Google DeepMind

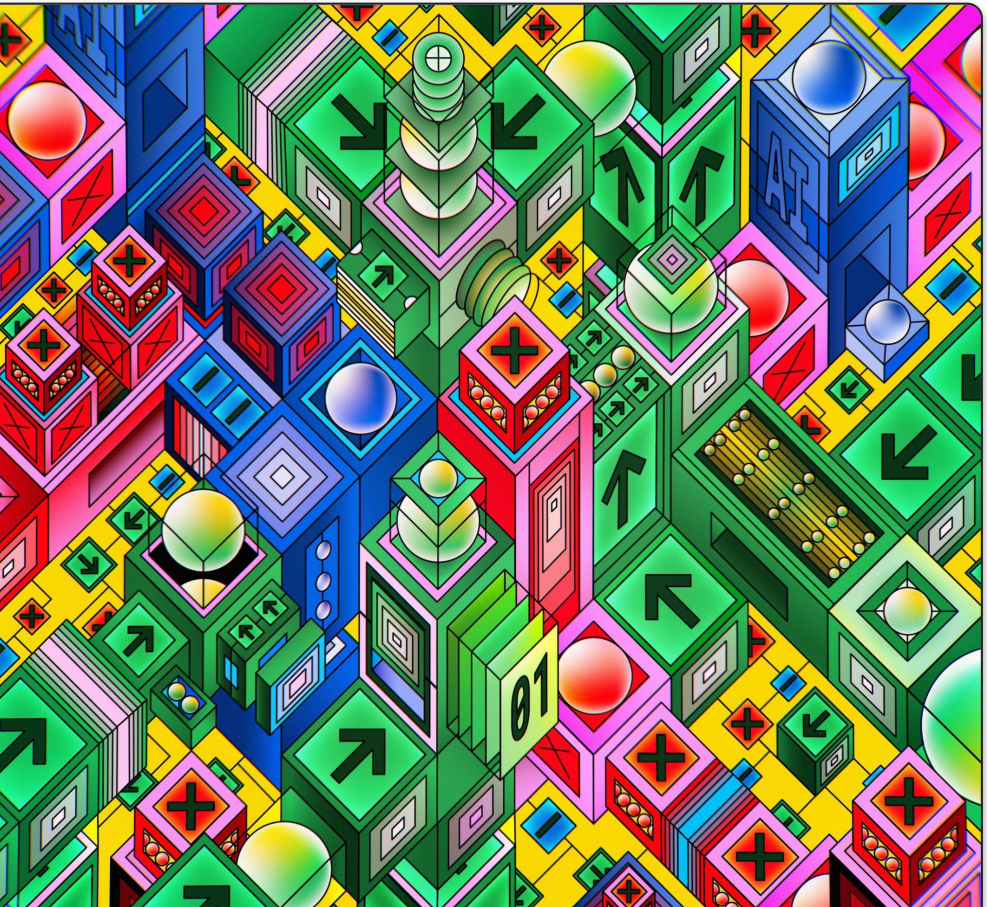

Modern Video Games as a Testbed for Developing

Generalist AI Agents

Joe Marino
Research Scientist
Google DeepMind

Scaling Instructable Agents Across Many Simulated Worlds
arXiv:2404.10179

00 Introduction


60
THE UNIVERSITY OF EDINBURGH
Edinburgh Computer Science & AI @ 60




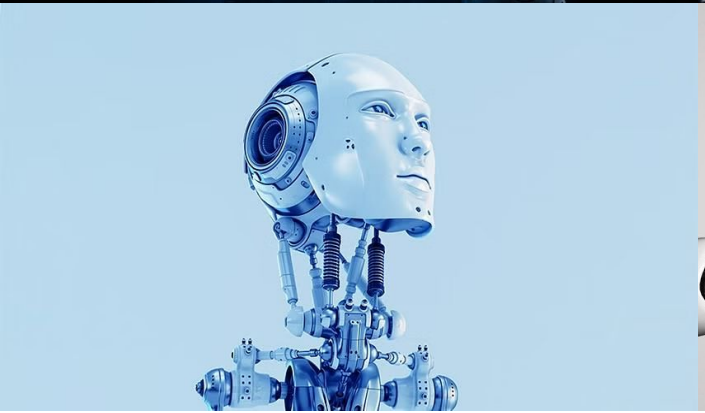
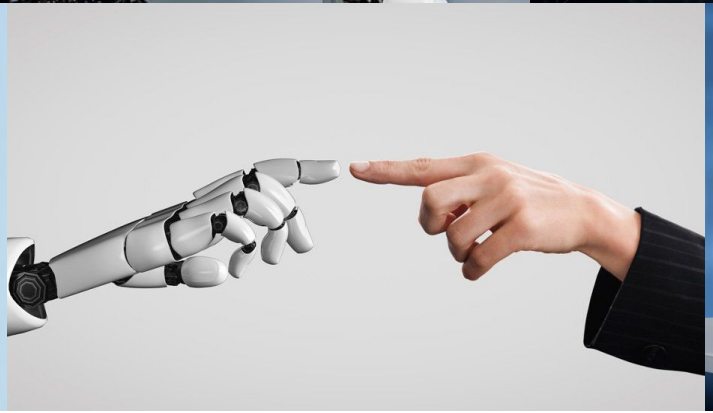


AI in the 2020s - *Humanoid Robots?*

Figure 1

1X

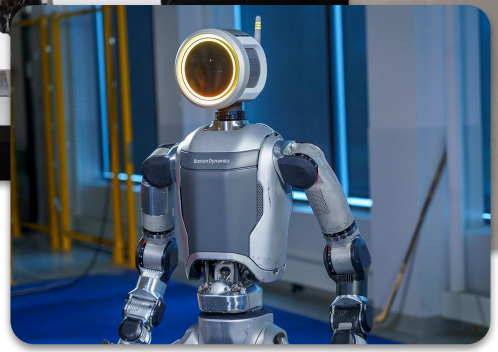

Boston Dynamics Atlas

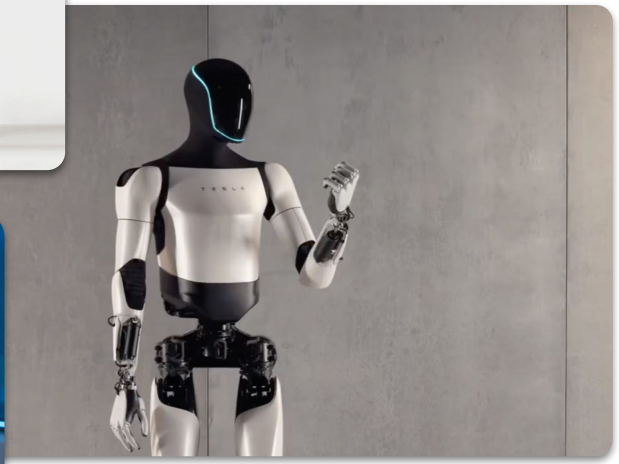

Tesla Optimus

AI in the 2020s - *"Generative AI"*

Large (Vision &) Language Models

- Programming Assistant
- Idea Brainstorming
- General-purpose Visual Question Answering
- ...

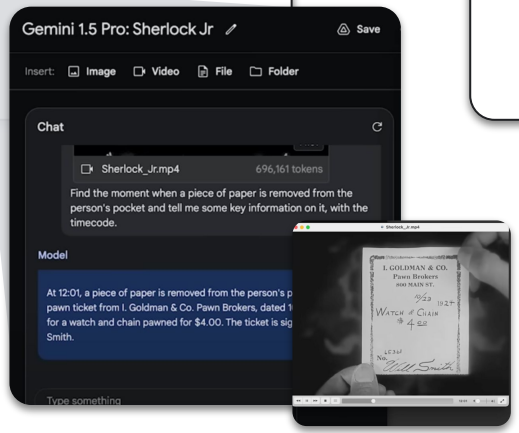

Image, Video, Audio Models

- Storyboarding
- Personalized Images, Music, etc.
- Feature-length films?
- ...

AI in the 2020s - *"Generative AI"*

Large (Vision &) Language Models

- Programming Assistant
- Idea Brainstorming
- General-purpose Visual Question Answering
- ...

Image, Video, Audio Models

- Storyboarding
- Personalized Images, Music, etc.
- Feature-length films?
- ...

Current models offer a largely ***disembodied*** version of AI.

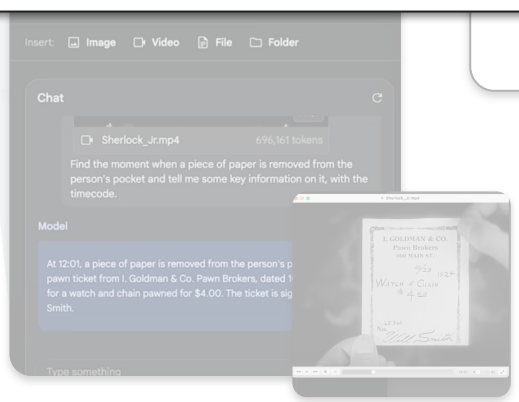

Embodiment

embodiment: *the degree to which a system can affect / control its sensors*

embodiment x ***large sensor space*** → *the system can reach many states of the environment (tasks)*

Current systems can bring about many useful states in the ***digital realm****,*
but remain largely unable to control the ***physical realm****.*

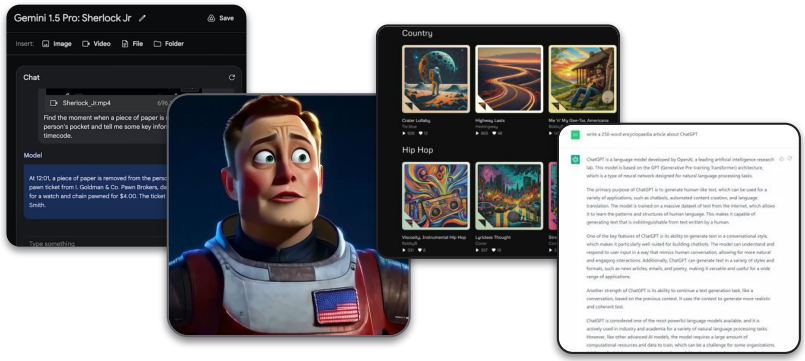

Biological / Neuroscience Perspective

Common Theme

*The purpose of the brain is to coordinate the body's **movement**.*

evolution x neuroscience

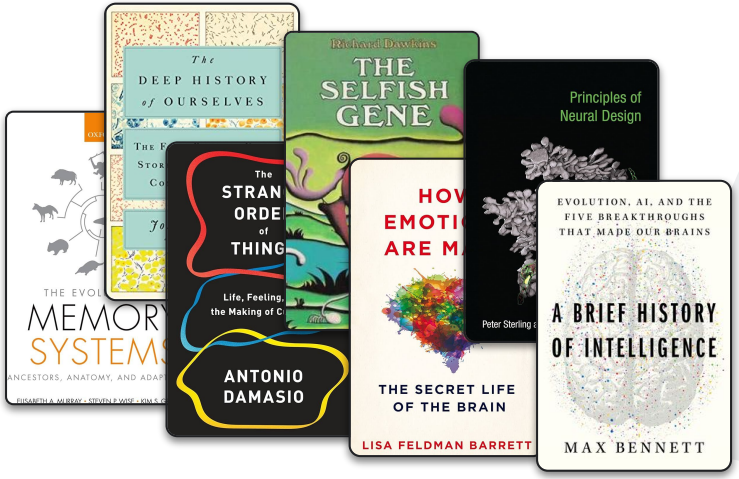

5 Breakthroughs in the Evolutionary History of Our Brains

Steering
600 million years ago

Learning from direct experience of valence

Reinforcement Learning
500 million years ago

Learning from temporally-extended experience

Mental Simulation
200 million years ago

Learning from internally-simulated experience

Mentalizing
15 million years ago

Learning from observing others' experience

Language
100 thousand years ago

Learning from others' internally-simulated experience

Biological / Neuroscience Perspective

Common Theme

*The purpose of the brain is to coordinate the body's **movement**.*

evolution x neuroscience

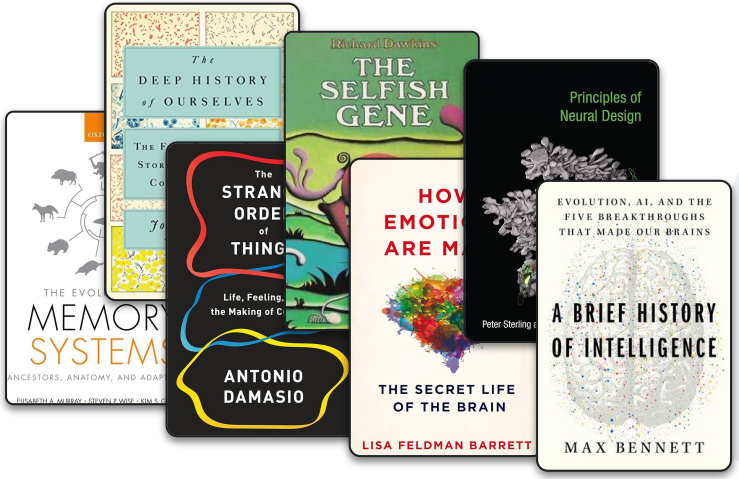

5 Breakthroughs in the Evolutionary History of Our Brains

Steering
600 million years ago
Learning from direct experience of valence

Reinforcement Learning
500 million years ago
Learning from temporally-extended experience

Mental Simulation
200 million years ago
Learning from internally-simulated experience

Mentalizing
15 million years ago
Learning from observing others' experience

Language
100 thousand years ago
Learning from others' internally-simulated experience

Challenges of Physical Embodiment

Why haven't consumer robots taken off?

- *Difficult to **scale** experiments.*
  - *Upfront and marginal costs are significant.*
- *Difficult to gather diverse **data**.*
  - *Almost entirely limited to sterile lab environments.*
- *Limited **experimental control**.*
  - *Research is even more challenging.*
- ***Safety** considerations.*
  - *Huge liability for making mistakes (i.e., progress).*

As a result,

- *current robotic systems live in 'impoverished' sensory environments,*
- *robotics research progress has been slower.*

Games & Simulated Environments in AI

Atari
2012-2013

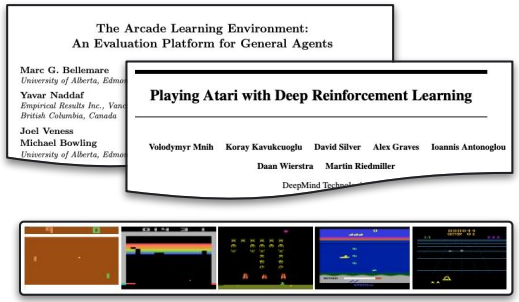

Minecraft, Doom, Sega, etc.
mid/late-2010s

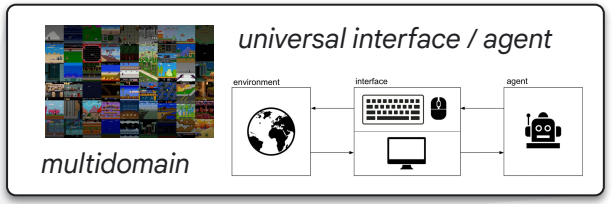

Dota 2, Starcraft II
late-2010s

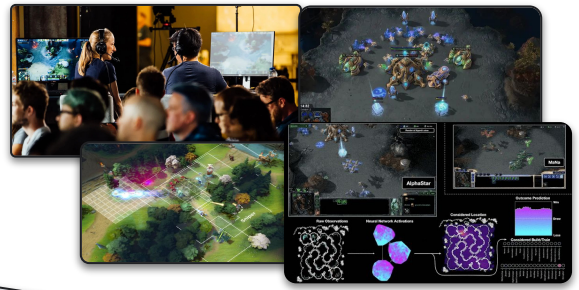

Pre-DL
early-2010s

LLMs/VLMs
early-2020s

Increasingly Complex

- *Observations*
- *Actions*
- *Dynamics*

but...

- *Largely single-domain*
- *Bespoke action spaces*
- ***Limited number of tasks***
- ***Limited task specification***

Limitations of the mid-2010s Era Approaches

Vision / Action Diversity

Limited Trajectory Diversity
→ Task = Win the Game

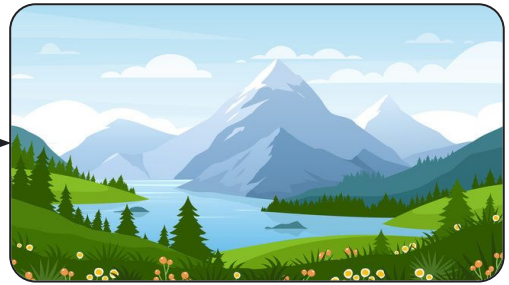

Trajectory (Task) Diversity

Limited Visual / Action Diversity

Limitations of the mid-2010s Era Approaches

01 SIMA

SIMA: *Scalable Instructable Multiworld Agent*

*A **single agent** with a **universal interface** that can be **instructed via language** to perform **any task** in **any 3D visual environment**...*

Data

Dataset

text

Training

Pretrained Models

Agents

Data Collection

text

SIMA Agent

Scalable, Instructable, Multiworld Agent

(Human) Evaluation

text

Commercial Video Games

TEARDOWN

NO MAN'S SKY

SATISFACTORY

HYDRONEER

Goat Simulator 3

ProcTHOR

WorldLab

VALHEIM

Wobbly Life

Construction Lab

Playhouse

Environments

Research Environments

pick up leaves

Evaluation

SIMA ***Environments***

Desiderata

- *3D visual environment*
- *First-person*
- *Complex, open world (~sandbox)*

Ideally

- *Interesting to humans*
- *Not violence focused*
- *Diverse*

Commercial Video Games

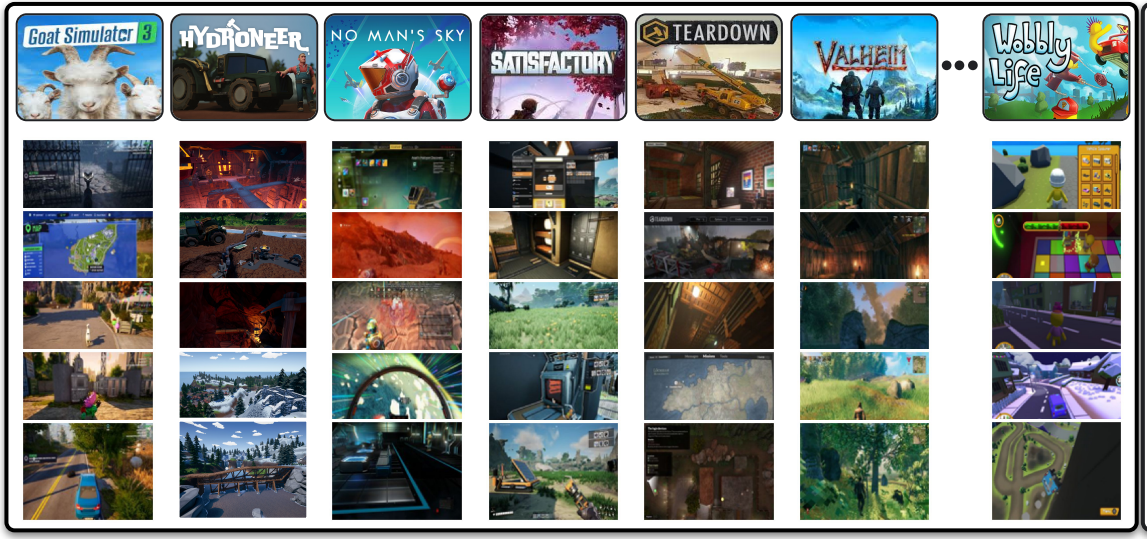

Research Environments

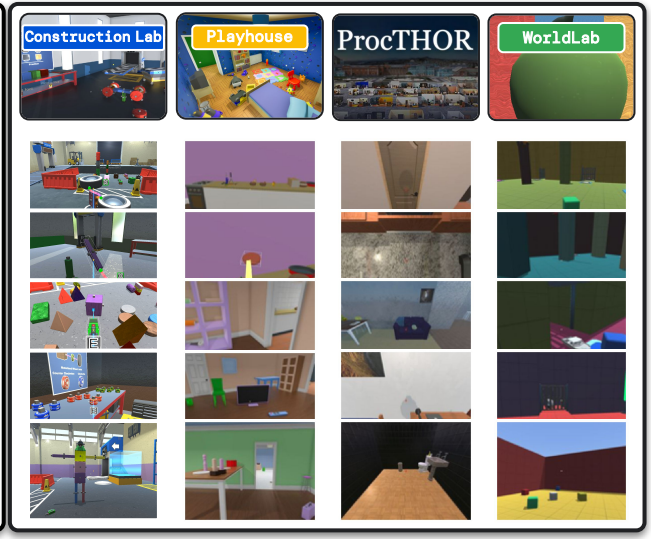

SIMA *Data*

*Focus on collecting **human data** to initially scale up.*

Data Collection

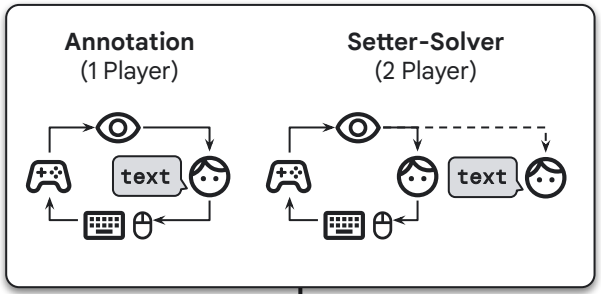

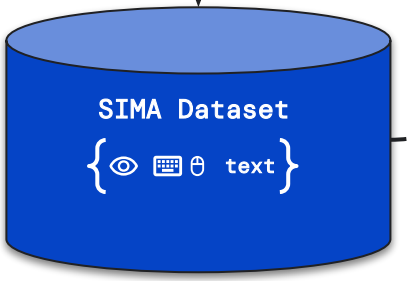

Skill Clustering

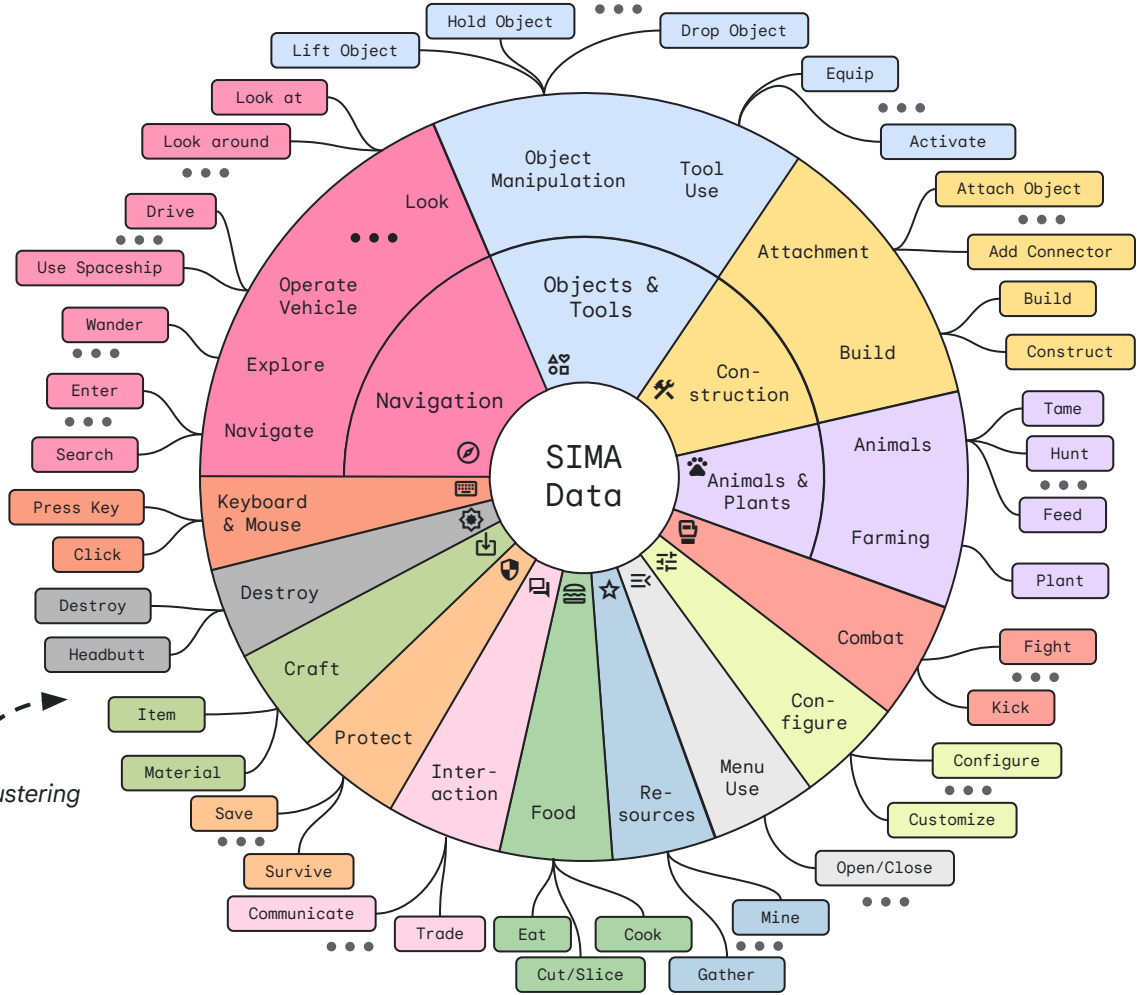

SIMA *Agent*

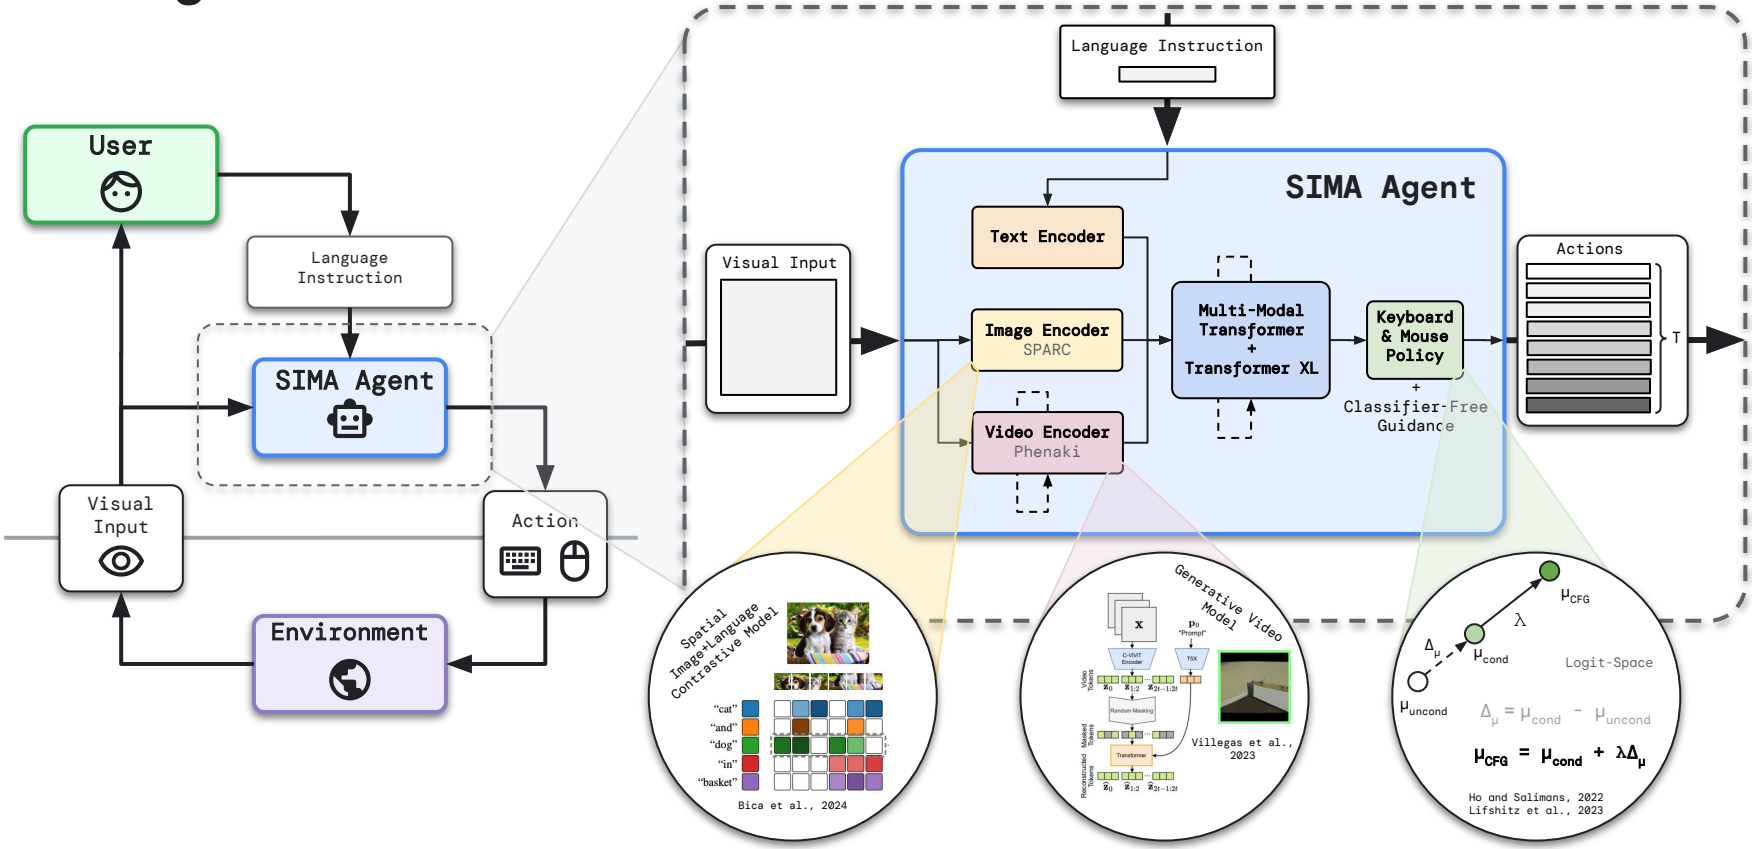
User
Language Instruction
SIMA Agent
Visual Input
Action
Environment
Language Instruction
Visual Input
SIMA Agent
Text Encoder
Image Encoder
SPARC
Video Encoder
Phenaki
Multi-Modal Transformer + Transformer XL
Keyboard & Mouse Policy
+
Classifier-Free Guidance
Actions
T
Spatial Image+Language Contrastive Model
"cat"
"and"
"dog"
"in"
"basket"
Bica et al., 2024
Generative Video Model
Villegas et al., 2023
Logit-Space
Ho and Salimans, 2022
Lifshitz et al., 2023

SIMA ***Evaluation***

Eval Basics

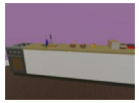

starting state

"put the frying pan on the stove"

instruction

* *Ground Truth*
* *OCR*
* *Human Eval*

→ {True, False}

success criterion

"Task"

$Task_0$ $Task_1$ ••• $Task_N$

"Challenge"

Ground Truth

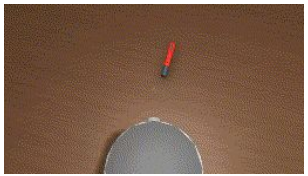

`marker.y_pos > 0` → *True*

Accurate, Automatic

Research Environments

Optical Character Recognition (OCR)

`OCR == 'raspberries'` → *True*

Reasonably Accurate, Automatic

Text-Based State Changes

Human Evaluation

`agent == upstairs` → *True*

Flexible / Widely Applicable

Costly

02 Results

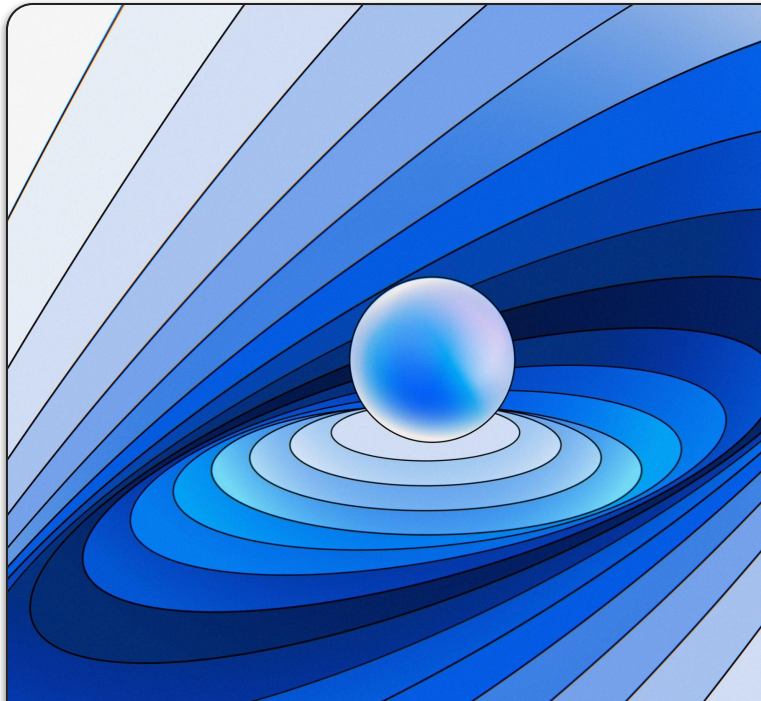

Qualitative Results - *Domain-General Behaviors*

3D First-Person Environments + Shared Interface → *Commonalities in* ***Spatial Navigation*** *&* ***Interaction***

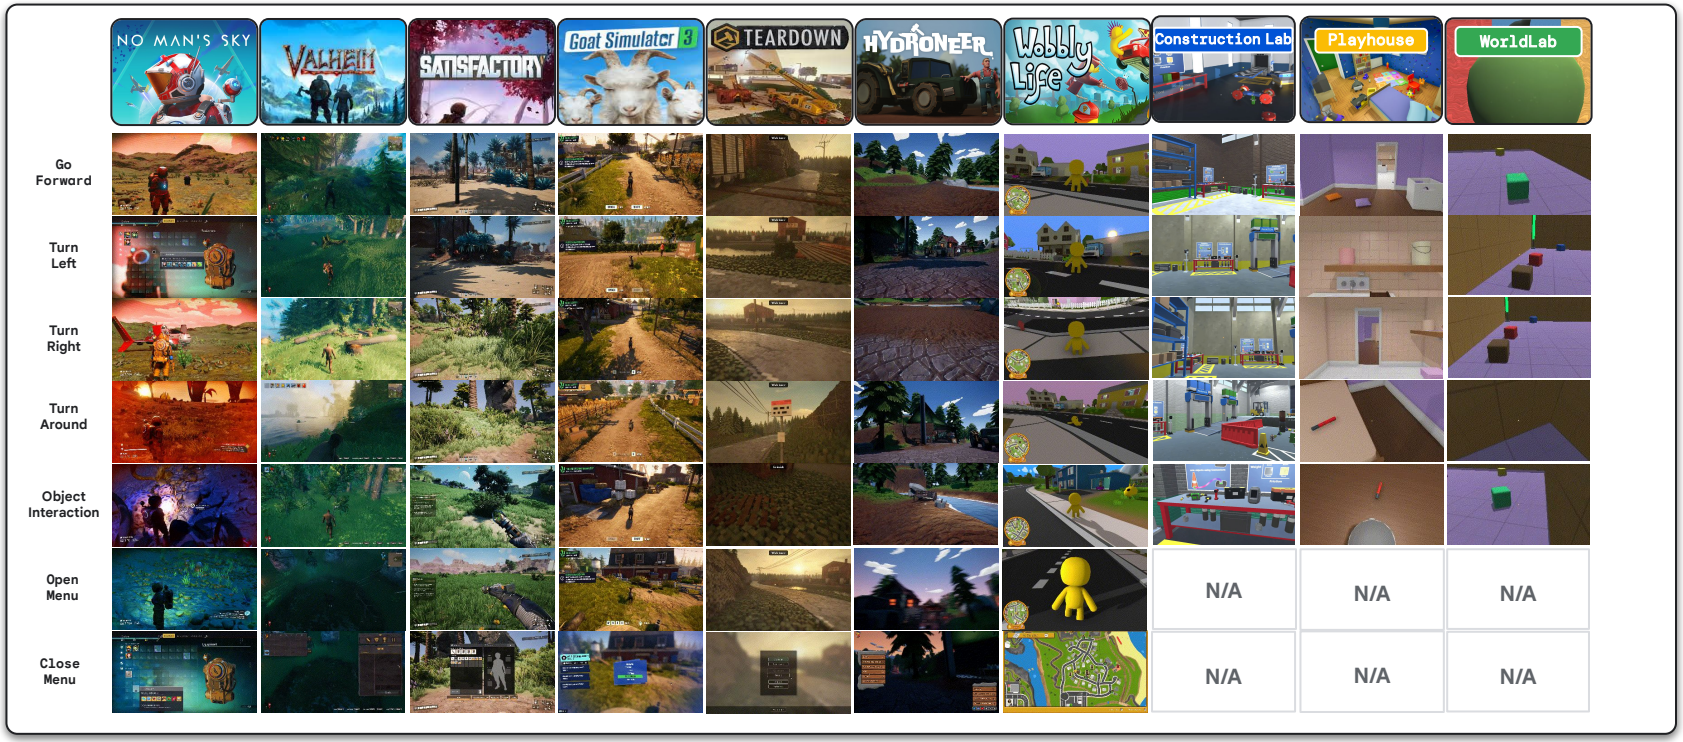

Qualitative Results - *Domain-Specific Behaviors*


Go to the Spaceship


Feed the Creature

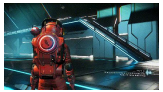
Use the Jetpack to Go Upstairs


Shoot Asteroids


Pick Up Mushrooms


Chop Down a Tree


Craft a Club


Fight the Greyling


Climb the Ladder


Drive the Tractor


Find a Green Object


Go to the Goat Castle


Pick Up Leaves

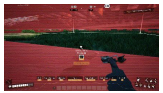
Build a HUB


Mine Copper Ore


Open the Craft Bench Menu

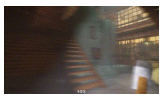
Go Upstairs


Go Through the Gate


Get in the Boat


Get in the Blue Car


Go Through the Bouncy Castle

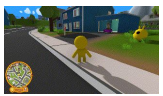
Get in the Purple Car


Talk to the Person


Click 'Accept Job'


Get in the Truck


Pick up the Bucket


Pick up a Tool Bag


Get the Wooden Cart


Pick Up a Box


Pick Up a Motor Joint

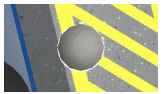
Find the Connector Tool


Attach a Blue Sphere to a Brown Cube

Qualitative Results - *Commonalities Across Domains*

Go to / get in a vehicle

Go to the Spaceship

Drive the Tractor

Get in the Blue Car

Get in the Purple Car

Get in the Boat

Get in the Truck

Absolute Performance – *By Environment*

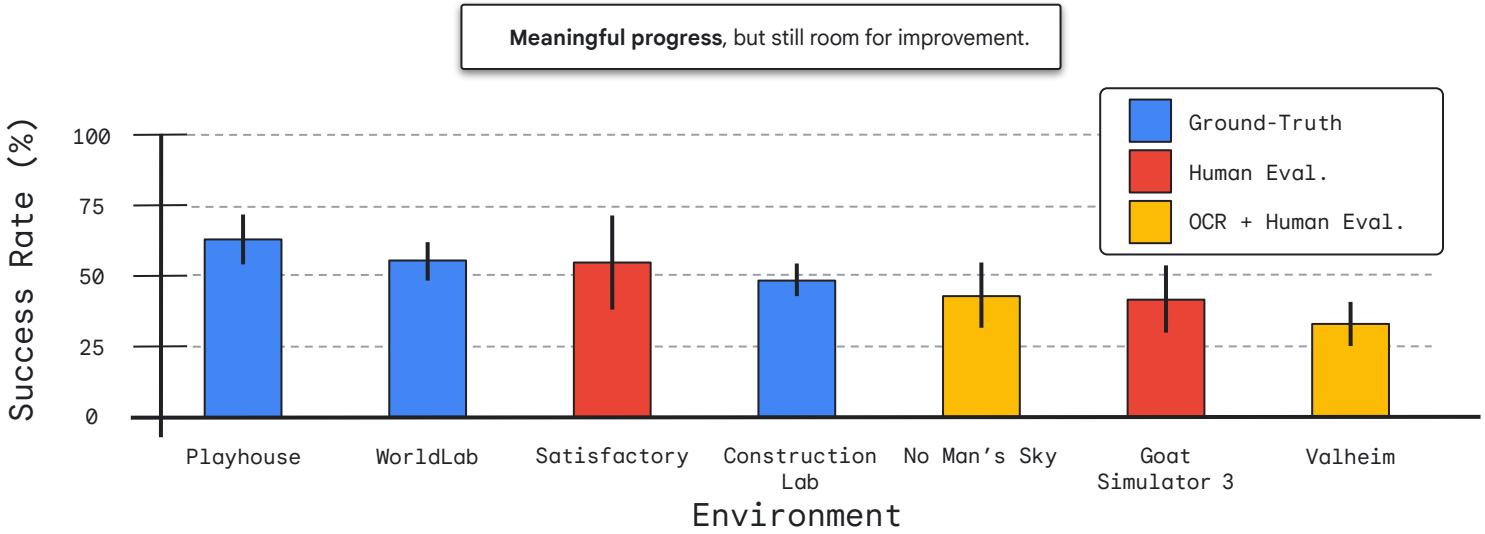

Absolute Performance – *By Skill Category*

Hypothesis: The agent struggles with tasks that require more precise control.

Ablations / Baselines

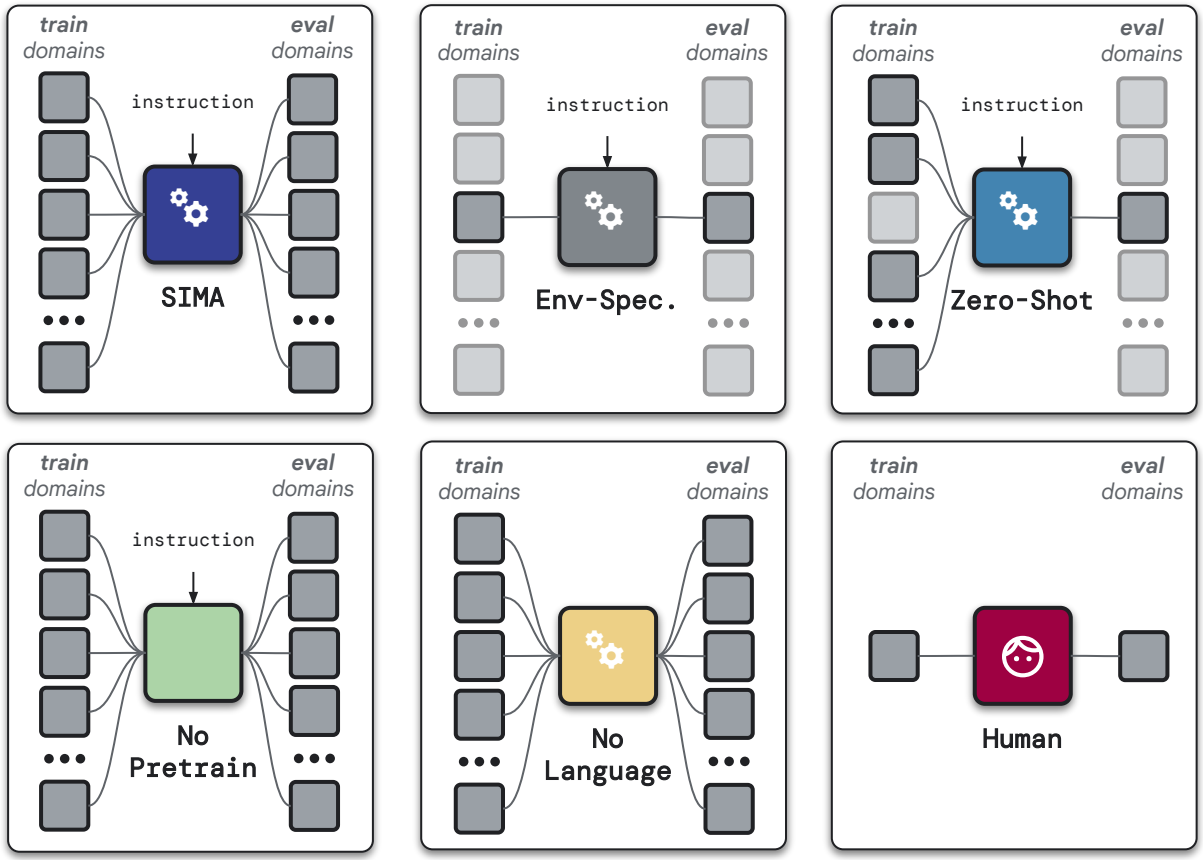

Relative Performance – *Aggregated Over Environments*

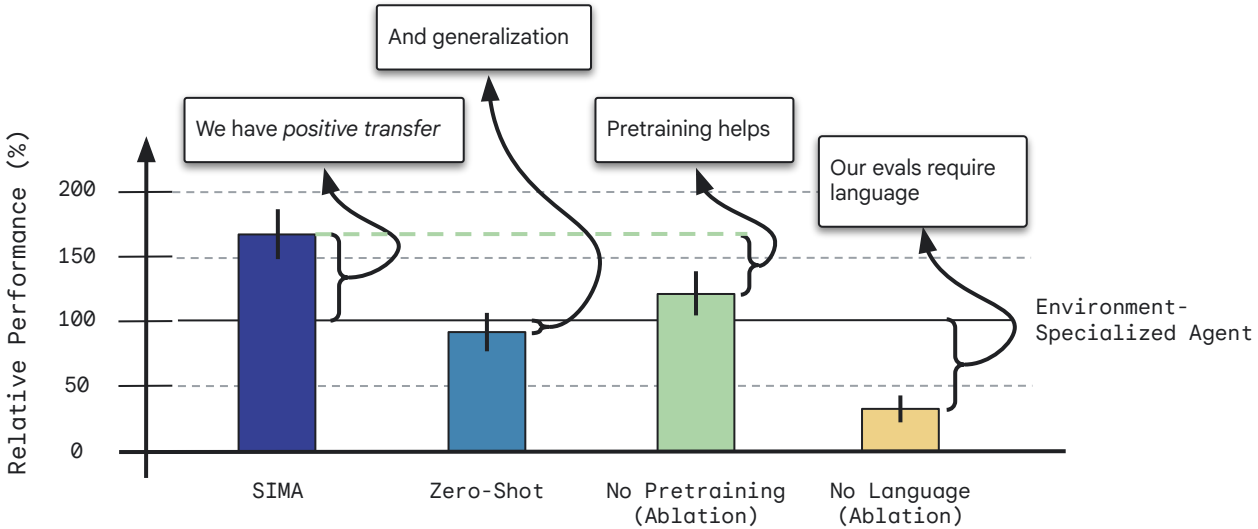

Relative Performance – *Per-Environment*

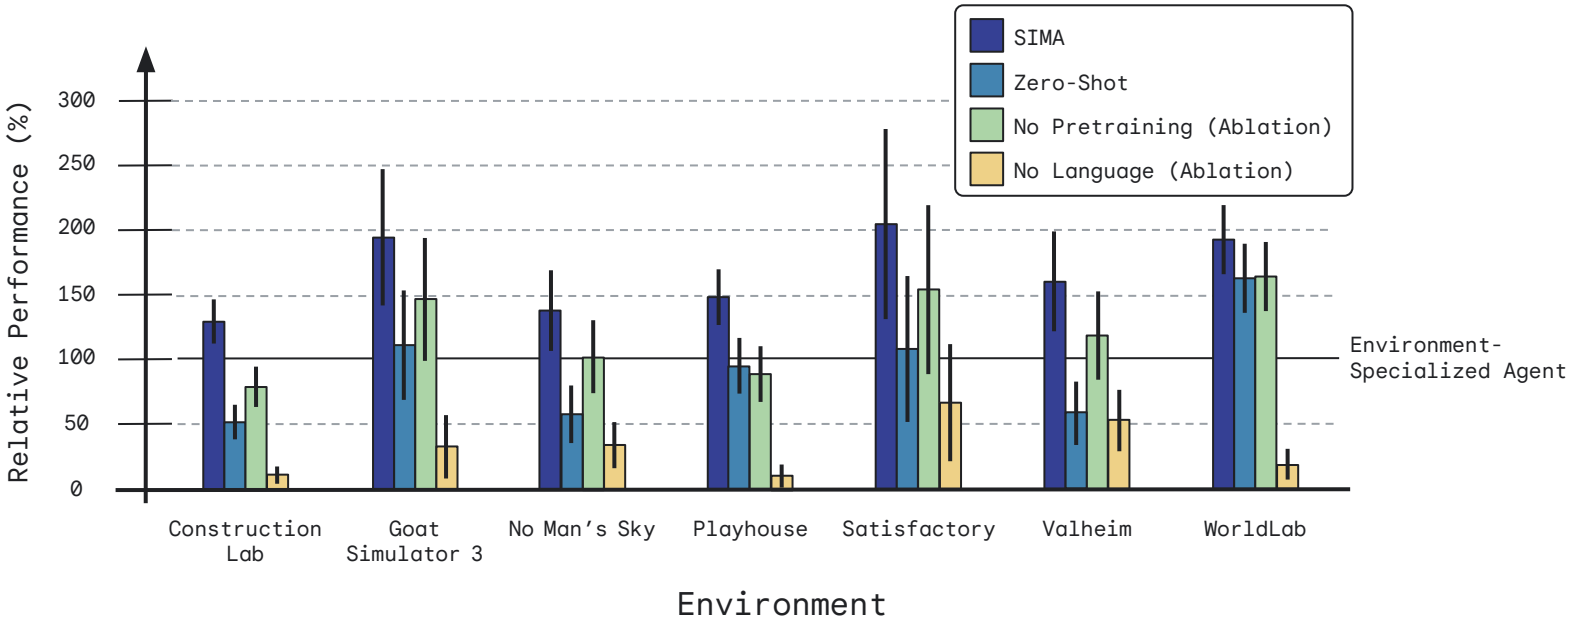

Relative Performance – *Per-Environment*

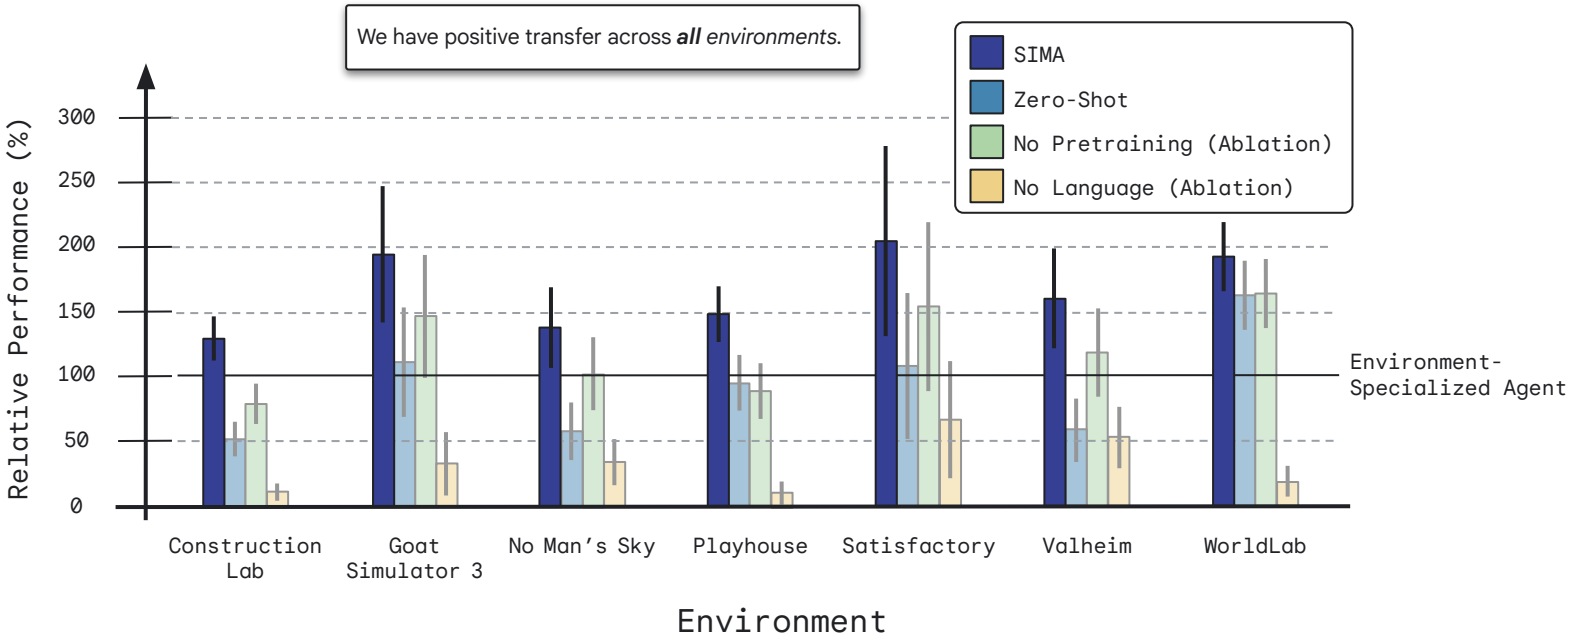

Relative Performance – *Per-Environment*

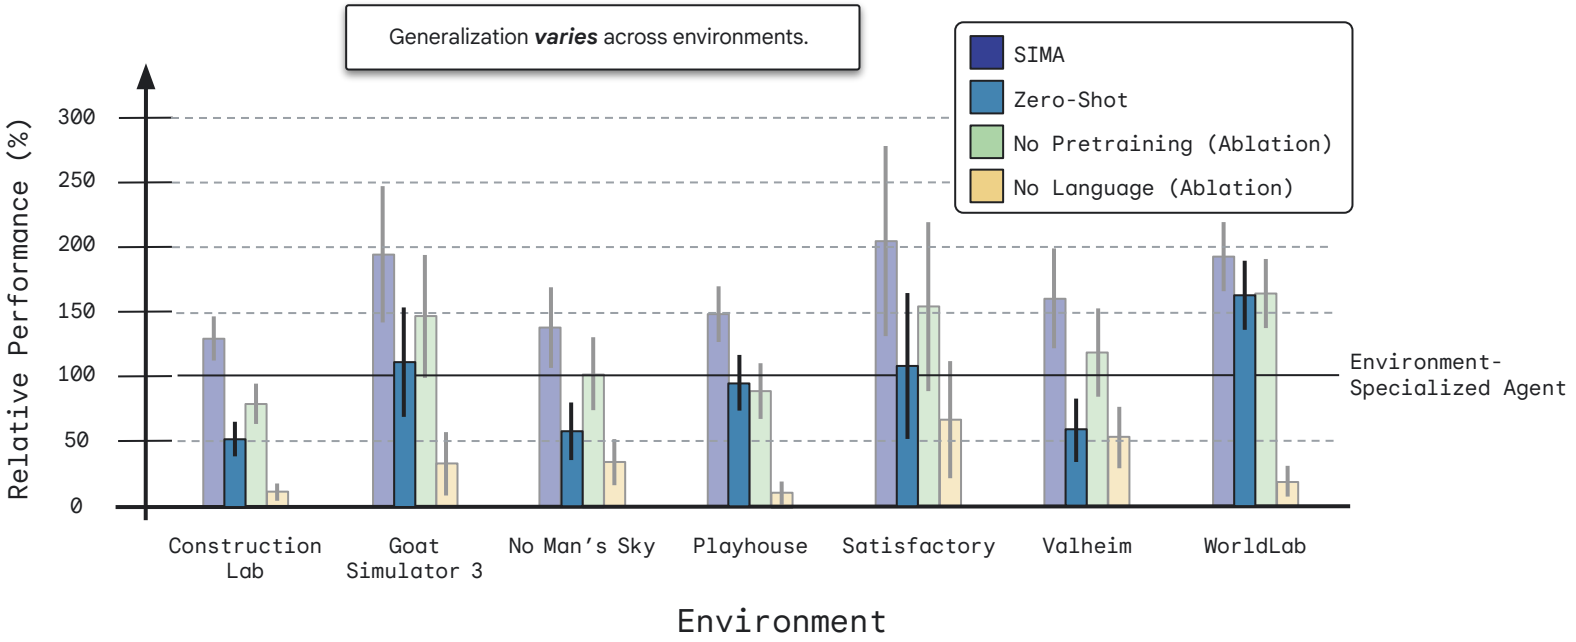

Relative Performance – *Per-Environment*

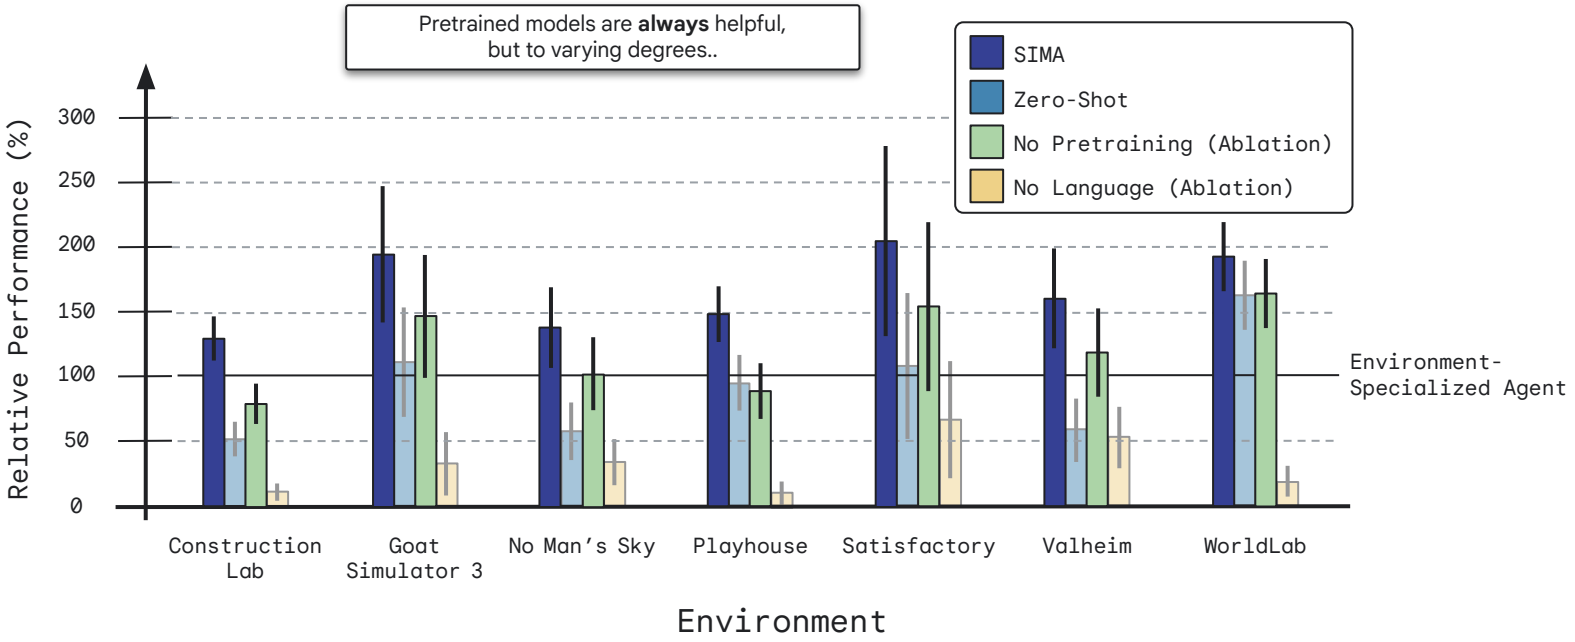

Relative Performance – *Per-Environment*

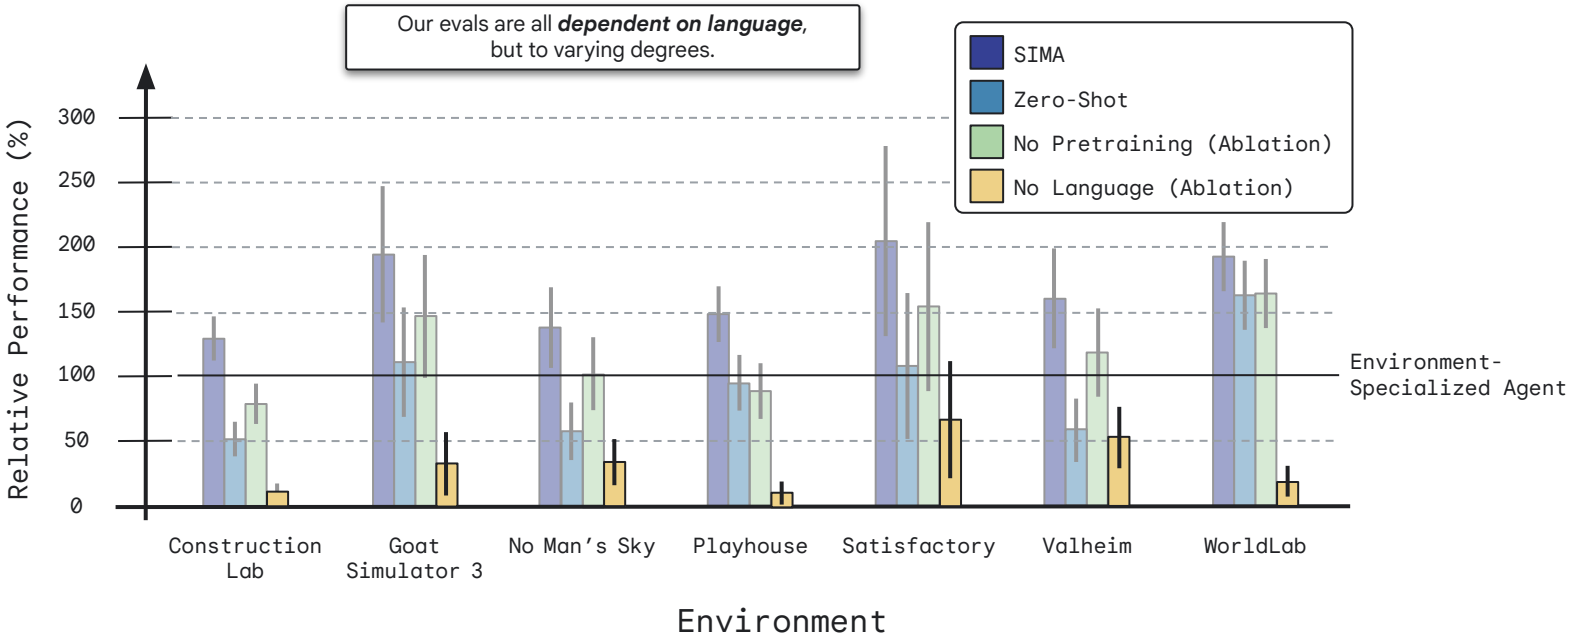

Classifier-Free Guidance

Evaluated on Playroom, Construction Lab, and WorldLab only.

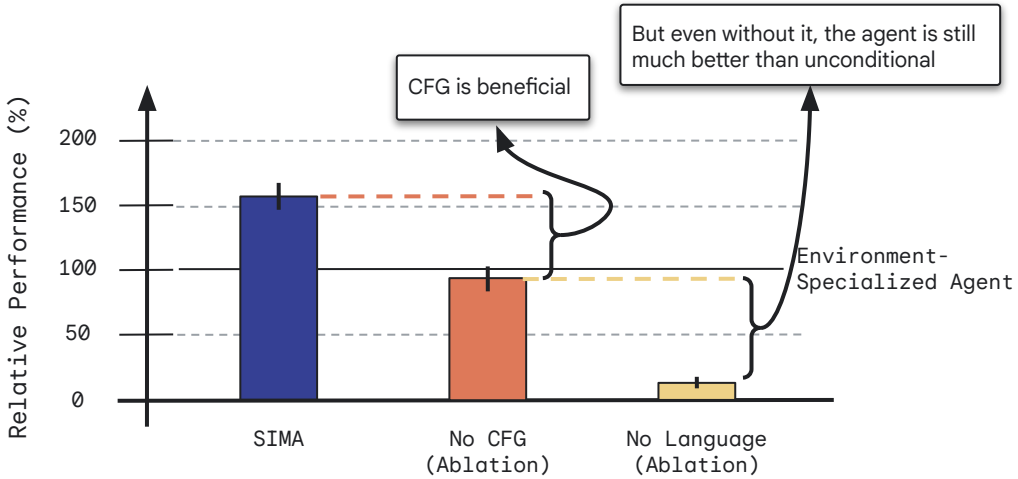

Human Baseline Comparison

Evaluated on a subset of No Man's Sky only.

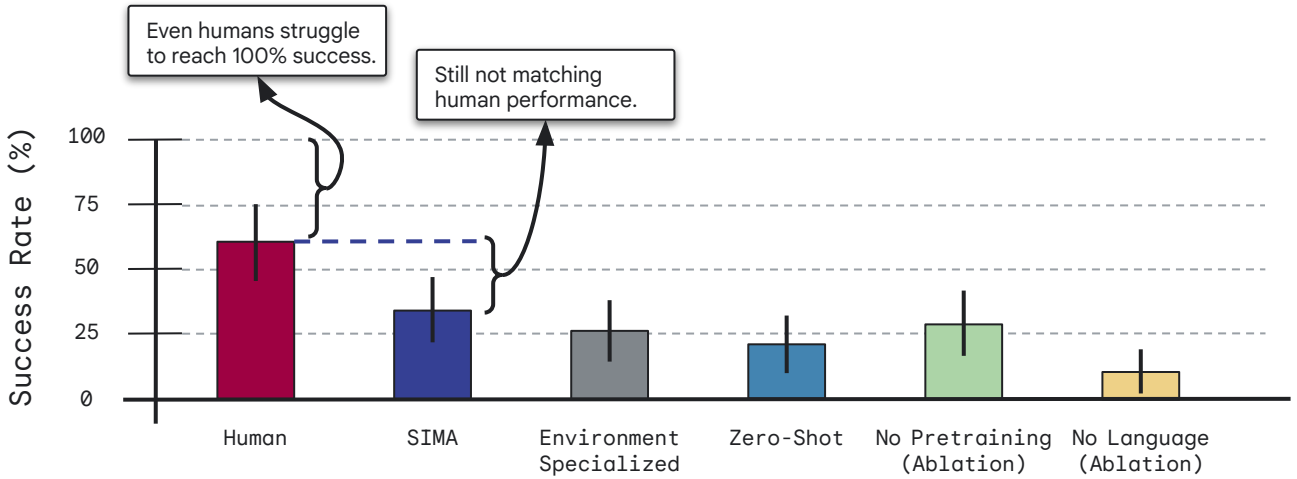

03 Closing Remarks

Looking Forward

*If we want to achieve the 'north star' of **general-purpose humanoid robots**...*

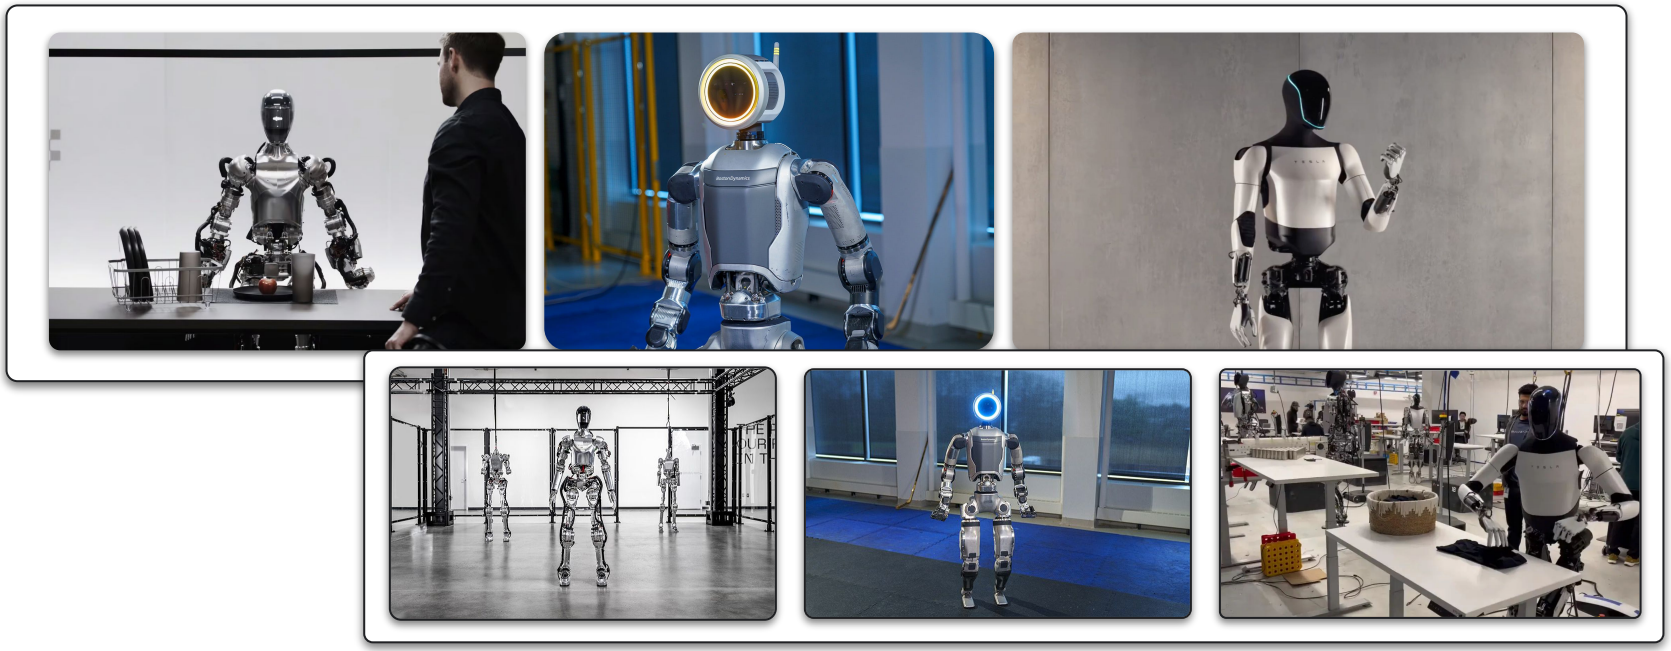

*...then overcoming the challenges of basic research may involve **relying, in part, on simulation**.*

Looking Forward

*Previous works have **lacked diversity** in their...*

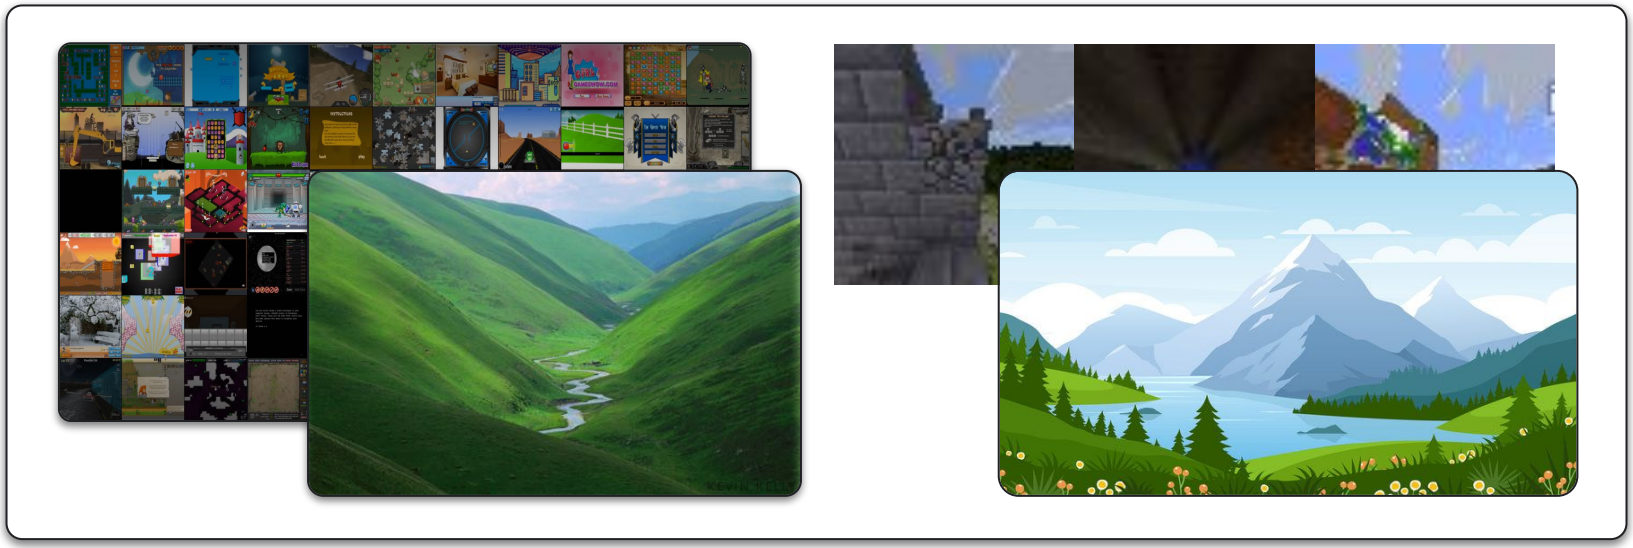

...trajectories (tasks)...

...and/or their observations / affordances.

Looking Forward

SIMA is the next generation *in simulation-based agent research...*

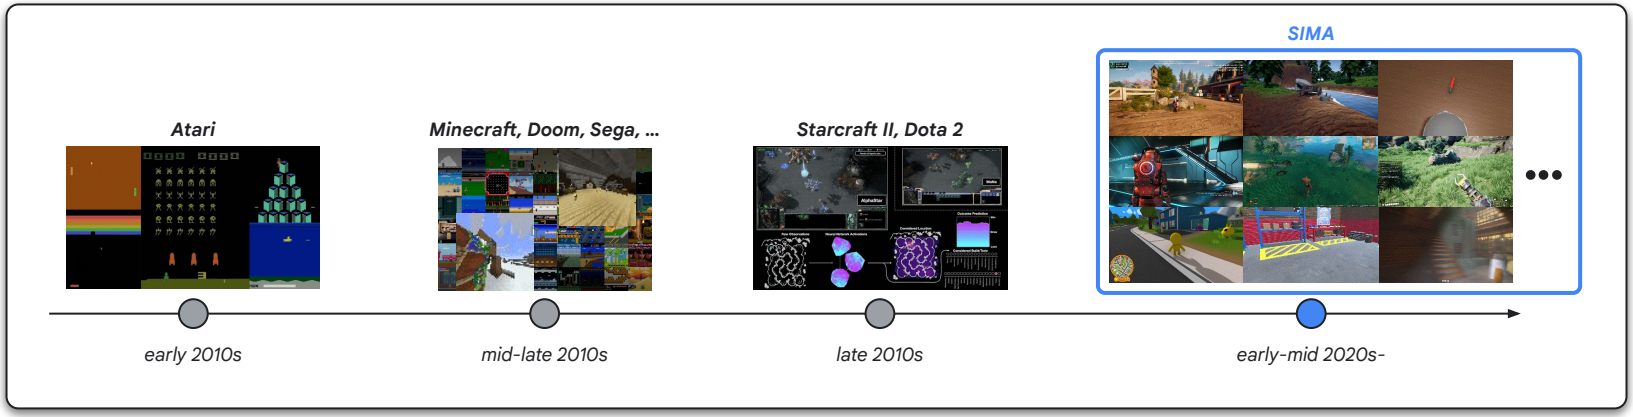

Looking Forward

*...and advancing simulation-based agent research is likely to yield meaningful **insights into creating AGI**.*

The SIMA agent is at the "pre-training" phase of general-purpose behavior.

*We are getting **glimpses of impressive general capabilities, but not yet a fully general-purpose agent**.*

Analogous to GPT2/3-era models.

c. 2019

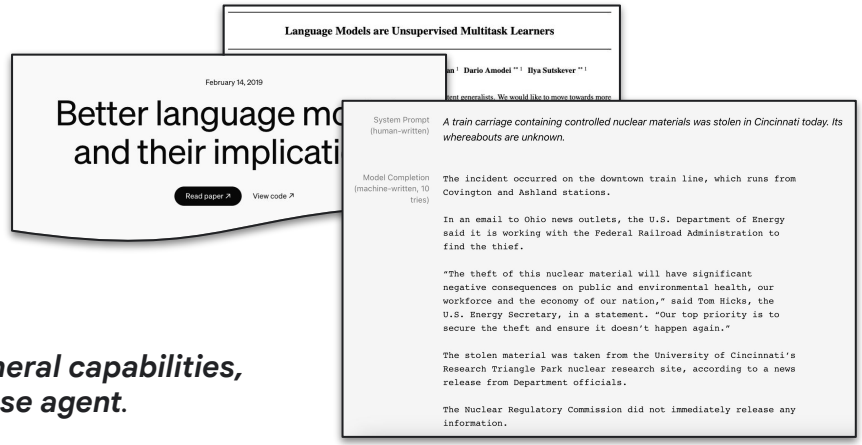

Looking Forward

*...and advancing simulation-based agent research is likely to yield meaningful **insights into creating AGI**.*

The SIMA agent is at the "pre-training" phase of general-purpose behavior.

Analogous to GPT2/3-era models.

c. 2019

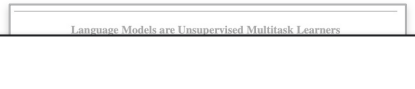
Language Models are Unsupervised Multitask Learners

*We may be a handful of innovations away from developing **general-purpose agents, capable of performing any task in any simulated 3D environment**.*

*We are getting **glimpses of impressive general capabilities, but not yet a fully general-purpose agent**.*

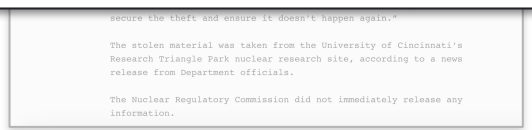
secure the theft and ensure it doesn't happen again."

The stolen material was taken from the University of Cincinnati's Research Triangle Park nuclear research site, according to a news release from Department officials.

The Nuclear Regulatory Commission did not immediately release any information.

If [an agent] is able to master 10,000 diverse simulated realities, it may well generalize to our physical world, which is simply the 10,001st reality.

Jim Fan

The SIMA Team

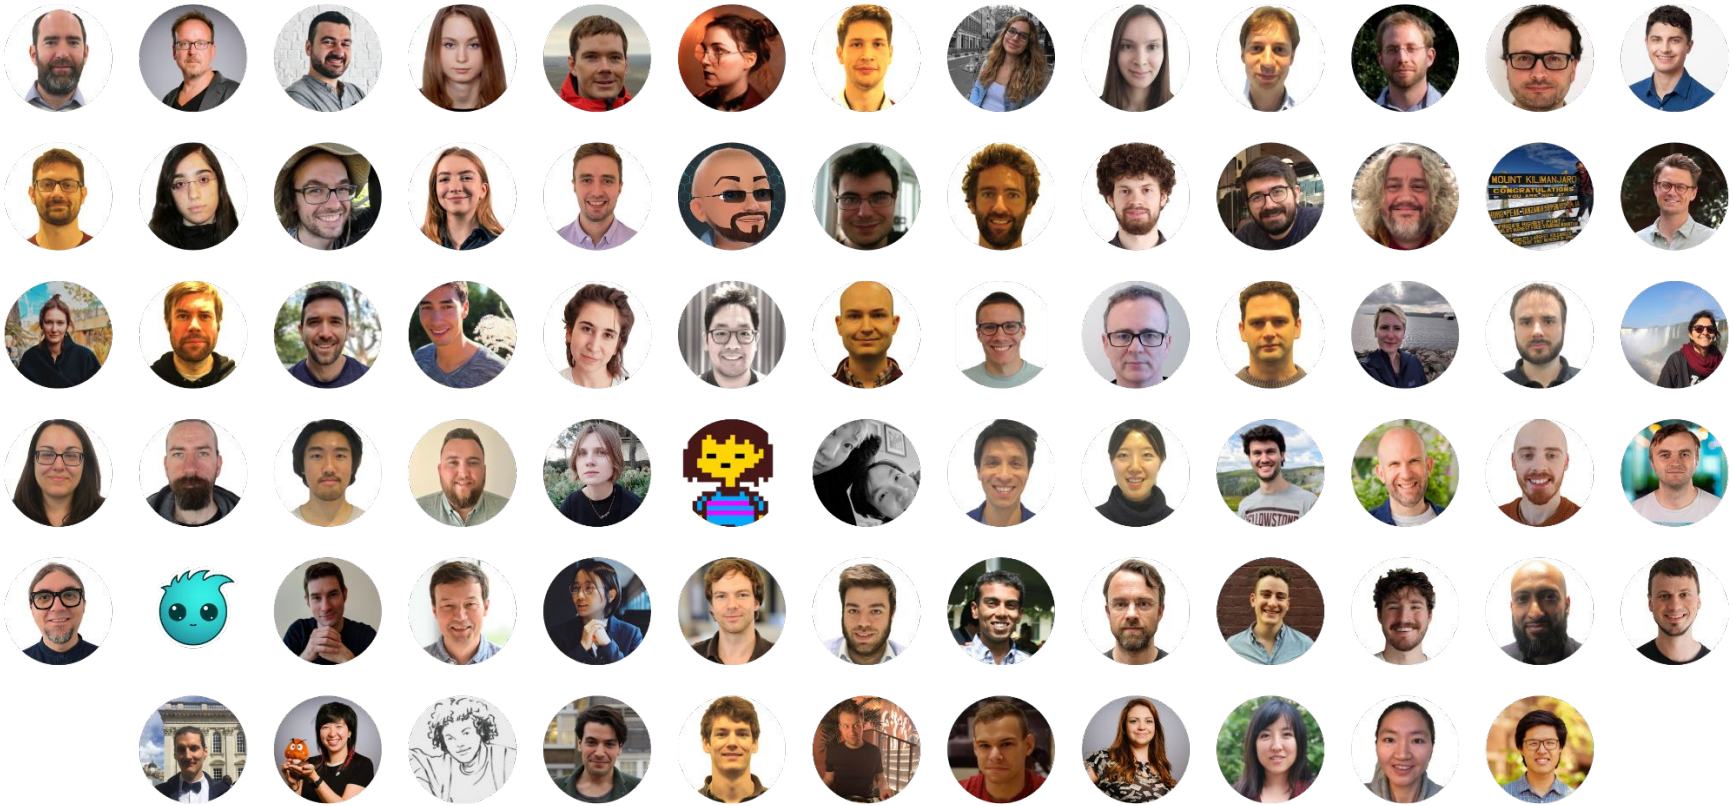

Tech Report + Blog Post

Google DeepMind

Scaling Instructable Agents Across Many Simulated Worlds

SIMA Team:[1] Maria Abi Raad, Arun Ahuja, Catarina Barros, Frederic Besse, Andrew Bolt, Adrian Bolton, Bethanie Brownfield, Gavin Buttimore, Max Cant, Sarah Chakera, Stephanie C. Y. Chan, Jeff Clune[1,3], Adrian Collister, Vikki Copeman[2], Alex Cullum, Ishita Dasgupta, Dario de Cesare, Julia Di Trapani, Yani Donchev, Emma Dunleavy, Martin Engelcke, Ryan Faulkner, Frankie Garcia, Charles Gbadamosi, Zhitao Gong, Lucy Gonzales[2], Karol Gregor, Kshitij Gupta[2], Arne Olav Hallingstad, Tim Harley, Sam Haves, Felix Hill, Ed Hirst, Drew A. Hudson, Jony Hudson, Steph Hughes-Fitt, Danilo J. Rezende, Mimi Jasarevic, Laura Kampis, Rosemary Ke, Thomas Keck, Junkyung Kim, Oscar Knagg, Kavya Kopparapu, Andrew Lampinen, Shane Legg, Alexander Lerchner, Marjorie Limont, Yulan Liu, Maria Loks-Thompson, Joseph Marino, Kathryn Martin Cussons[2], Loic Matthey, Siobhan Mcloughlin, Piermaria Mendolicchio, Hamza Merzic, Anna Mitenkova, Alexandre Moufarek, Valeria Oliveira, Yanko Oliveira, Hannah Openshaw, Renke Pan, Aneesh Pappu, Alex Platonov, Ollie Purkiss, David Reichert, John Reid, Pierre Harvey Richemond, Tyson Roberts, Giles Ruscoe, Jaume Sanchez Elias, Tasha Sandars[2], Daniel P. Sawyer, Tim Scholtes, Guy Simmons, Daniel Slater, Hubert Soyer, Heiko Strathmann, Peter Stys, Allison C. Tam[2], Denis Teplyashin, Tayfun Terzi, Davide Vercelli, Bojan Vujatovic, Marcus Wainwright, Jane X. Wang, Zhengdong Wang, Daan Wierstra[2], Duncan Williams, Nathaniel Wong, Sarah York, Nick Young

[1]Google DeepMind unless otherwise noted, authors listed in alphabetical order, contributions listed at end of report, [2]work performed while at Google DeepMind, [3]University of British Columbia

Building embodied AI systems that can follow arbitrary language instructions in any 3D environm… a key challenge for creating general AI. Accomplishing this goal requires learni… in perception and embodied actions, in order to accomplish comple… Multiworld Agent (SIMA) project tackles this by training ag… a diverse range of virtual 3D environments, includi… commercial video games. Our goa…

[arXiv:2404.10179](https://arxiv.org/abs/2404.10179)

Google DeepMind Blog

RESEARCH

A generalist AI agent for 3D virtual environments

13 MARCH 2024

By the SIMA Team

Share

We present new research on a Scalable Instructable Multiworld… Agent (SIMA) that can follow natural-language instruct… out tasks in a variety of video game settings.

Video games are a key proving ground for … the real world, games are rich learni… settings and ever-changi…

Thank you.

Joe Marino
Research Scientist
josephmarino@google.com

Google DeepMind