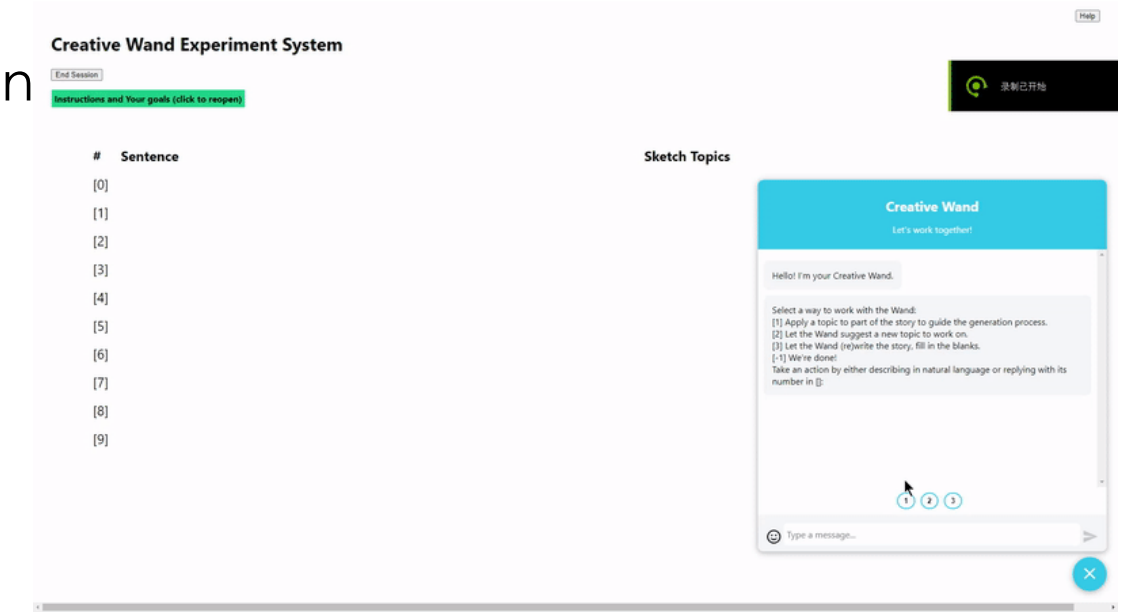

Example

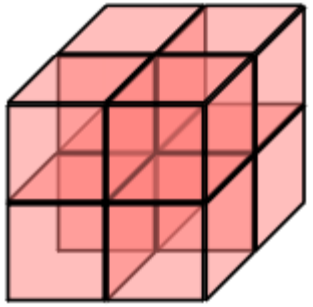

All communications

VS

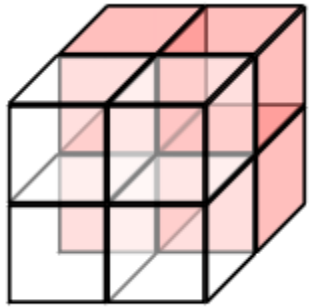

AI-initiated only

(Human Initiated)	Elaboration	Reflection
Local Scope	Human write a sentence	H. provide why writing a sentence so that A. can verify
Global Scope	Human apply topic control	H. Ask a question about current story so that A. can answer

(Agent Initiated)	Elaboration	Reflection
Local Scope	Agent suggest a sentence	Agent point out that a sentence is out of topic
Global Scope	Agent suggest a topic	Agent point out that the whole story is a cliché

Disabled

+Regenerate/Undo

The study

185 participants from Prolific

Pre-study questionnaire:

- Creative background
  - A - General computer-aided authoring (from 1 to 4)
  - B - Game content authoring (from 1 to 4)
  - C - Familiarity with AI (from 1 to 3)
- Tutorial for the experimental system

A. Level of confidence with using a computer to author contents:
   - A1. I do not use computers to create things.
   - A2. I have used computers to create things, but for the past year, I have not done it once a week.
   - A3. I use computers to create things more than once a week, but I'm doing it not for the job (for example, for interest).
   - A4. I use computers to create things for my job.

B. Level of confidence with using a computer to create games:
   - B1. I never used a computer to create anything related to games.
   - B2. I've done some work in the realm of games, but for the past year, I have not done it once a week.
   - B3. I create content for games out of interest, for more than once a week.
   - B4. I create content for games for my job.

C. Familiarity with AI:
   - C1. All I know is no more than it being a buzz word.
   - C2. I have experience using something with "AI technologies" with it.
   - C3. I understand how recent AI technologies work.

Post-study

- Based on Creative Support Index (Cherry, 2014)
- Free-text justifications on perceived difference.

Q1. (Expressiveness) This system made it easiest for me to express and share my goals, given to me in instructions, with the AI system.
Q2. (Enjoyment) I enjoyed interacting with this system most.
Q3. (Exploration) This system was most helpful for exploring different ideas and possibilities
Q4. (Immersion) This system made me feel the most absorbed in the task to the point that I forgot I was working with the system.
Q5. (Collaboration) This system best allowed me to achieve the goal assigned to me.
Q6. (Results worth effort) This system provides the overall best quality stories by the time I was done.
Q7. Which system tends to get the better response for the same type of requests?

Participants come with diversity in their creative backgrounds

	n	Q1	Q2	Q3	Q4	Q5	Q6
Level of		**onfidence with using a computer to author contents (A1**					
A1	3	*Too few participants*					
A2	44	61.4%	56.8%	54.5%	54.5%	61.4%	56.8%
A3	62	66.1%	67.7%	72.6%	67.7%	66.1%	69.4%
A4	76	60.5%	57.9%	60.5%	63.2%	53.9%	56.6%
Level o		**confidence with using a computer to create games (B1-**					
B1	94	63.8%	63.8%	60.6%	66.0%	64.9%	67.0%
B2	65	52.3%	50.8%	66.2%	52.3%	46.2%	49.2%
B3	21	81.0%	76.2%	61.9%	76.2%	71.4%	61.9%
B4	5	*Too few participants*					
		Familiarity with AI (C1-C3)					
C1	29	48.3%	55.2%	48.3%	51.7%	48.3%	58.6%
C2	107	66.4%	64.5%	67.3%	67.3%	61.7%	58.9%
C3	49	61.2%	55.1%	61.2%	57.1%	61.2%	65.3%

Table 2: Rate of preference on Full system, grouped by answers to demographic
than 20 participants are shown. * means different distribution with $p < 0.01$, + fc

A. Level of confidence with using a computer to author contents:
   - A1. I do not use computers to create things.
   - A2. I have used computers to create things, but for the past year, I have not done it once a week.
   - A3. I use computers to create things more than once a week, but I'm doing it not for the job (for example, for interest).
   - A4. I use computers to create things for my job.

B. Level of confidence with using a computer to create games:
   - B1. I never used a computer to create anything related to games.
   - B2. I've done some work in the realm of games, but for the past year, I have not done it once a week.
   - B3. I create content for games out of interest, for more than once a week.
   - B4. I create content for games for my job.

C. Familiarity with AI:
   - C1. All I know is no more than it being a buzz word.
   - C2. I have experience using something with "AI technologies" with it.
   - C3. I understand how recent AI technologies work.

Significant difference in preference between most of the groups we have enough data with

	n	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Welch's t-test		
Level of confidence with using a computer to author contents (A1-A4)									**A2**	**A3**	**A4**
A1	3	*Too few participants*									
A2	44	61.4%	56.8%	54.5%	54.5%	61.4%	56.8%	56.8%	N/A	*	
A3	62	66.1%	67.7%	72.6%	67.7%	66.1%	69.4%	71.0%	*	N/A	*
A4	76	60.5%	57.9%	60.5%	63.2%	53.9%	56.6%	57.9%		*	N/A
Level of confidence with using a computer to create games (B1-B4)									**B1**	**B2**	**B3**
B1	94	63.8%	63.8%	60.6%	66.0%	64.9%	67.0%	66.0%	N/A	*	
B2	65	52.3%	50.8%	66.2%	52.3%	46.2%	49.2%	53.8%	*	N/A	*
B3	21	81.0%	76.2%	61.9%	76.2%	71.4%	61.9%	61.9%		*	N/A
B4	5	*Too few participants*									
Familiarity with AI (C1-C3)									**C1**	**C2**	**C3**
C1	29	48.3%	55.2%	48.3%	51.7%	48.3%	58.6%	58.6%	N/A	*	+
C2	107	66.4%	64.5%	67.3%	67.3%	61.7%	58.9%	63.6%	*	N/A	
C3	49	61.2%	55.1%	61.2%	57.1%	61.2%	65.3%	59.2%	+		N/A

Table 2: Rate of preference on Full system, grouped by answers to demographics questions. Only data for groups with more than 20 participants are shown. * means different distribution with $p < 0.01$, + for $p < 0.1$.

What did we learn?

MI-CC tools should be customized to different types of users with different levels of creative expertise and AI familiarity.

Participants preferred significantly more a system having more Communications.

	Overall	Agent-Init. Only	Human-Init. Only	Elaboration Only	Reflection Only	Global Only	Local Only
Num. valid responses	185	31	32	30	32	27	33
Q1: Expressiveness	62.2%*	74.2%*	46.9%	56.7%	78.1%*	63.0%	54.5%
Q2: Enjoyment	60.5%*	74.2%*	43.8%	50.0%	81.2%*	59.3%	54.5%
Q3: Exploration	62.7%*	71.0%*	46.9%	56.7%	71.9%*	70.4%*	60.6%
Q4: Immersion	62.2%*	71.0%*	50.0%	60.0%	75.0%*	59.3%	57.6%
Q5: Collaboration	59.5%*	71.0%*	40.6%	56.7%	81.2%*	59.3%	48.5%
Q6: Result worth effort	60.5%*	64.5%+	53.1%	60.0%	71.9%*	66.7%+	48.5%
Q7: Better responses	61.6%*	67.7%*	56.2%	63.3%	78.1%*	59.3%	45.5%

Table 1: Rate of participants that preferred the Full System over the ablations. * represents a significance level of $p < 0.05$ on Full system preferred over the ablation; + for $p < 0.1$. No ablation was preferred statistically significantly.

They also think the system gives better responses, despite the underlying implementation being the same.

	Overall	Agent-Init. Only	Human-Init. Only	Elaboration Only	Reflection Only	Global Only	Local Only
Num. valid responses	185	31	32	30	32	27	33
Q1: Expressiveness	62.2%*	74.2%*	46.9%	56.7%	78.1%*	63.0%	54.5%
Q2: Enjoyment	60.5%*	74.2%*	43.8%	50.0%	81.2%*	59.3%	54.5%
Q3: Exploration	62.7%*	71.0%*	46.9%	56.7%	71.9%*	70.4%*	60.6%
Q4: Immersion	62.2%*	71.0%*	50.0%	60.0%	75.0%*	59.3%	57.6%
Q5: Collaboration	59.5%*	71.0%*	40.6%	56.7%	81.2%*	59.3%	48.5%
Q6: Result worth effort	60.5%*	64.5%+	53.1%	60.0%	71.9%*	66.7%+	48.5%
Q7: Better responses	61.6%*	67.7%*	56.2%	63.3%	78.1%*	59.3%	45.5%

Table 1: Rate of participants that preferred the Full System over the ablations. * represents a significance level of $p < 0.05$ on Full system preferred over the ablation; + for $p < 0.1$. No ablation was preferred statistically significantly.

What did we learn?

We should consider building MI-CC systems with a wider coverage of the design space of user-AI communication types.

Creative Support Index questions are correlated

Overall correlation

Stronger correlation between

- Q1 (expressiveness) and Q2 (enjoyment) and Q5 (collaboration)
- Q2 (enjoyment) and Q4 (immersion)

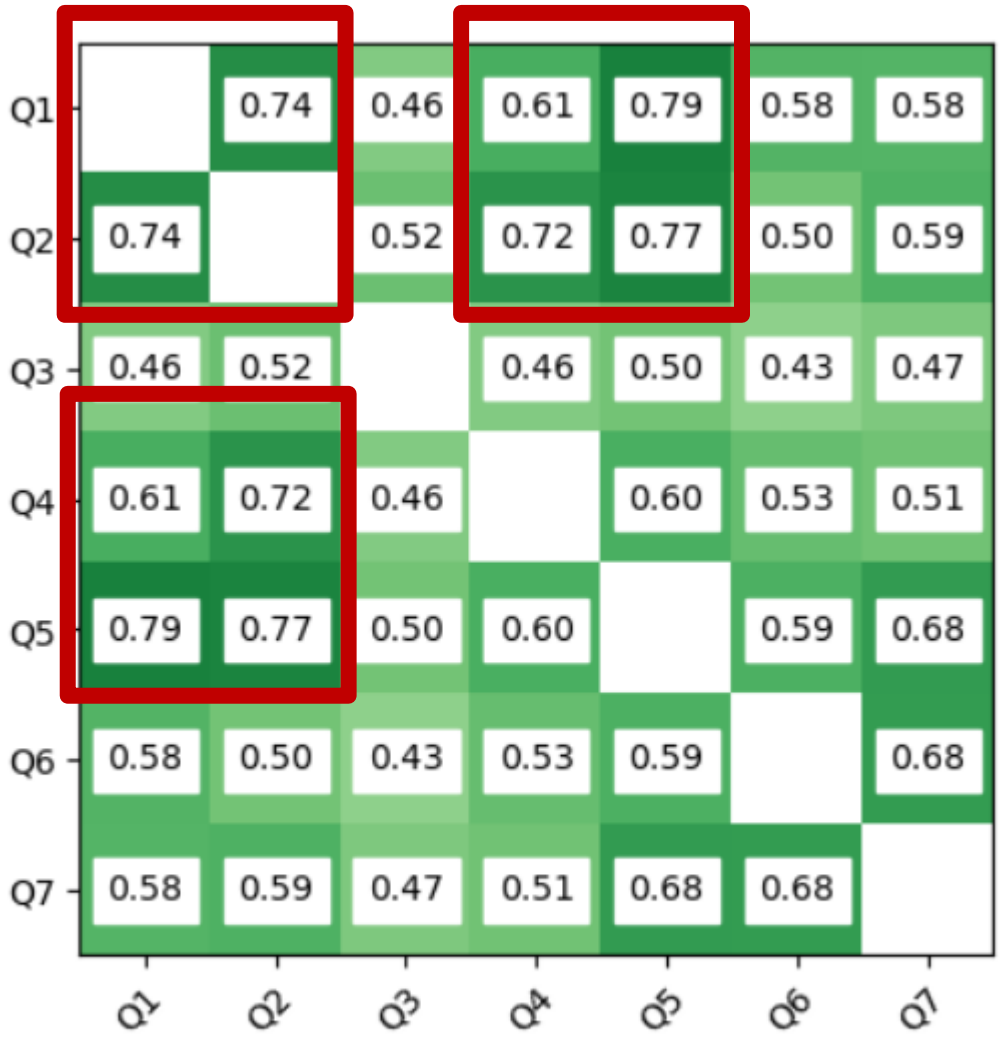

Figure 4: Correlation between questions in the survey.

What did we learn?

Focusing on enjoyable MI-CC experience can make them seem immersive, expressive, and collaborative too.

Salient voices from participants

Fish[4] is superior to Rabbit[5] in that it you can guide and interact with it and *it listens to feedback* and doesn't just write what it wants. Fish allowed you more control in guiding the story on topics before starting so it was more accurate and also more customizable. Rabbit felt more random with less options and control, it

Ability to exercise control

Ease of controlling, customizability

I had an easier time understanding the Fish system. And it appeared to understand the topics better based on my interaction. (P76)

Scrutability

Mechanical, functional and two-way understanding

Personalized Explainability is favored

Takeaways

We conducted a study with 185 participants and 7 MI-CC systems probing plausible subset of a design space. We found out that:

User-based customization is needed	Wider coverage is good	Personalized Explainability is favored

- And most importantly, there are much to explore **beyond prompts**!

What's next?

More axes for co-creative ontology?

Autonomous discovery of ways to work with creators?

How to make role hybrids?

How to extend beyond computational creativity?

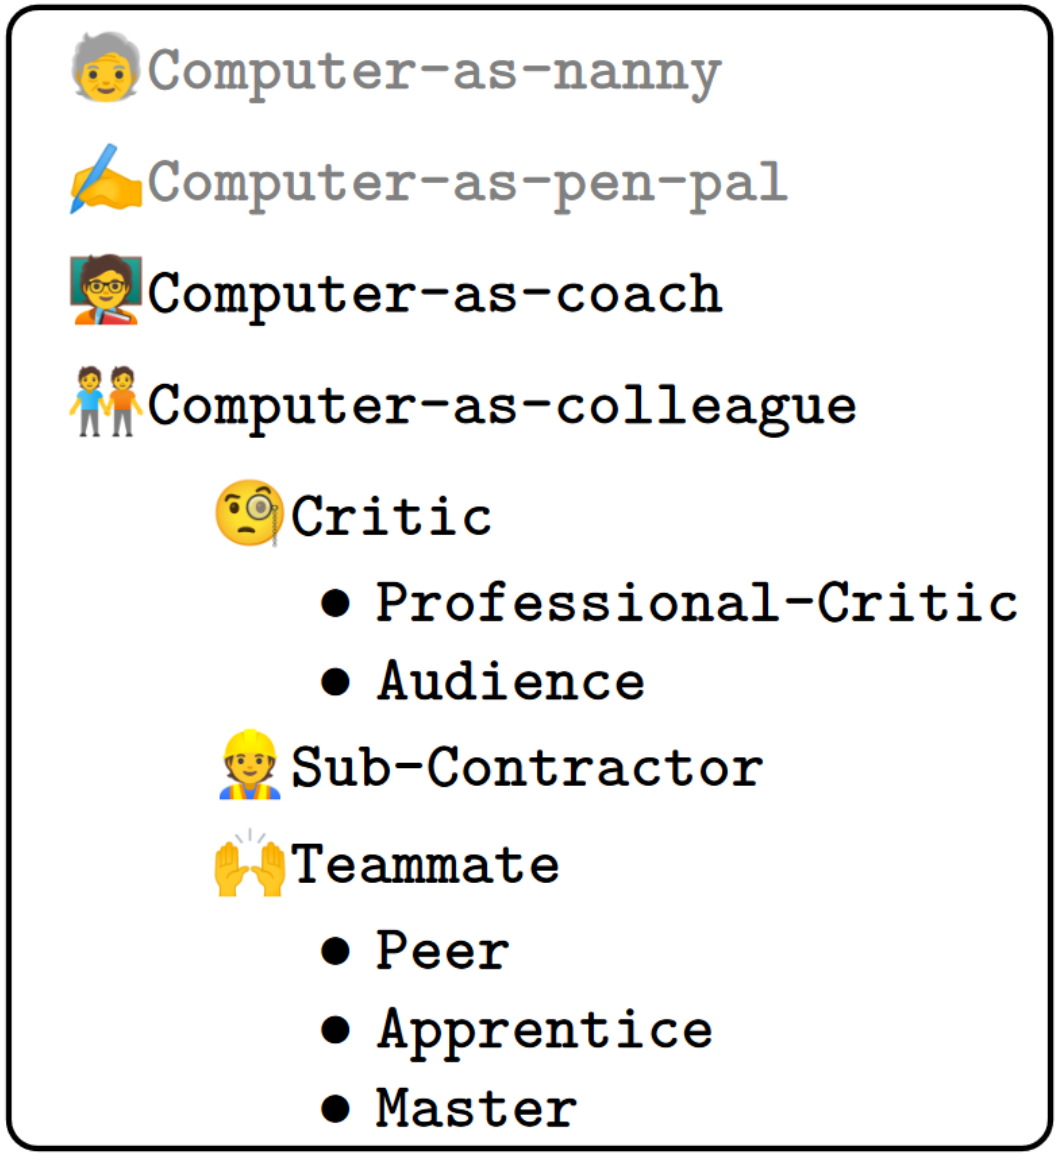

How would you expand co-creative paradigms?

Let AI probe best way to collaborate

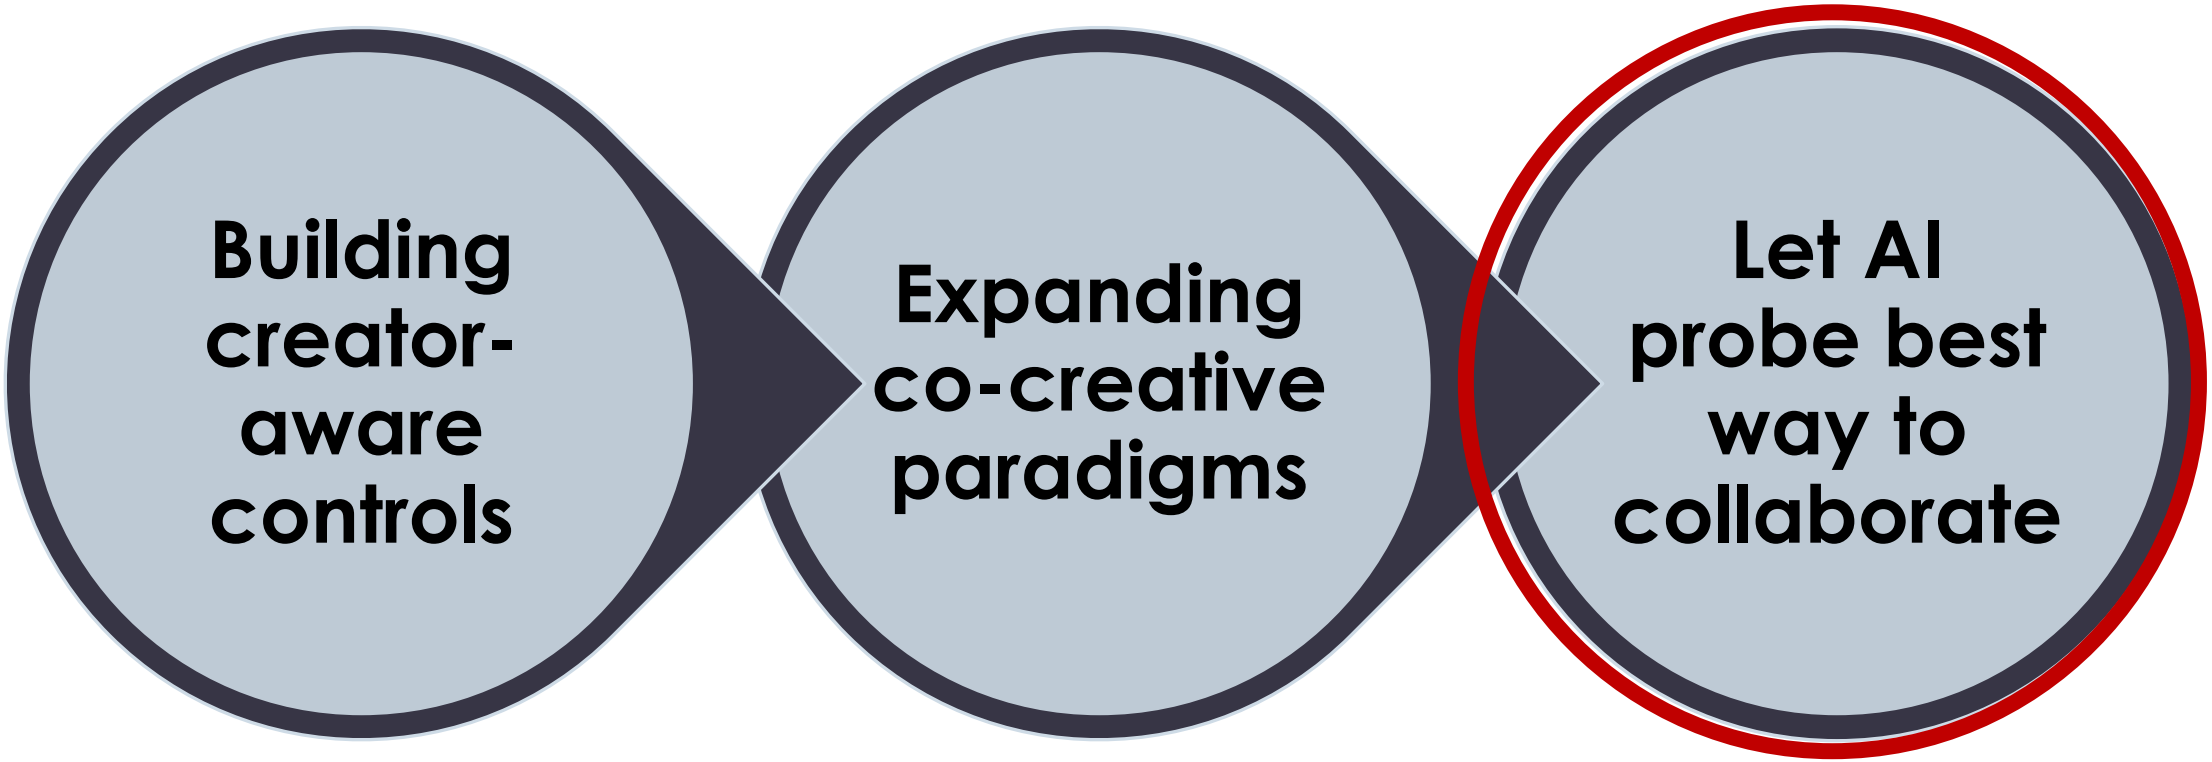

Let's take a deeper look into mixed initiative

Mixed-Initiative: Human initiative and a computational initiative cooperate towards a shared goal

- AI is capable of a wide range of tasks.
- How should AI collaborate with ***individual creators in mind***?
- **Let's explore how AI can learn to better collaborate!**

Human-AI Collaboration is non-trivial…

But when human learn things, **they will ask questions** on what they are unclear, so that they can learn more in less time!

Can we let AI do the same?

Lin, Z., Harrison, B., Keech, A., & Riedl, M. O. , Explore, exploit or listen: Combining human feedback and policy model to speed up deep reinforcement learning in 3D worlds (2017).

Challenge

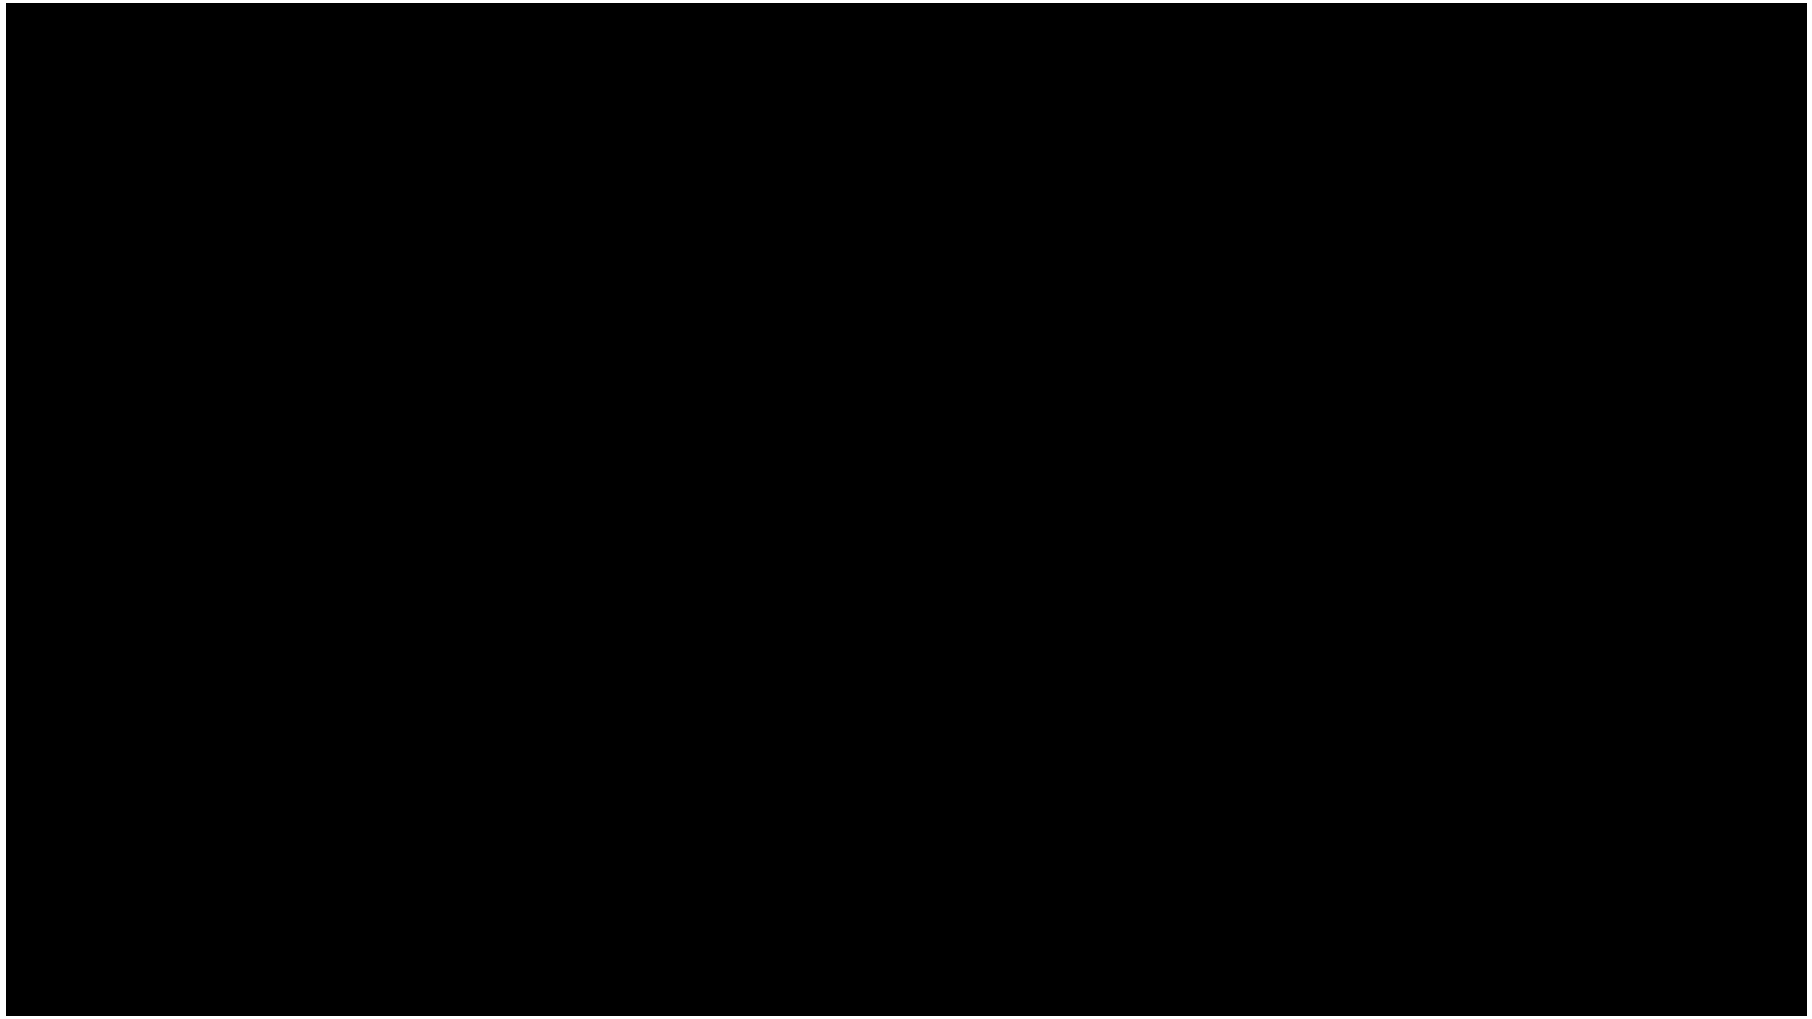

Formalization: Reinforcement Learning

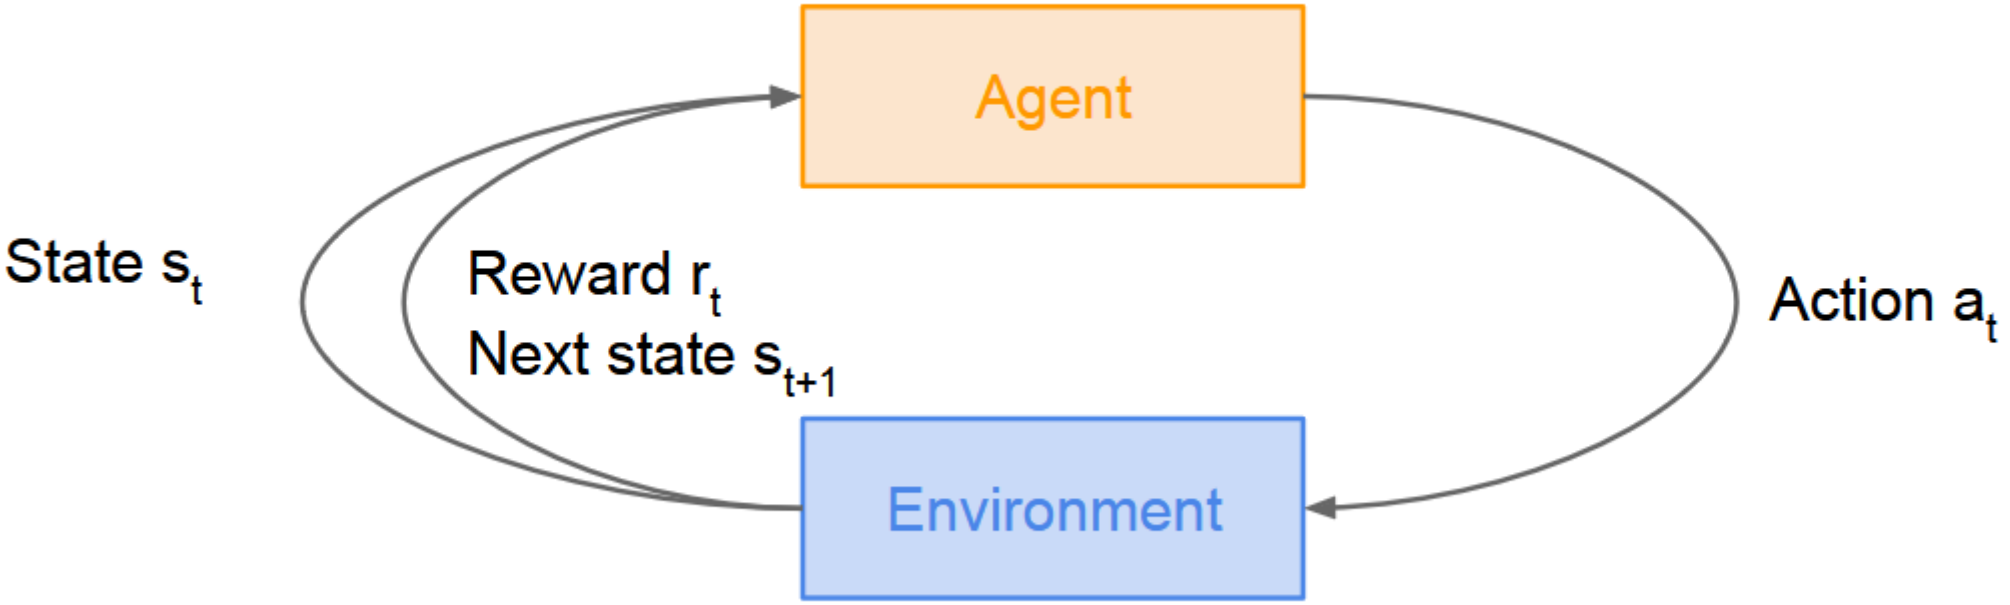

Slide credit: Fei-Fei Li & Justin Johnson & Serena Yeung

Training of RL in principle…

Explore

- Choose a random action to perform

Exploit

- Choose an action that the agent believe grants the best utility

How do we create a human-aware agent?

Explore	• Choose a random action to perform
Exploit	• Choose an action that the agent believe grants the best utility
Listen	• **Ask for advice on which action to take, when it is not certain**

When???

When should I ask for advice from human?

$$p_{\text{conf}} = \frac{-1}{\ln \sqrt{\frac{l}{l_{\max}}} - 1}$$

$$p_{\text{cons},t} = \begin{cases} \max(1, p_{\text{cons},t-1}) * f_1 * d & a_{\text{DQN}} = a_{\text{AA}} \\ p_{\text{cons},t-1} * f_2 * d & a_{\text{DQN}} \neq a_{\text{AA}} \end{cases}$$

$$p_{\text{explore}} = \begin{cases} 1 & t < r_{min} \\ e^{\ln 0.01 * \frac{t - t_{min}}{t_{max} - t_{min}}} & t_{min} \leq t < t_{max} \\ 0.01 & t_{max} \leq t \end{cases}$$

Confidence

"If I believe what I think, I should act more by myself"

Consistency

"If I heard what I thought, I should strengthen my belief more"

Training Curriculum

Strategically place random exploration sessions

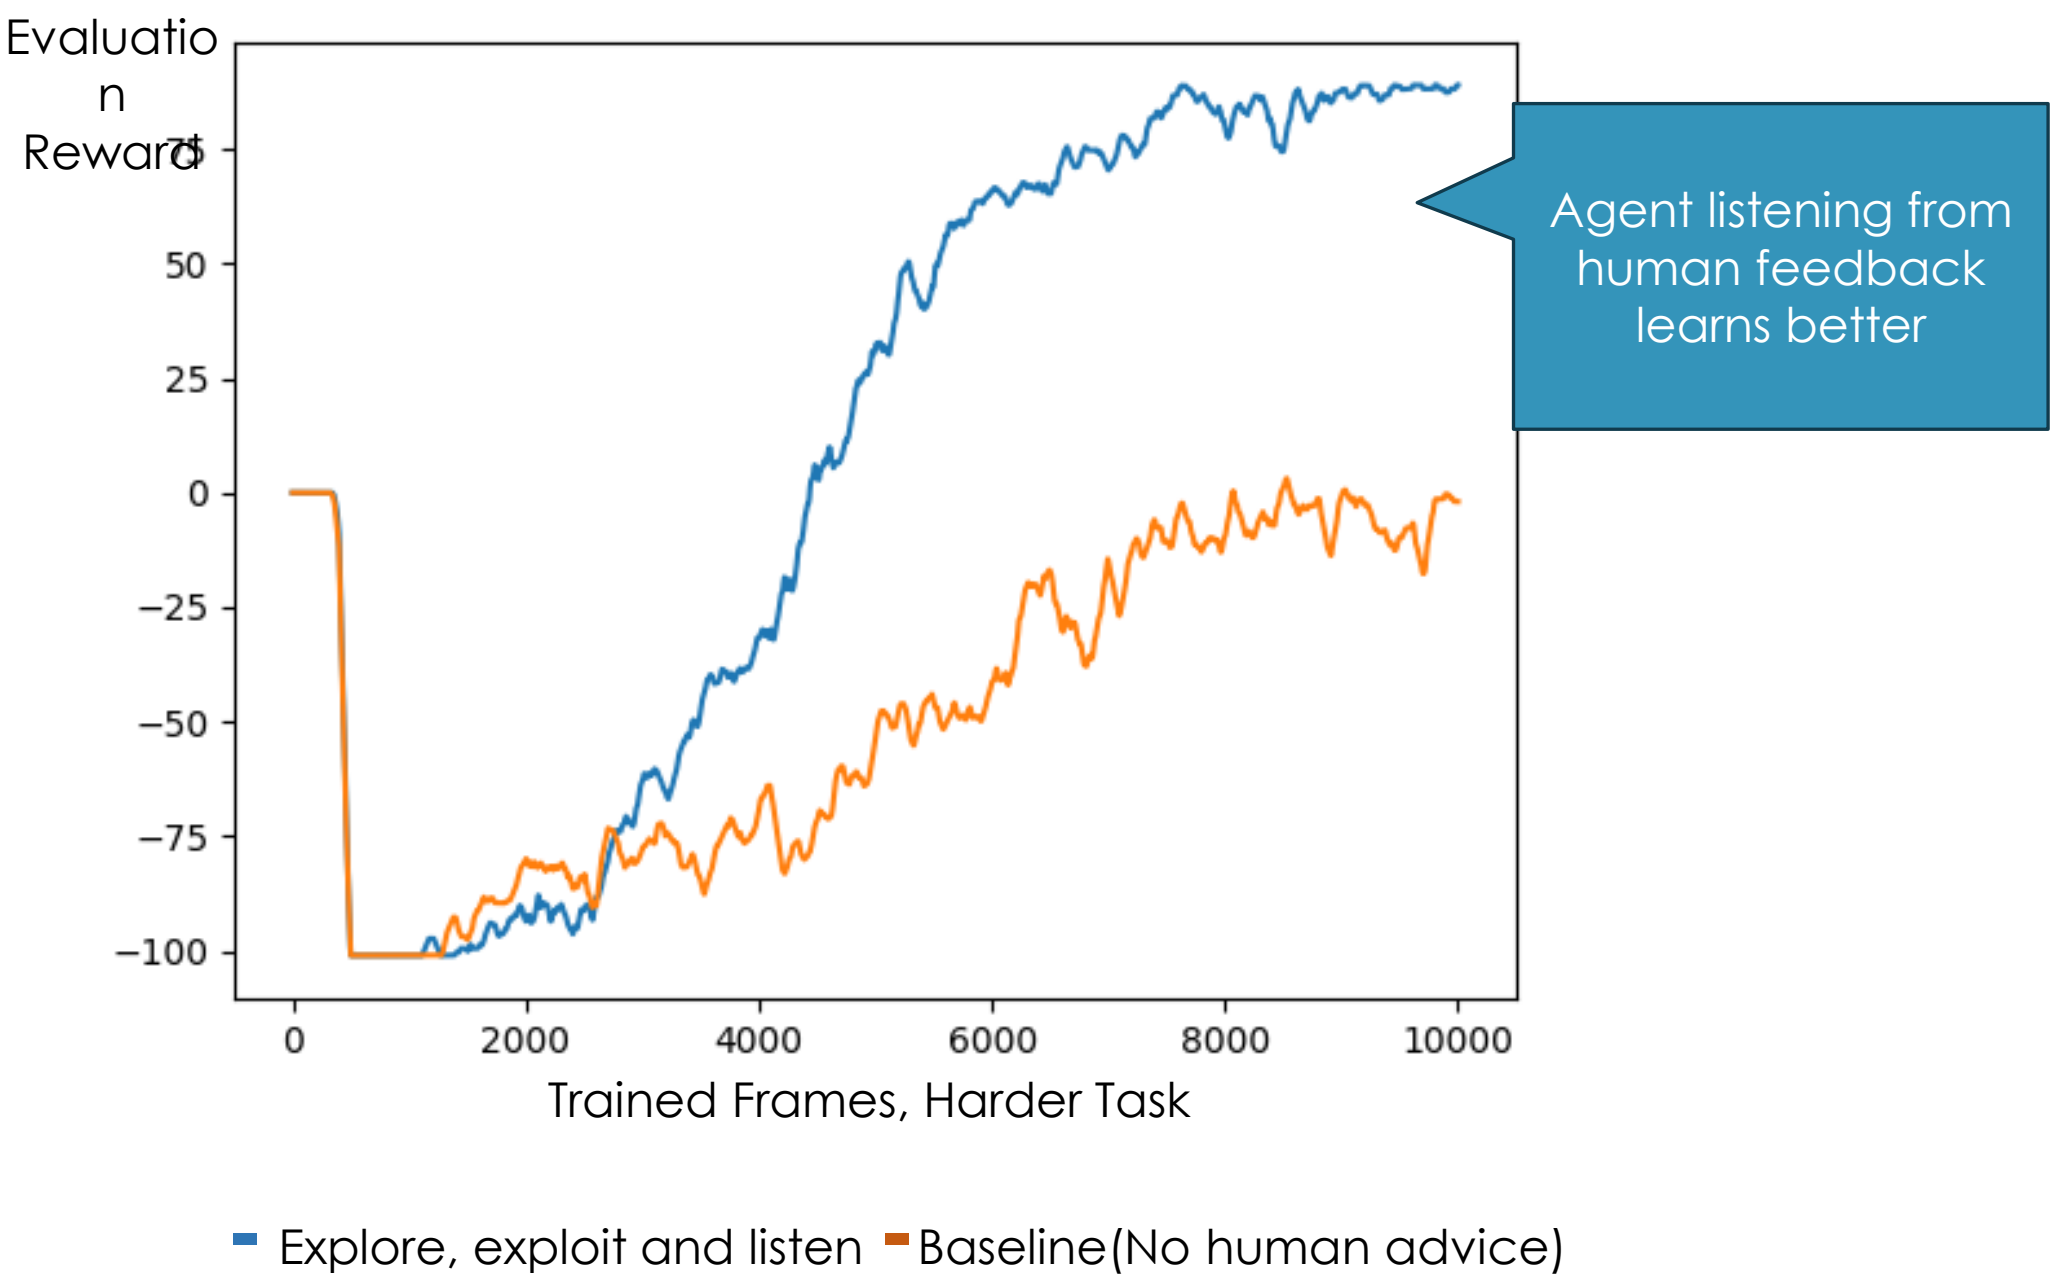
Evaluation Reward
75
50
25
0
−25
−50
−75
−100
0
2000
4000
6000
8000
10000
Trained Frames, Harder Task
Agent listening from human feedback learns better
Explore, exploit and listen
Baseline(No human advice)

Listening will help even with human make mistakes.

What have I learned?

By introducing human feedbacks into RL systems, we found performance boost especially in harder tasks

Human can help an AI agent learn, even if they are not an expert and make mistakes.

We can make an AI agent **aware of human and work with them**, in the process **improving its learning performance**.

Now, let's see how it works in MI-CC systems!

Lin, Z., Ehsan, U., Agarwal, R., Dani, S., Vashishth, V., & Riedl, M. (2024). Beyond Following: Mixing Active Initiative into Computational Creativity. EXAG @ AIIDE 2024.

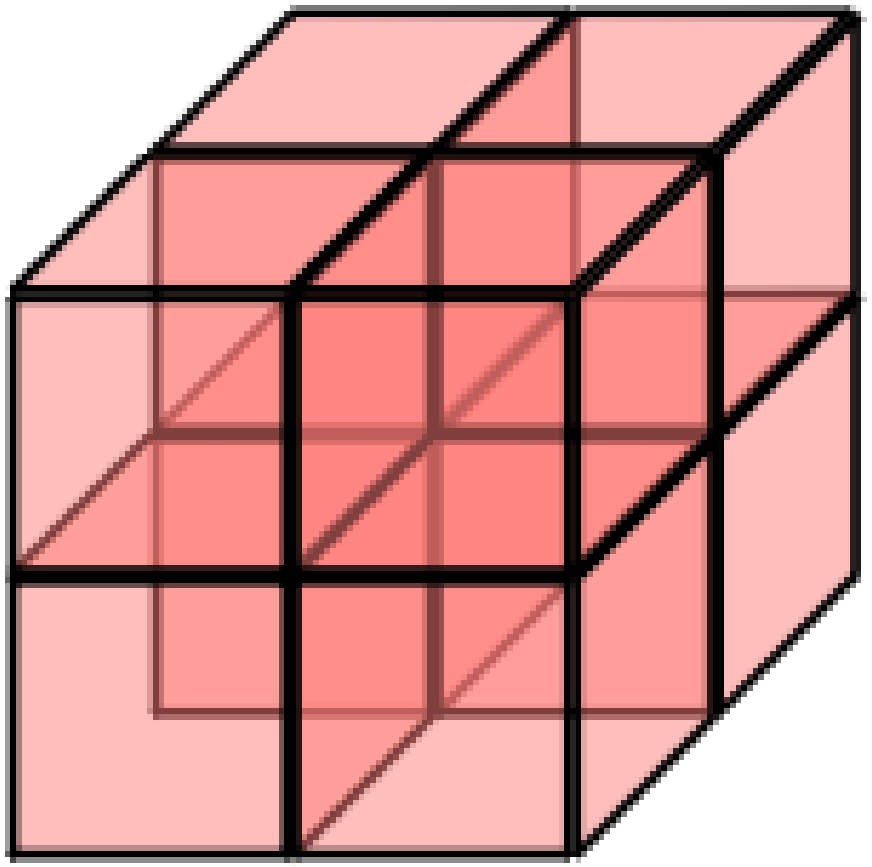

All communications

Generative AI are capable of a lot of things…

I have demonstrated the **diverse space** of communications in MI-CC systems, and the potential of such a system with **greater coverage of this design space, improving its creative support capability**.

Challenge

What should I do for ***the creator I'm working with***?

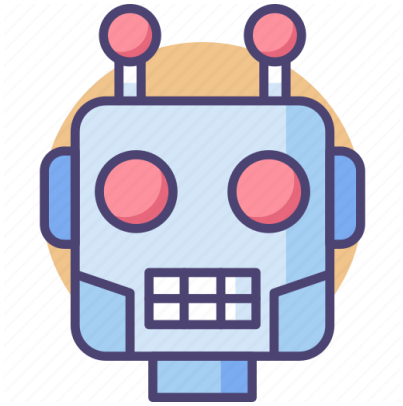

I should write more.

I should give suggestions

I should reflect on overall structures

…

Research Questions

How should a MI-CC agent actively learn from human creators and update its collaborative behavior?

How will this influence the perception of the creators and the computational creative experience?

Task design: Learn a delegation

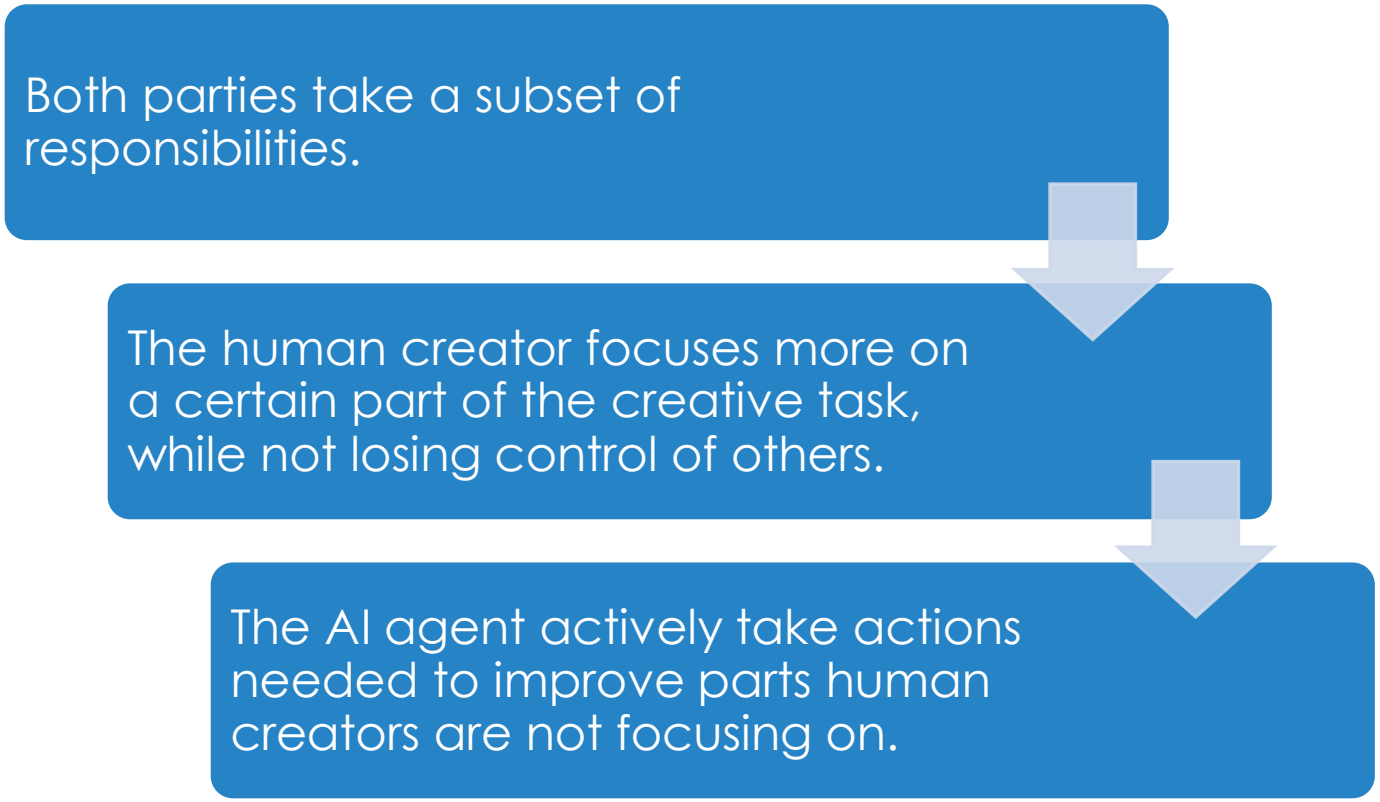

Captures *collaborative preference.*

Creators can evaluate.

Domain: Writing stories together

Story
Beginning
Development / Uprising
Climax
Conclusion

They are all good by itself. What should I do to best collaborate?

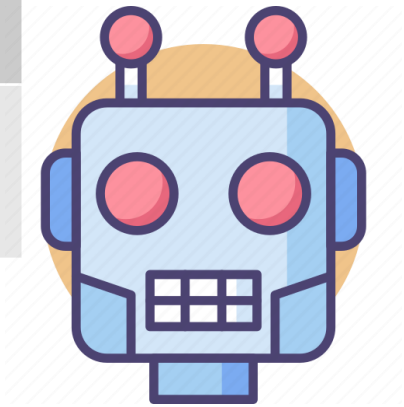

I should improve the first half of the story.

I should improve the second half of the story.

I should reflect and provide suggestions

Multi-Armed Bandits

Formalization: Multi-Armed Bandit (MAB)

- Many Communications the agent can choose from casino
- Using them to co-create will give you a reward
- Which one should I pull to maximize

Reward from human creators?

We built an agent equipped with multi-armed bandits

Thompson Sampling is chosen.

- Every action receives a score, sampled from Beta (α, β) each time. Action with highest score is chosen
- For each action:
  - α starts at 1 and increased by reward received.
  - β starts a d).

$$a = argmax(\mathcal{B}(\alpha_a, \beta_a))$$

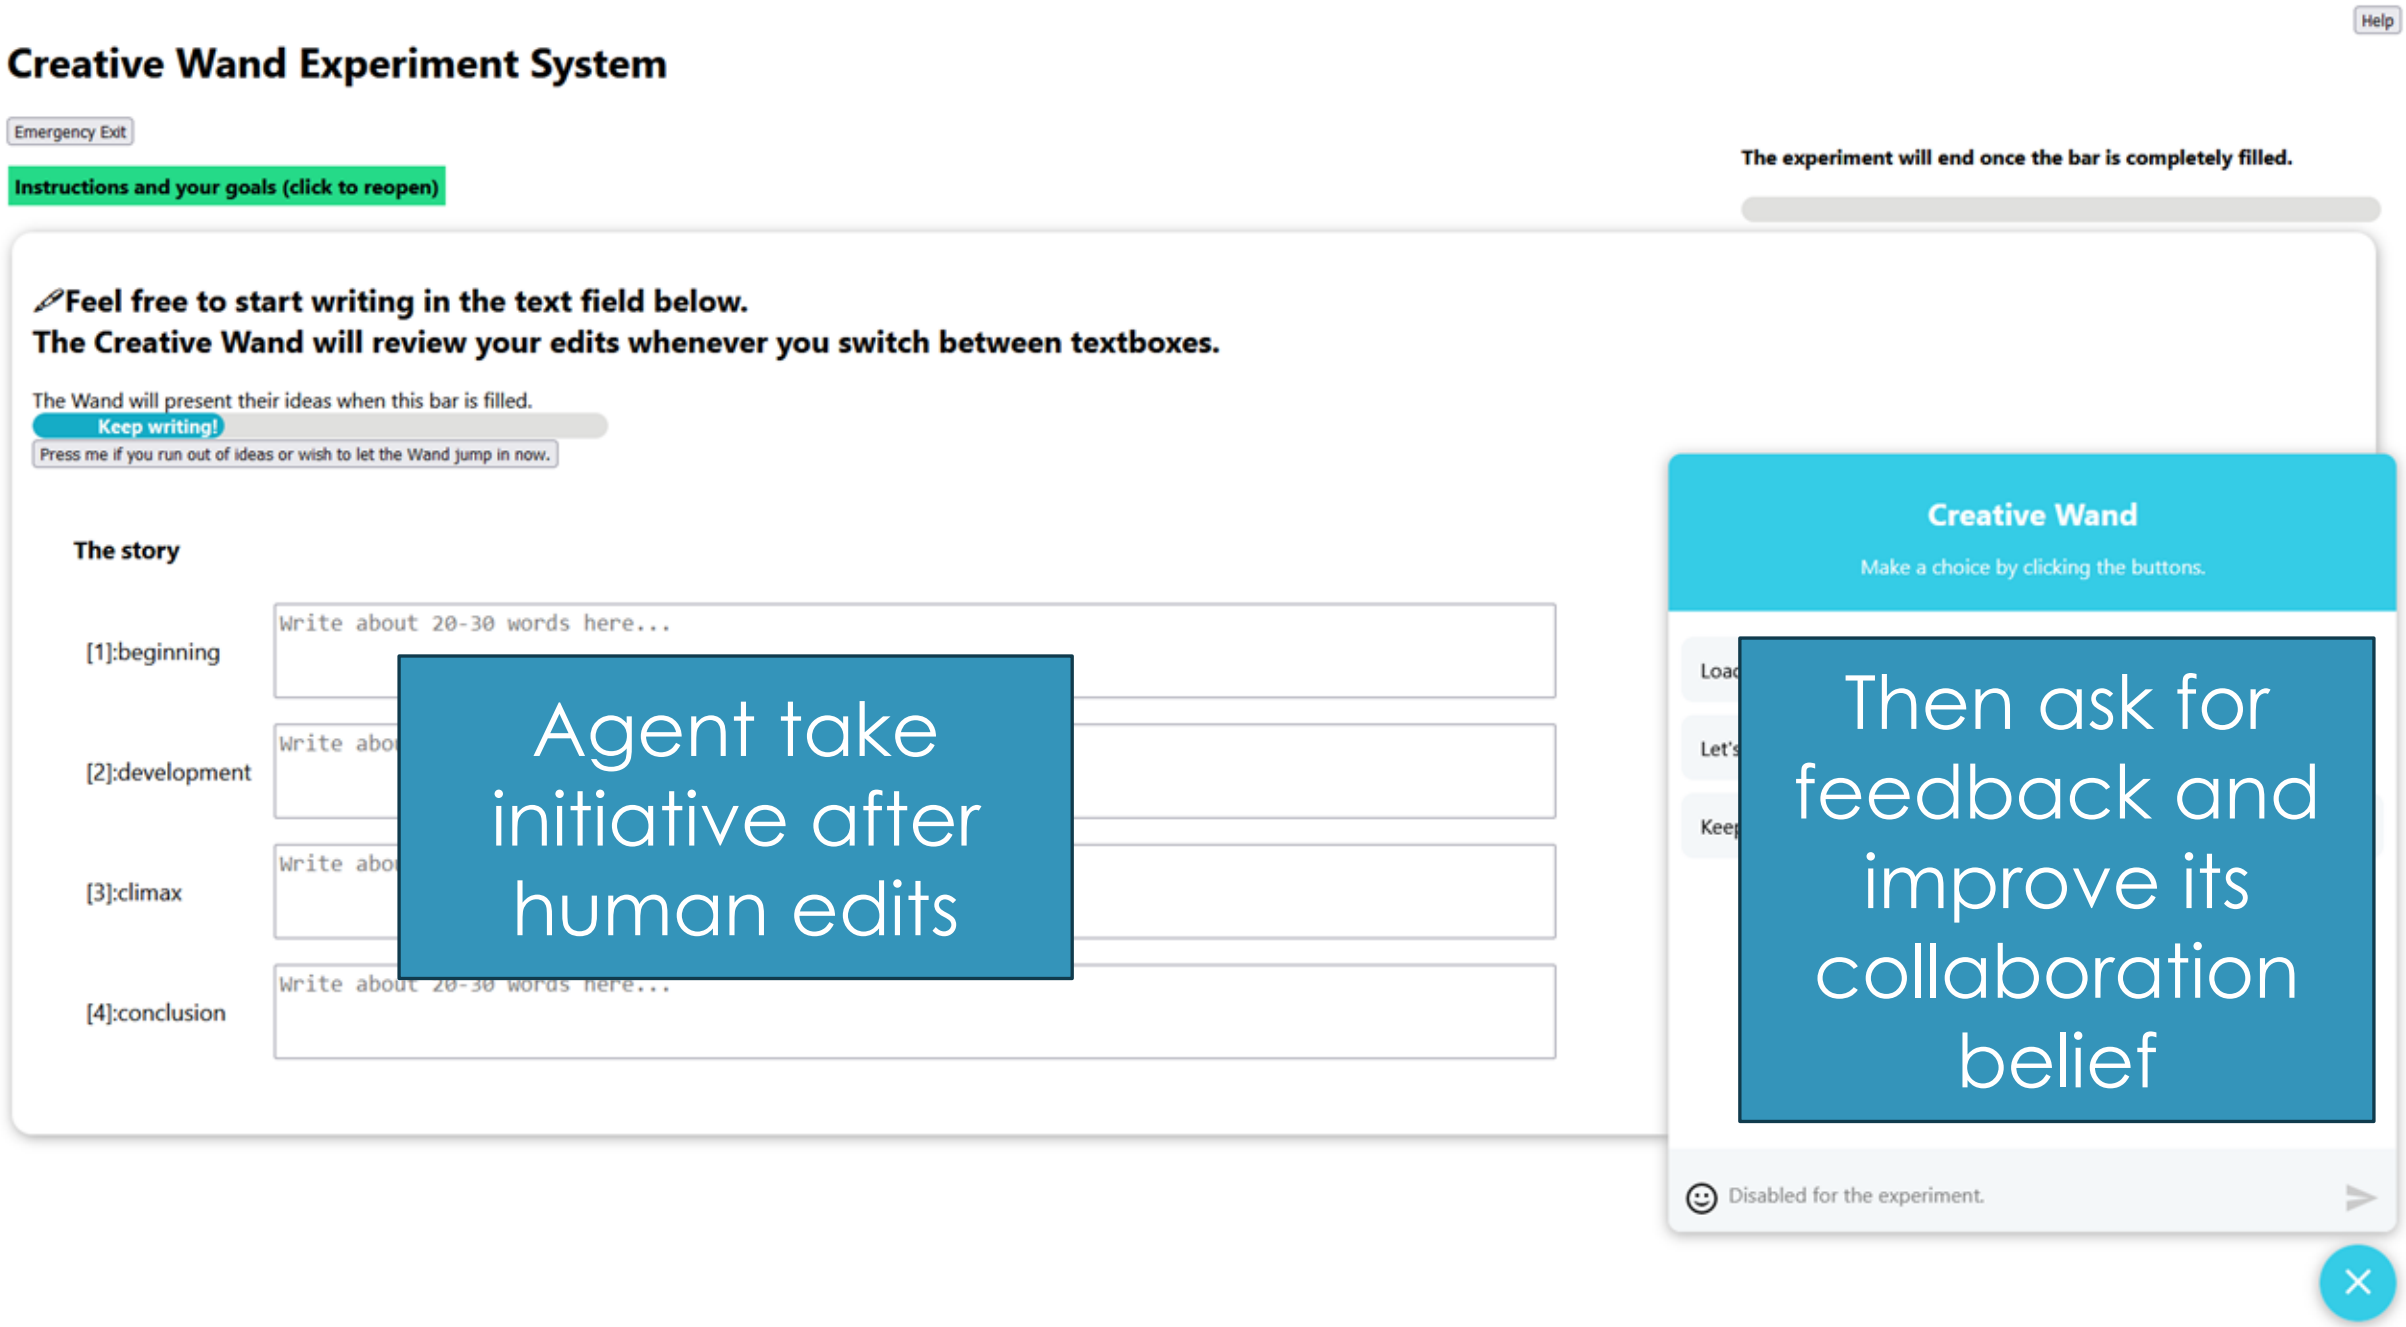
Help
Creative Wand Experiment System
Emergency Exit
Instructions and your goals (click to reopen)
The experiment will end once the bar is completely filled.
Feel free to start writing in the text field below.
The Creative Wand will review your edits whenever you switch between textboxes.
The Wand will present their ideas when this bar is filled.
Keep writing!
Press me if you run out of ideas or wish to let the Wand jump in now.
The story
[1]:beginning
Write about 20-30 words here...
[2]:development
[3]:climax
[4]:conclusion
Agent take initiative after human edits
Creative Wand
Make a choice by clicking the buttons.
Then ask for feedback and improve its collaboration belief
Disabled for the experiment.

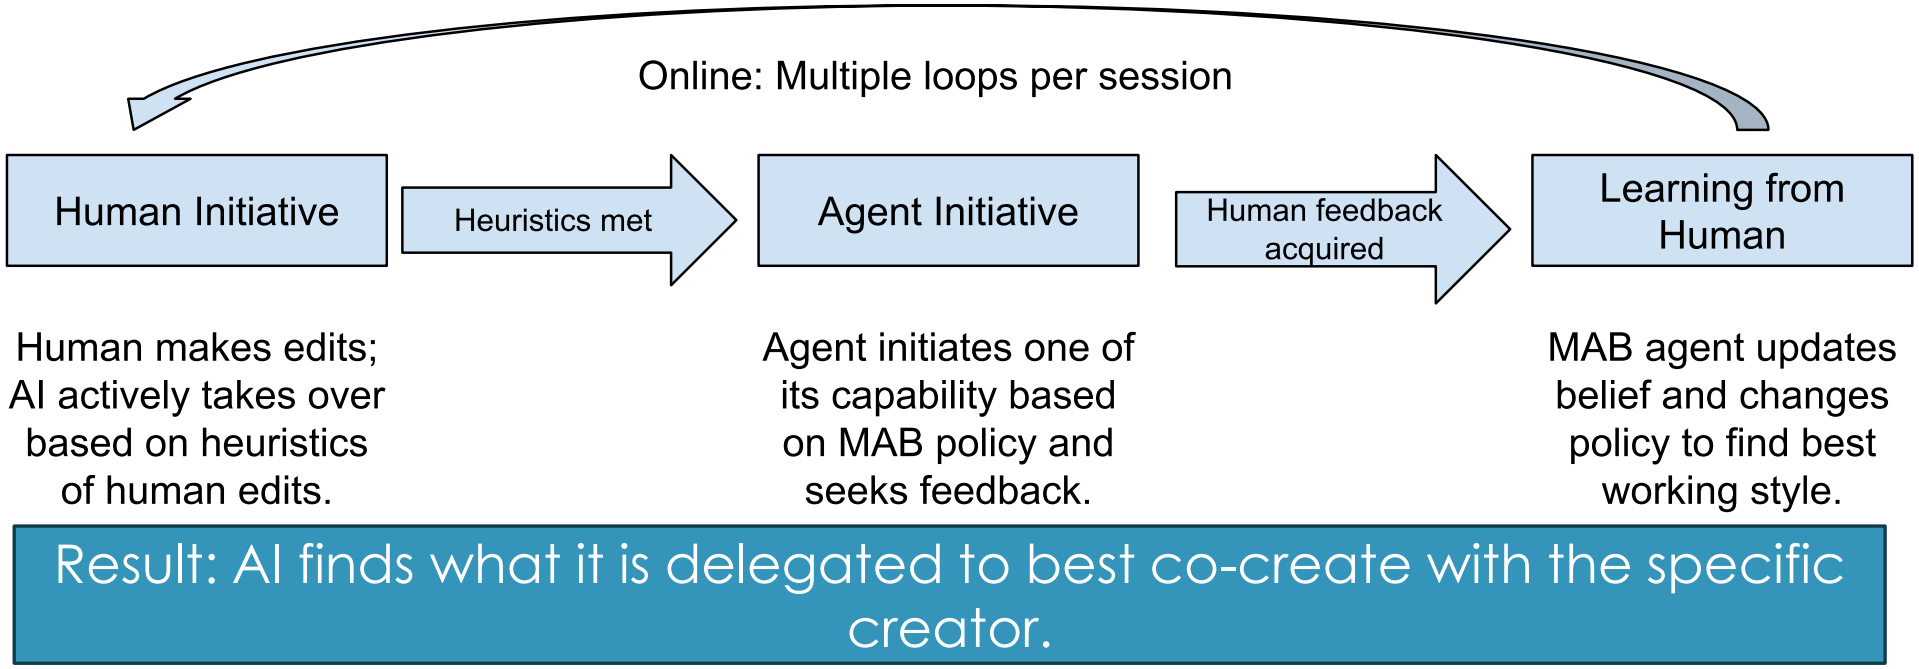

INTEGRATING ACTIVE LEARNING

Human Initiative

Human creators contribute to the story by making edits in **any text fields**, **regardless of the task** they are assigned to.

This phase ends when the **agent decides to take the initiative**, determined by a heuristic based on edits made by human creators.

Agent Initiative

Agent queries the Multi-Armed Bandit Agent for an action to take

Agent then communicate with human creators according to the action chosen.

Learning from human creators

The system seek either a positive or negative feedback, in the form of two option buttons.

- Action feedback - Should I keep working on the story in this way?
- Content feedback - How about the changes and updates I brought to you?

Study Methods

Mixed-method research

39 participants recruited from Prolific returned a valid response

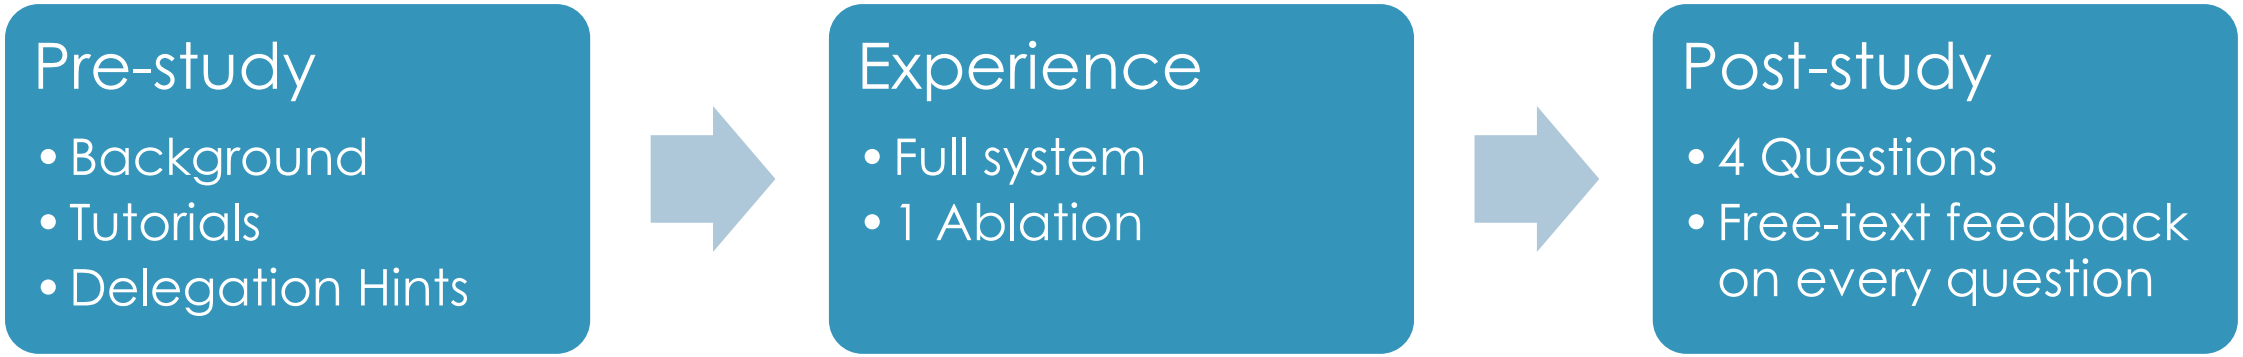

Two systems for participants to evaluate

Full system "The Echo Wand"

- Everything included, including the Multi Armed Bandit agent

Baseline "The Harmony Wand"

- MAB agent disabled: always chooses a random action.

Key findings

Likert Scale Scores: 1 = Most Negative, 5 = Most Positive

	1	2	3	4	5	Average	Median
Q1: CAD skills	1	1	2	19	16	4.23	4
Q2: Writing skills	1	0	7	20	11	4.03	4
Q3: Frequency of using AI	0	0	16	11	12	3.90	4
Q4: Understanding of AI Tech.	0	5	14	19	1	3.41	4

Table 9.1: Participants of the study, grouped by their answer to the creative background questions at the beginning of the survey. See section 9.4 for the full question texts.

Background of the participants

Which system learned to collaborate with the participants under that arrangement?

- The "Full" system with **learning capabilities** being perceived significantly better at learning the delegation than the baseline.
- Our MAB-based model is effective **from the human creator perspective**

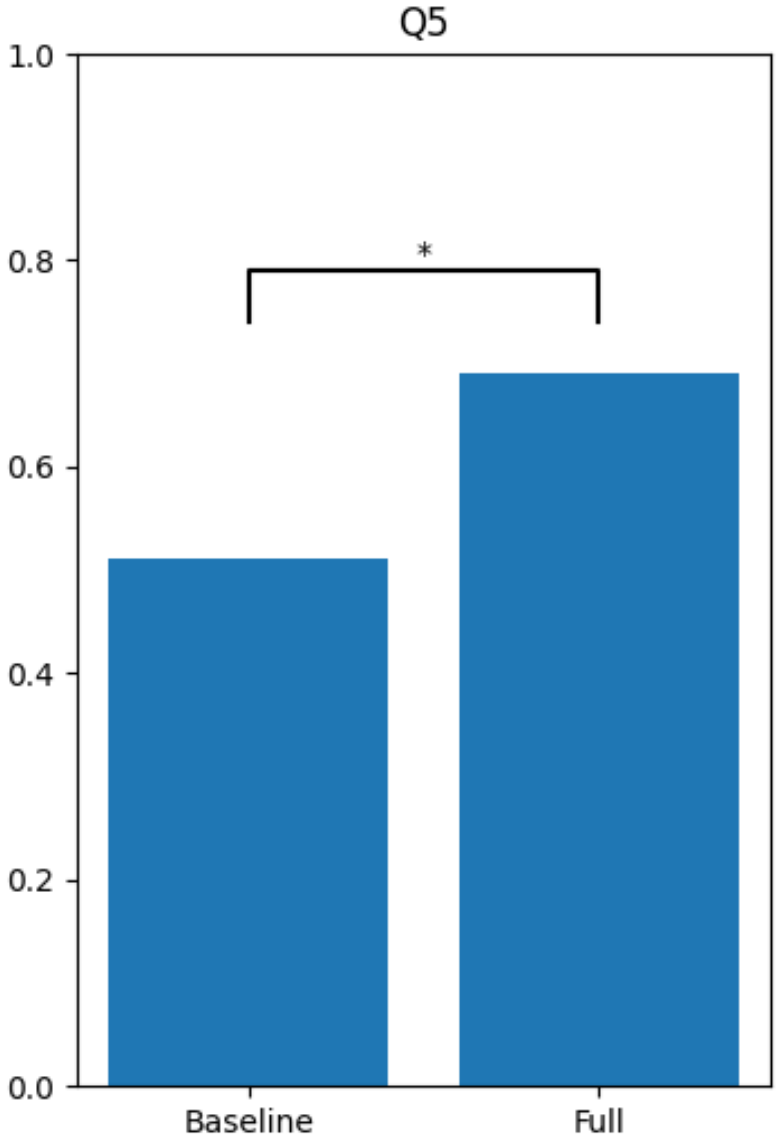

Which system would you recommend?

- The **only difference is the learning capabilities**!
- This statistically significant improvement in preference illustrating the **potential of our method in enhancing MI-CC experience, making it better for human creators**.

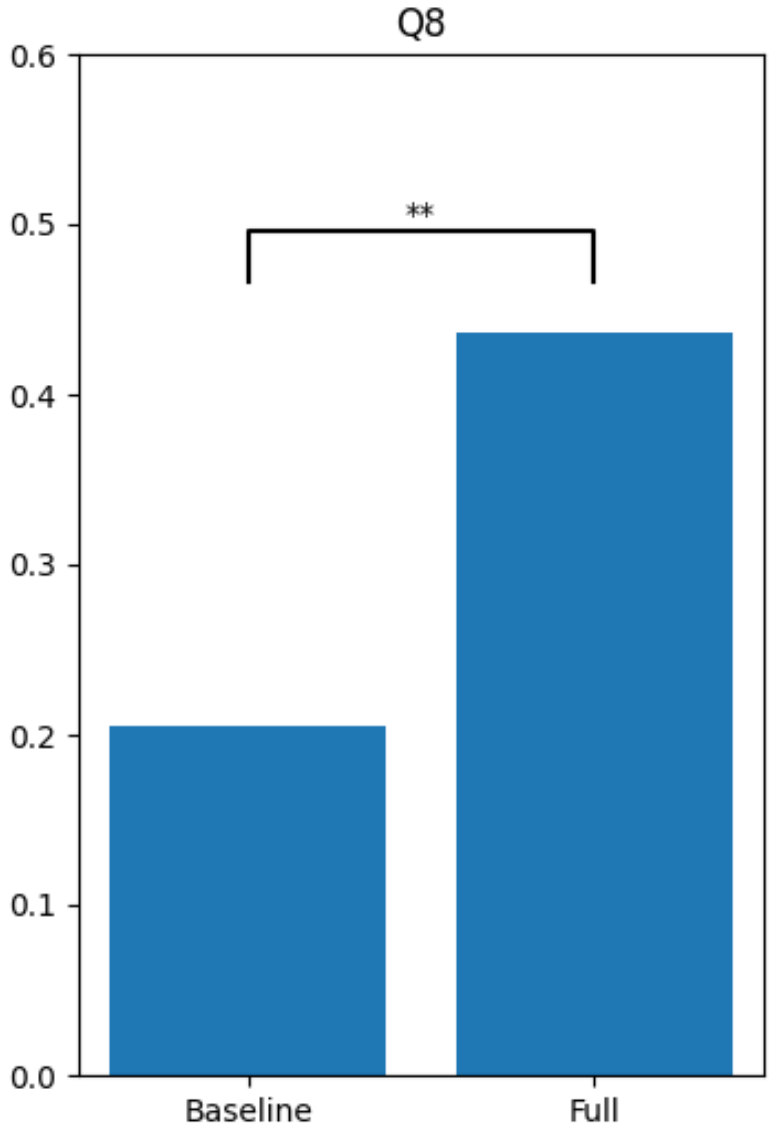

Key Qualitative Findings

The learning agent is favored and collaborates well.

Good contents also give sense of learning and collaborating

Diversity is important, always doing “best action” is not ideal

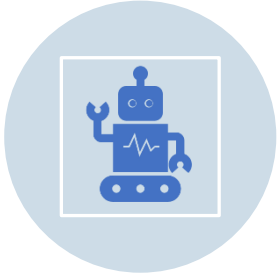

Creator seek control even when AI gives plausible candidates.

A MI-CC system that **understands the intents of the human creators** and **follows them by learning** is in overall **favored and collaborates well** with the creators.

Full system "better about learning that I specifically wanted help with" (P34) and "more useful helper" (P32)

Baseline system "did less of the work … did not necessarily learn what its role was expected to be (P19)"

Aligns with Quantitative findings

Good content suggestions may **give people the feeling** that the system is learning how to collaborate with them, **regardless of how AI is actu-ally collaborating** with them.

"This one learned from me because it was able to build off of my original foundation of my story that I typed." (P25)

I could see Echo Wand adding more detail and building out more cre-atively than with Harmony Wand. (P18)

LM capability and choice of them is coupled.

Diversity and Controllability

Baseline “had much more interesting suggestions … balanced the second two sections to match my intro and build up, unlike [Full system] who almost refused to work on them. (P36)

Diversity is also important, it may not be the best strategy for a learning agent to pick the “best options”, and sometimes the agent may want to intentionally surprise their teammates.

AI “finish the story that I started with.” (P28)

[I] was in control of the final text to accept changes or not, or to make my own. (P27)

Creator control is important, and creators may want their ideas to be included even when AI can also provide plausible candidates.

Mental model – From both directons

The success of our **Full** system on learning comes from its ability to **learn a model of how the creators wish to collaborate** with them.

Creators are interpreting the capabilities of our AI agent learning as an attempt of it to **learn a mental model of themselves**.

Expectations are important.

What did we learn?

We now see how a MI-CC system is capable of **listening from human feedback and improving itself.**

This capability is **well recognized by the participants and led to better satisfaction** overall.

Now & Future

"This control was clearly not designed with me in mind!"	• I tailored creator-aware controls
"I want to control them in diverse ways that work for me."	• I expanded the information exchange paradigm
"There are so many kinds of control, which one should I use?"	• I enabled an agent to learn from the specific human creator in a MI-CC setting

I've shown my story of improving experiences of human creators.

Is this the *end of the story*?

Into the future

How should these agents handle novelties?

How should we educate human to learn these systems?

How should we facilitate sharing of mental models?

How do we make this process fun?

Human-Awareness to Human-Centered

How will you let AI go *beyond following*?

Thank you! More to come ☺

https://www.zhiyulin.info

zlin34@ucsc.edu

UC SANTA CRUZ

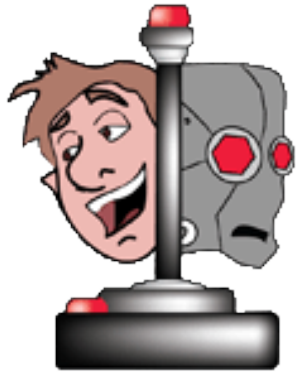

Georgia Tech

Shoutouts to all my great colleagues, collaborators and friends!