Simulation-based design
in games, user interfaces and beyond

Prof. Perttu Hämäläinen
Aalto SCI & Aalto ARTS
perttu.hamalainen@aalto.fi

My Focus: Movement & Games

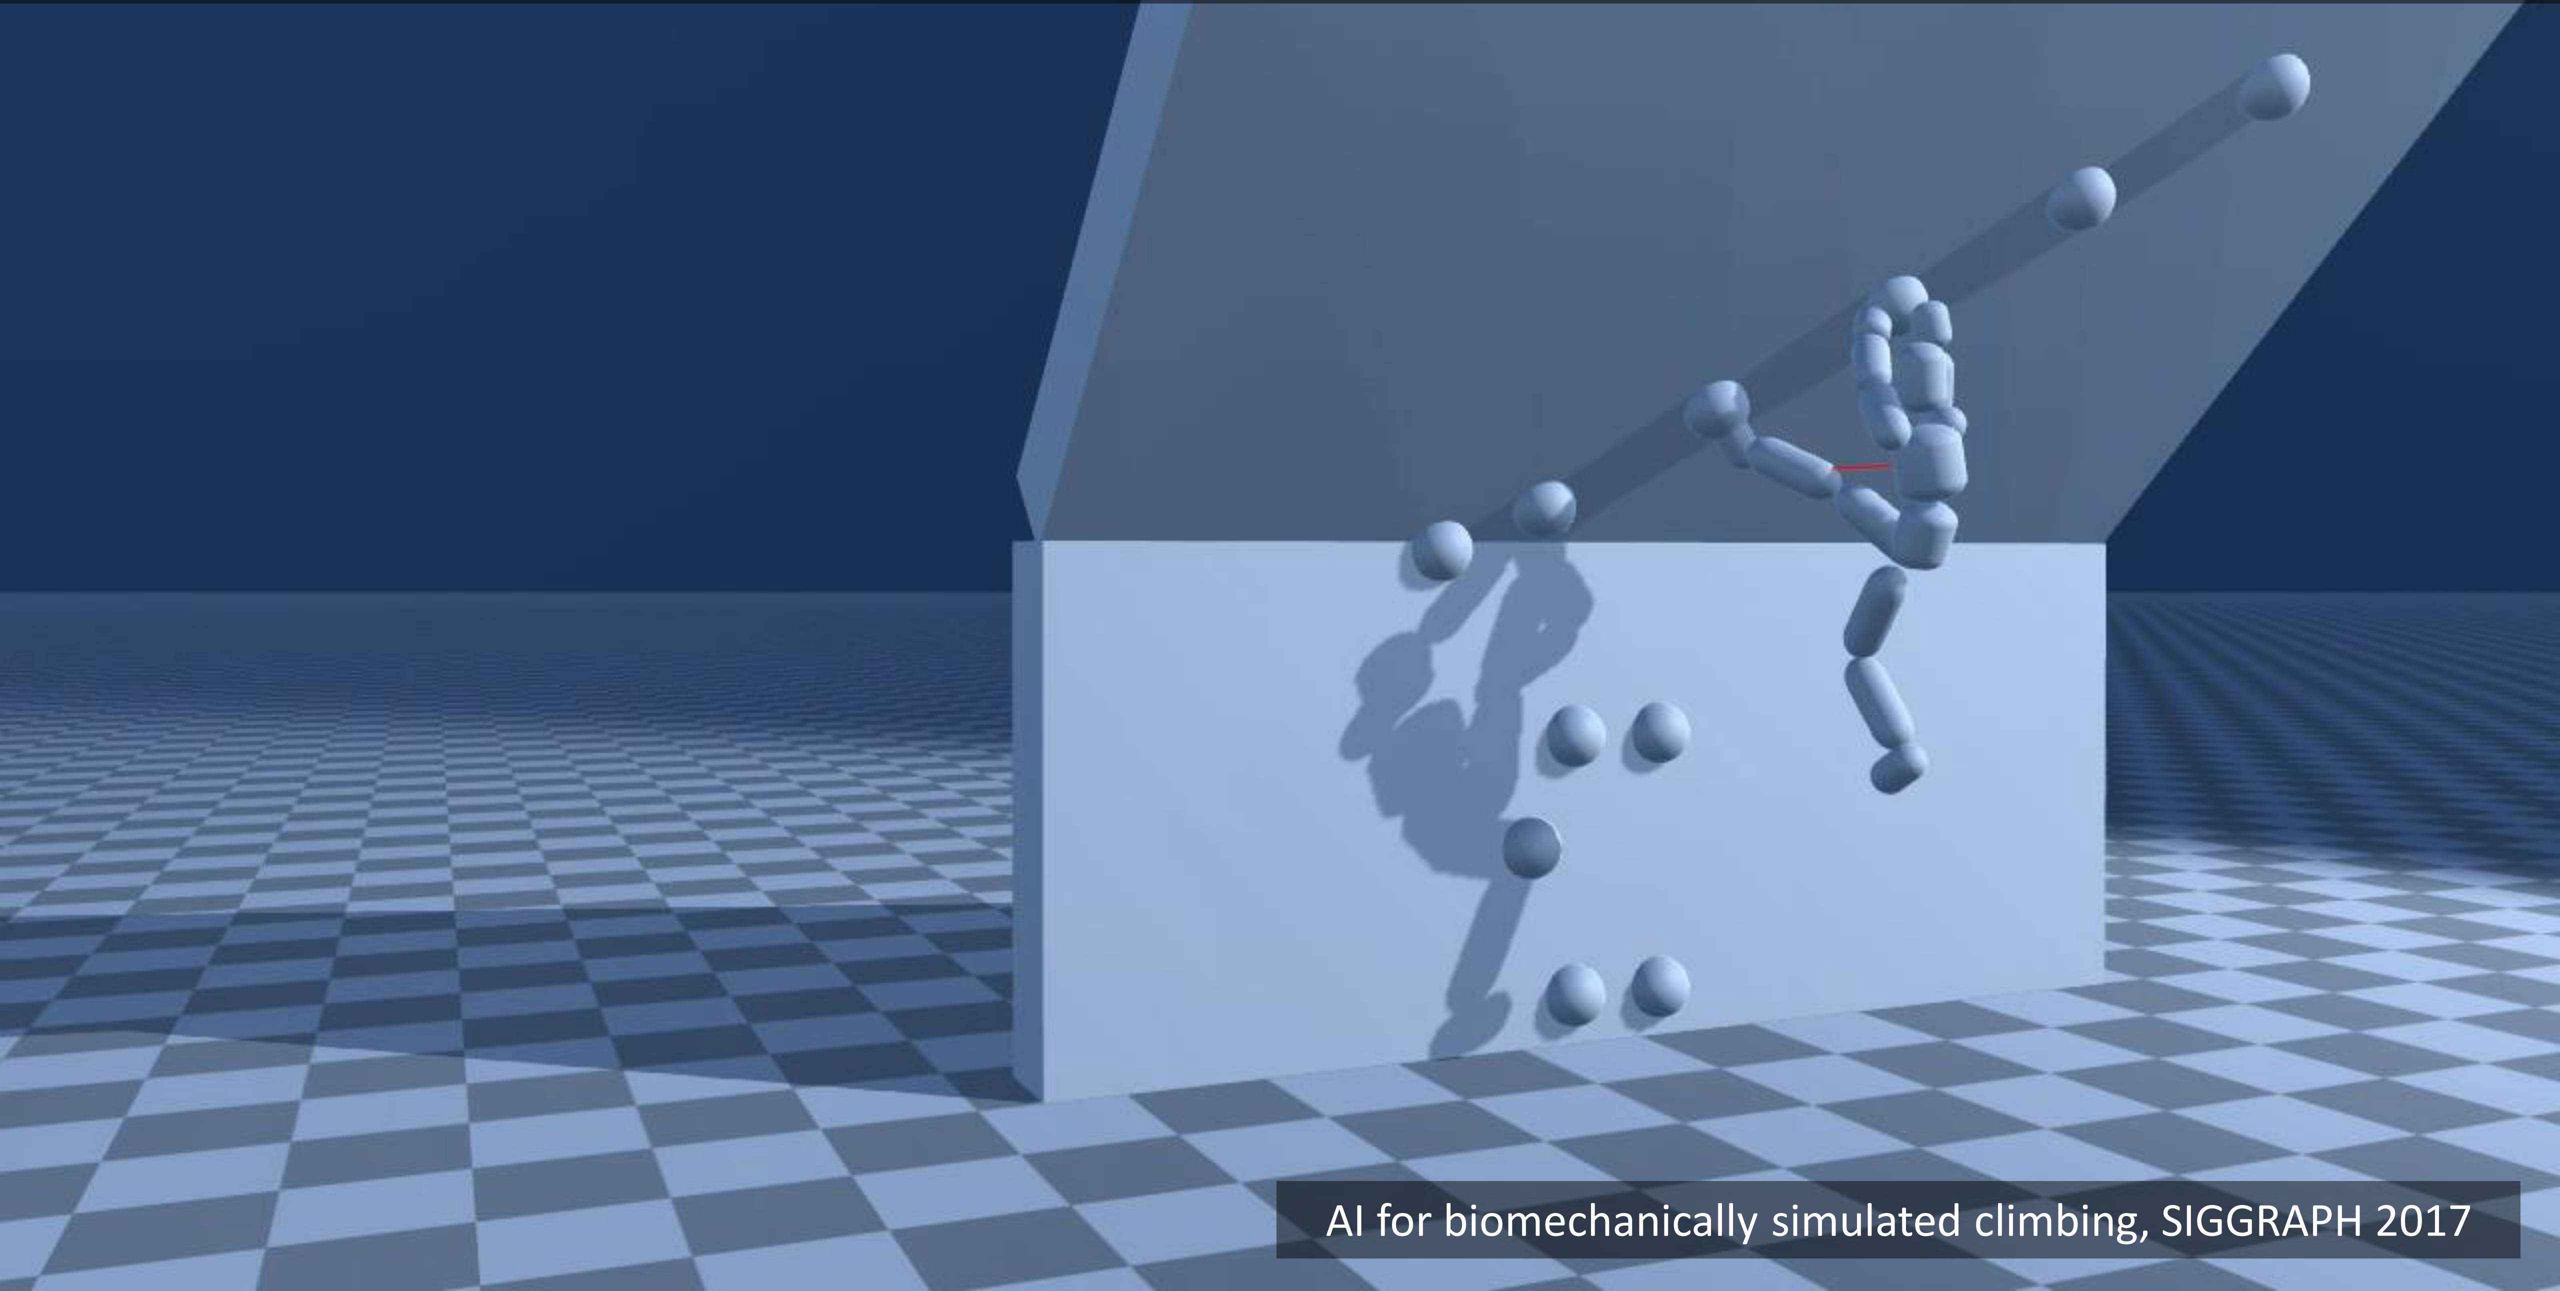

AI for biomechanically simulated climbing, SIGGRAPH 2017

My Focus: Movement & Games

ValoClimb, a.k.a., Augmented Climbing Wall, CHI 2016

Locations

Our products can already be found in almost **300** locations and in over **50** countries. Check the map below to find the closest one to you!

How to produce good design reliably and repeatedly?

Design is iterative, i.e., slow and expensive

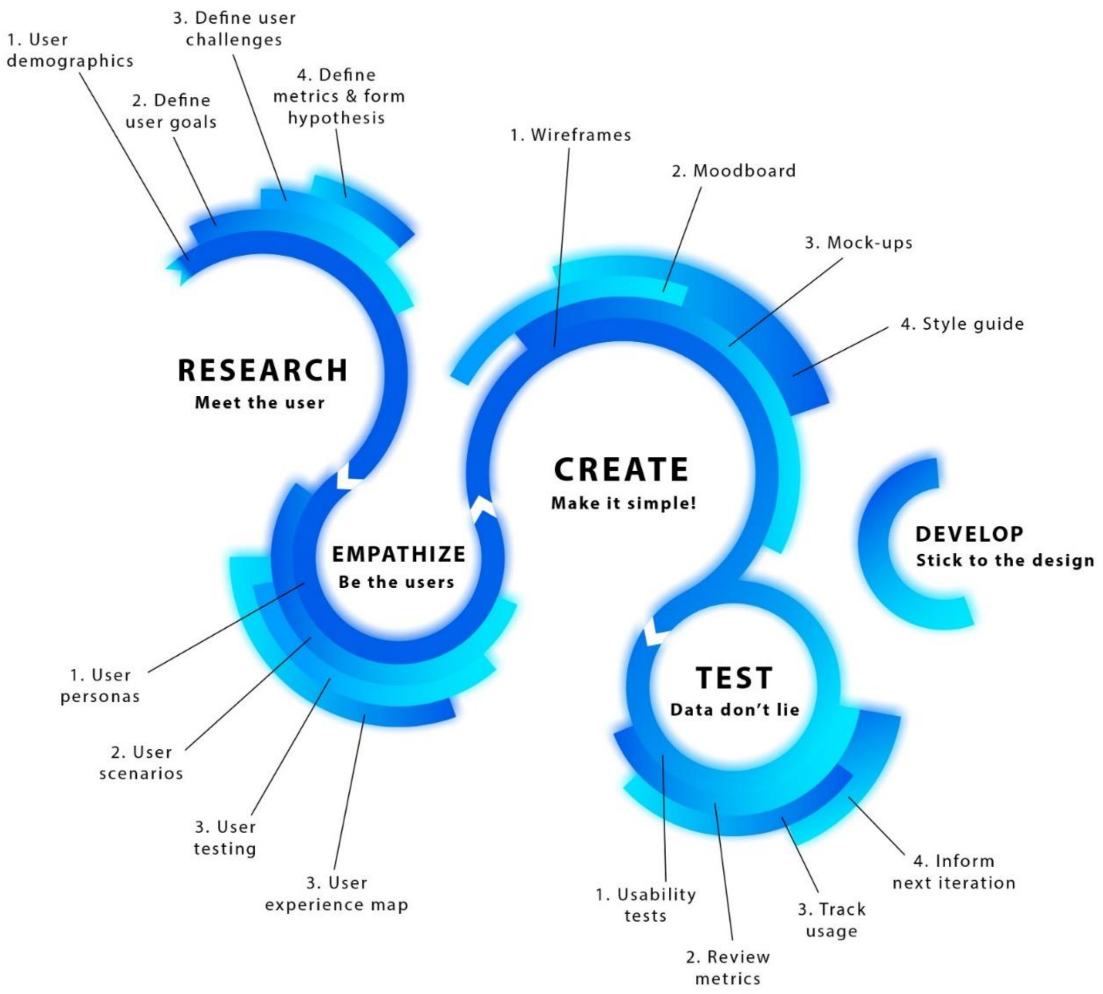

https://medium.com/@alexwyrick/ux-design-process-a66e6796857e

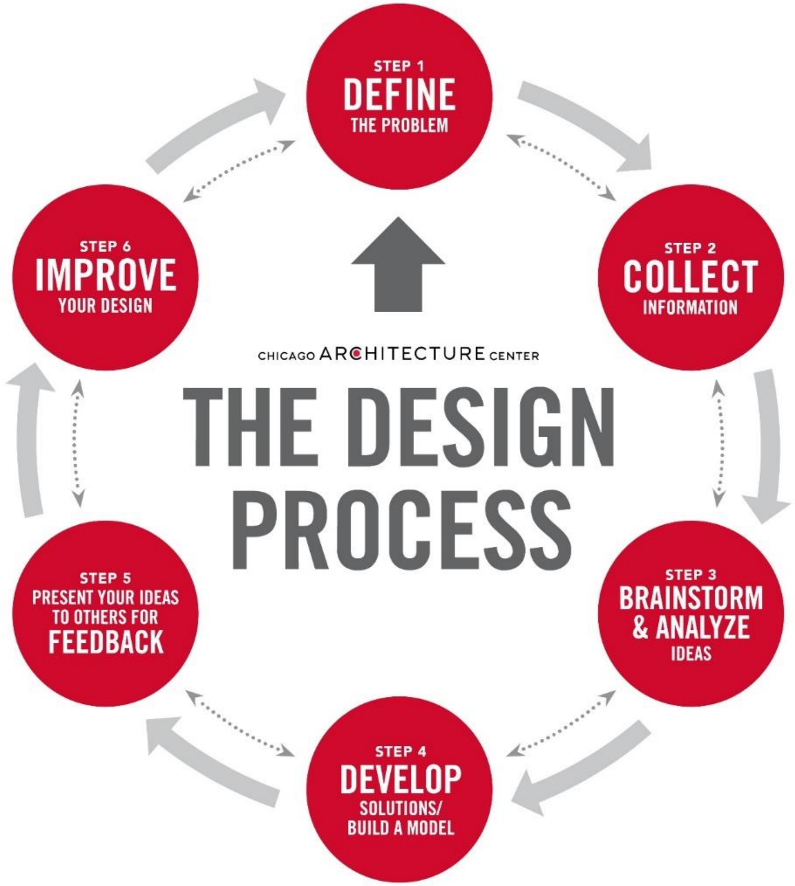

https://www.discoverdesign.org/handbook

Users are unpredictable

BLOG 08.02.2019

DIGITAL TWIN IS THE NEW BLACK

What used to take weeks or months, can now be done in minutes or hours. Watch our video to see the beauty of Digital Twin prototyping and testing.

Digitalization | New business models | Startup | SME | Large company | Research organization | Finland | ICT & Digitalization | Connected Intelligent Industries Finland | Digital Trust Finland

AUTHOR

REIJO SMOLANDER

DIGITAL TWINS TO SUPPORT YOUR WHOLE PRODUCT LIFECYCLE

Develop a unique digital twin to support your specific needs throughout the whole product lifecycle. Create the virtual model already at inception, utilize the same model in the sales and marketing of the product, start training earlier on, and eventually connect the real machine to the digital twin.

Problem: Still Slow and Expensive

Future of simulation-based design:

Digital twin simulation of both ***designs*** and ***users***

How to simulate and design?

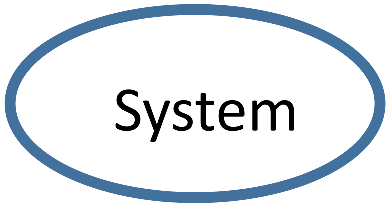
System


User
System

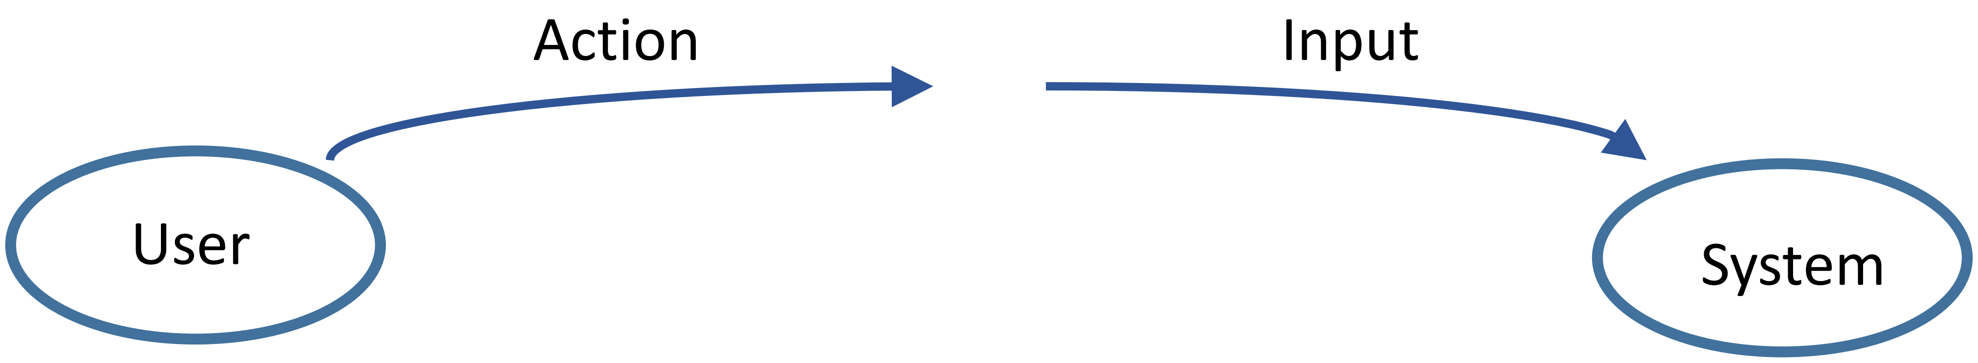
Action
Input
User
System

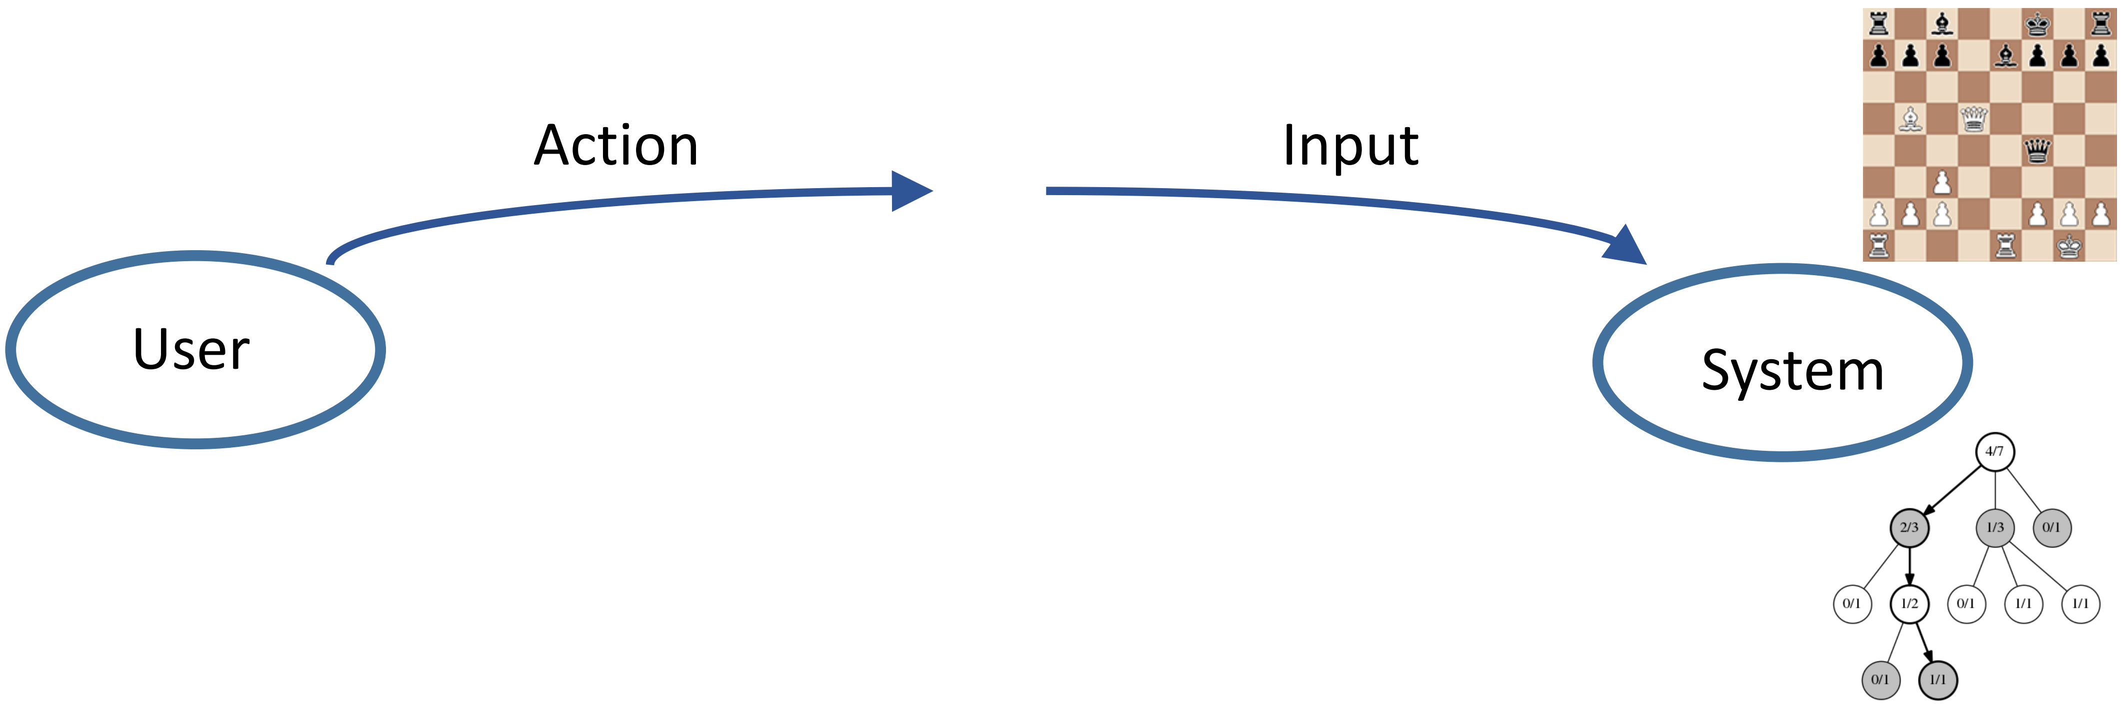
Action
Input
User
System

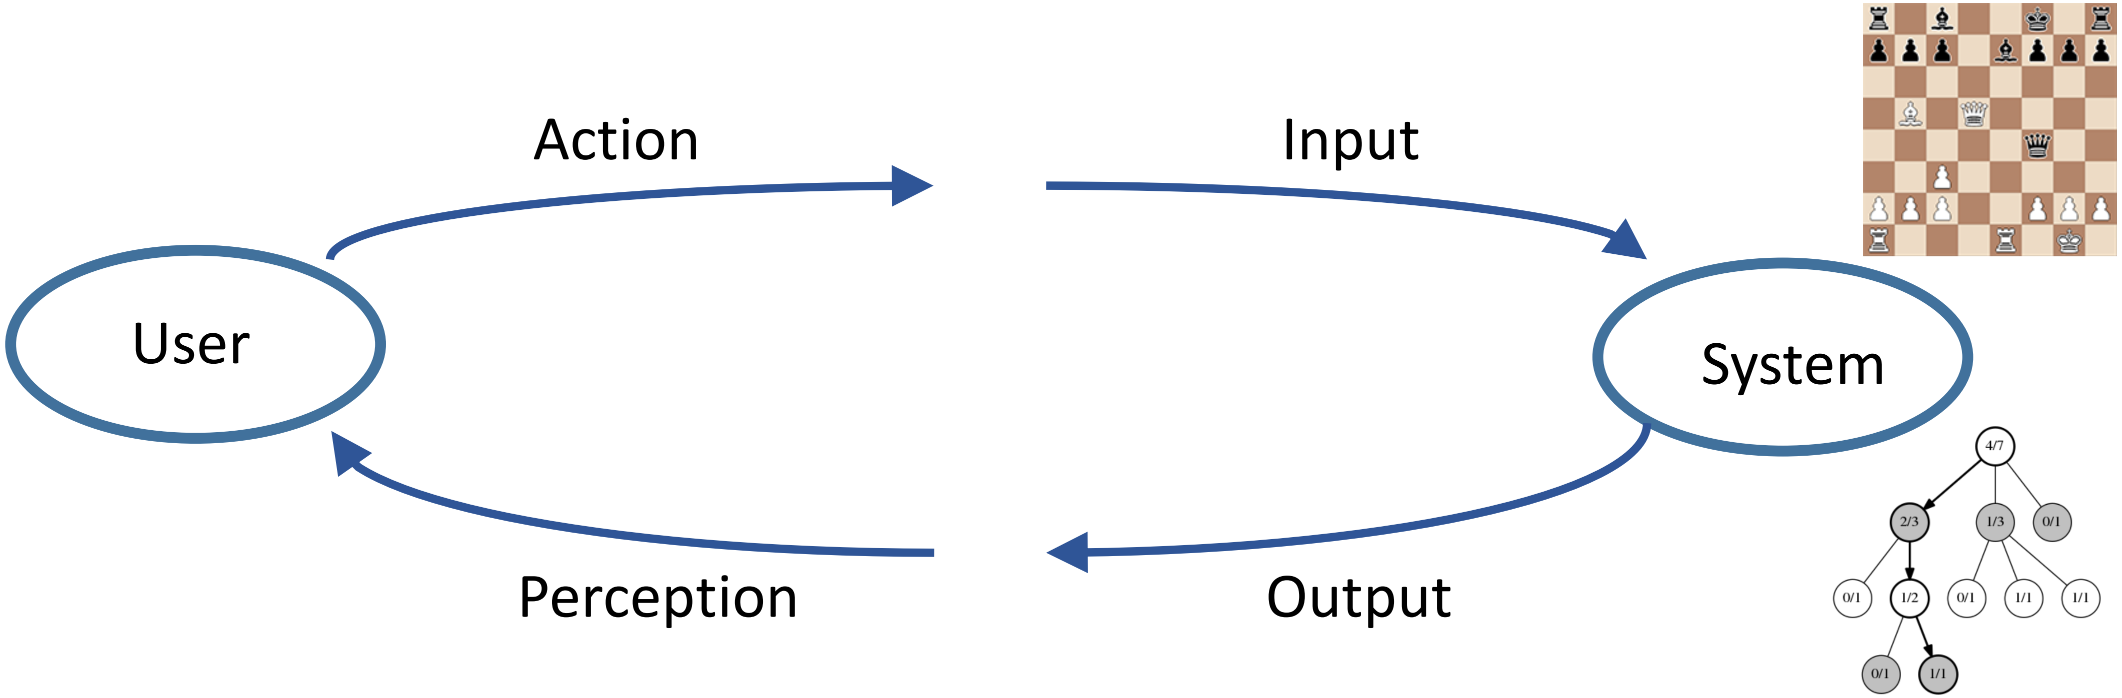
Action
Input
User
System
Perception
Output

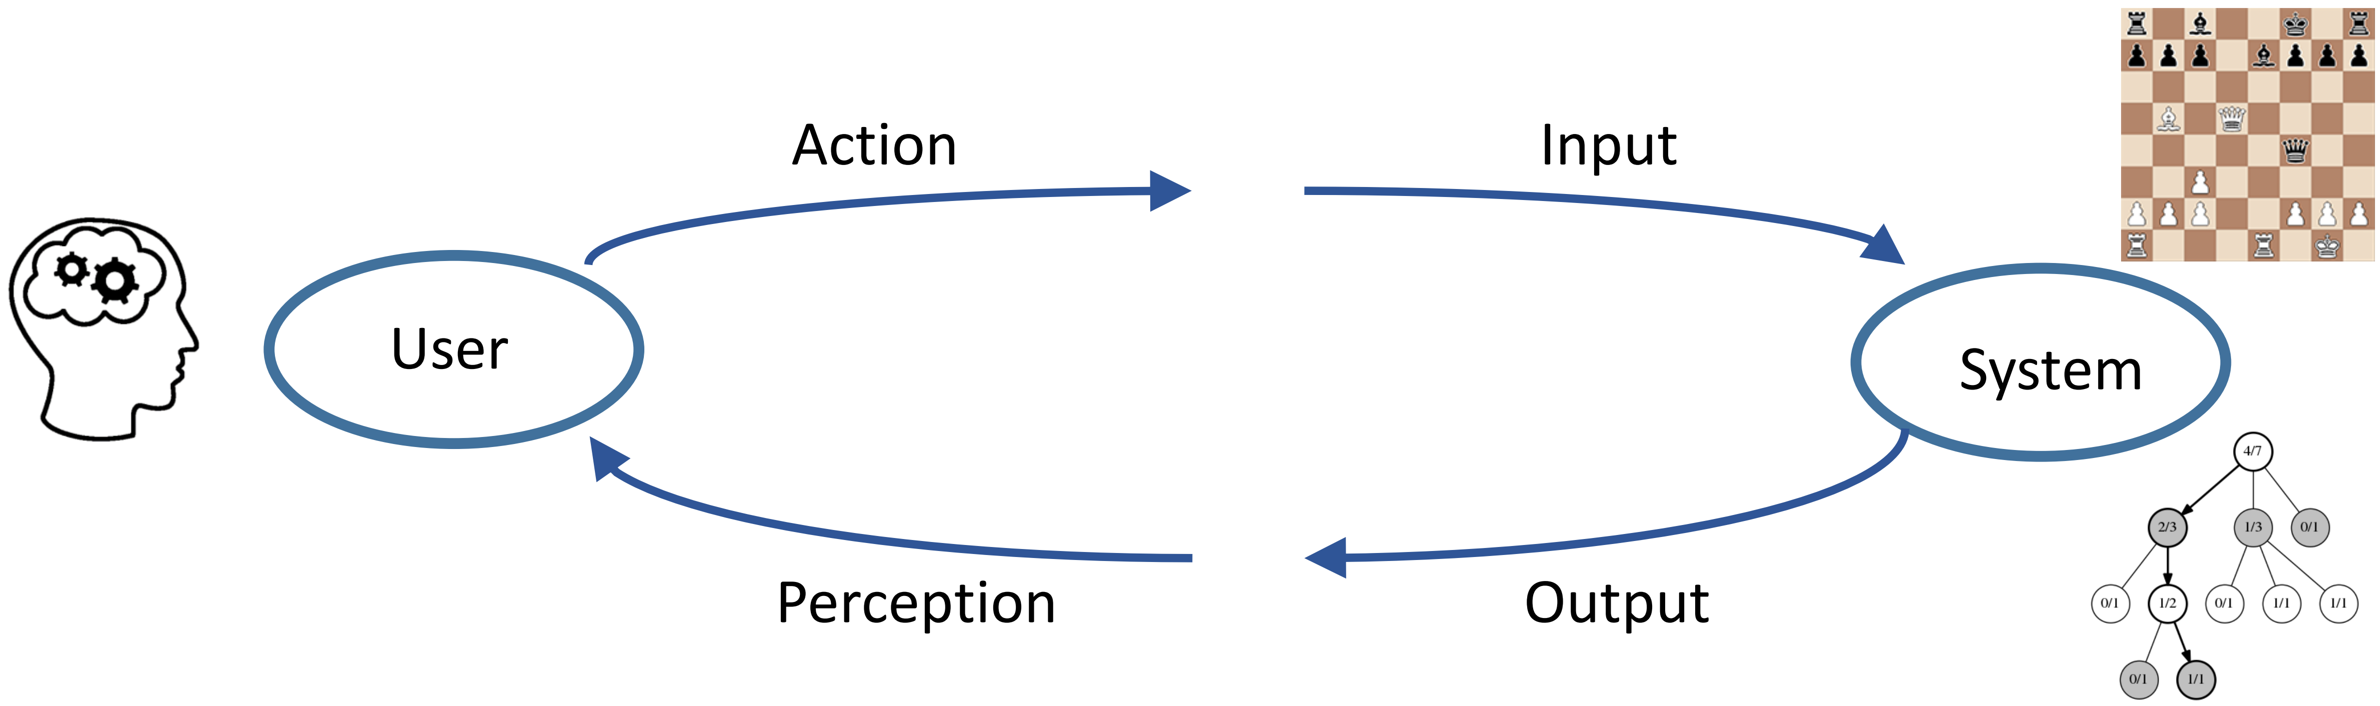
Action
Input
User
System
Perception
Output
4/7
2/3
1/3
0/1
0/1
1/2
0/1
1/1
1/1
0/1
1/1

REVIEW

Computational rationality: A converging paradigm for intelligence in brains, minds, and machines

Samuel J. Gershman,[1*] Eric J. Horvitz,[2*] Joshua B. Tenenbaum[3*]

After growing up together, and mostly growing apart in the second half of the 20th century, the fields of artificial intelligence (AI), cognitive science, and neuroscience are reconverging on a shared view of the computational foundations of intelligence that promotes valuable cross-disciplinary exchanges on questions, methods, and results. We chart advances over the past several decades that address challenges of perception and action under uncertainty through the lens of computation. Advances include the development of representations and inferential procedures for large-scale probabilistic inference and machinery for enabling reflection and decisions about tradeoffs in effort, precision, and timeliness of computations. These tools are deployed toward the goal of computational rationality: identifying decisions with highest expected utility, while taking into consideration the costs of computation in complex real-world problems in which most relevant calculations can only be approximated. We highlight key concepts with examples that show the potential for interchange between computer science, cognitive science, and neuroscience.

Science 349.6245 (2015)

Models of computational rationality are built on a base of inferential processes for perceiving, predicting, learning, and reasoning under uncertainty (*1*–*3*). Such inferential processes operate on representations that encode probabilistic dependencies among variables capturing the likelihoods of relevant states in the world. In light of incoming streams of perceptual data, Bayesian updating procedures or approximations are used to propagate information and to compute and revise probability distributions over states of variables. Beyond base processes for evaluating probabilities, models of computational rationality require mechanisms for reasoning about the feasibility and implications of actions. Deliberation about the best action to take hinges on an ability to make predictions about how different actions will influence likelihoods of outcomes and a consideration of the value or utilities of the outcomes (*4*). Learning procedures make changes to parameters of probabilistic models so as to better explain perceptual data and provide more accurate inferences about likelihoods to guide actions in the world.

Last, systems with bounded computational power must consider important tradeoffs in the precision and timeliness of action in the world.

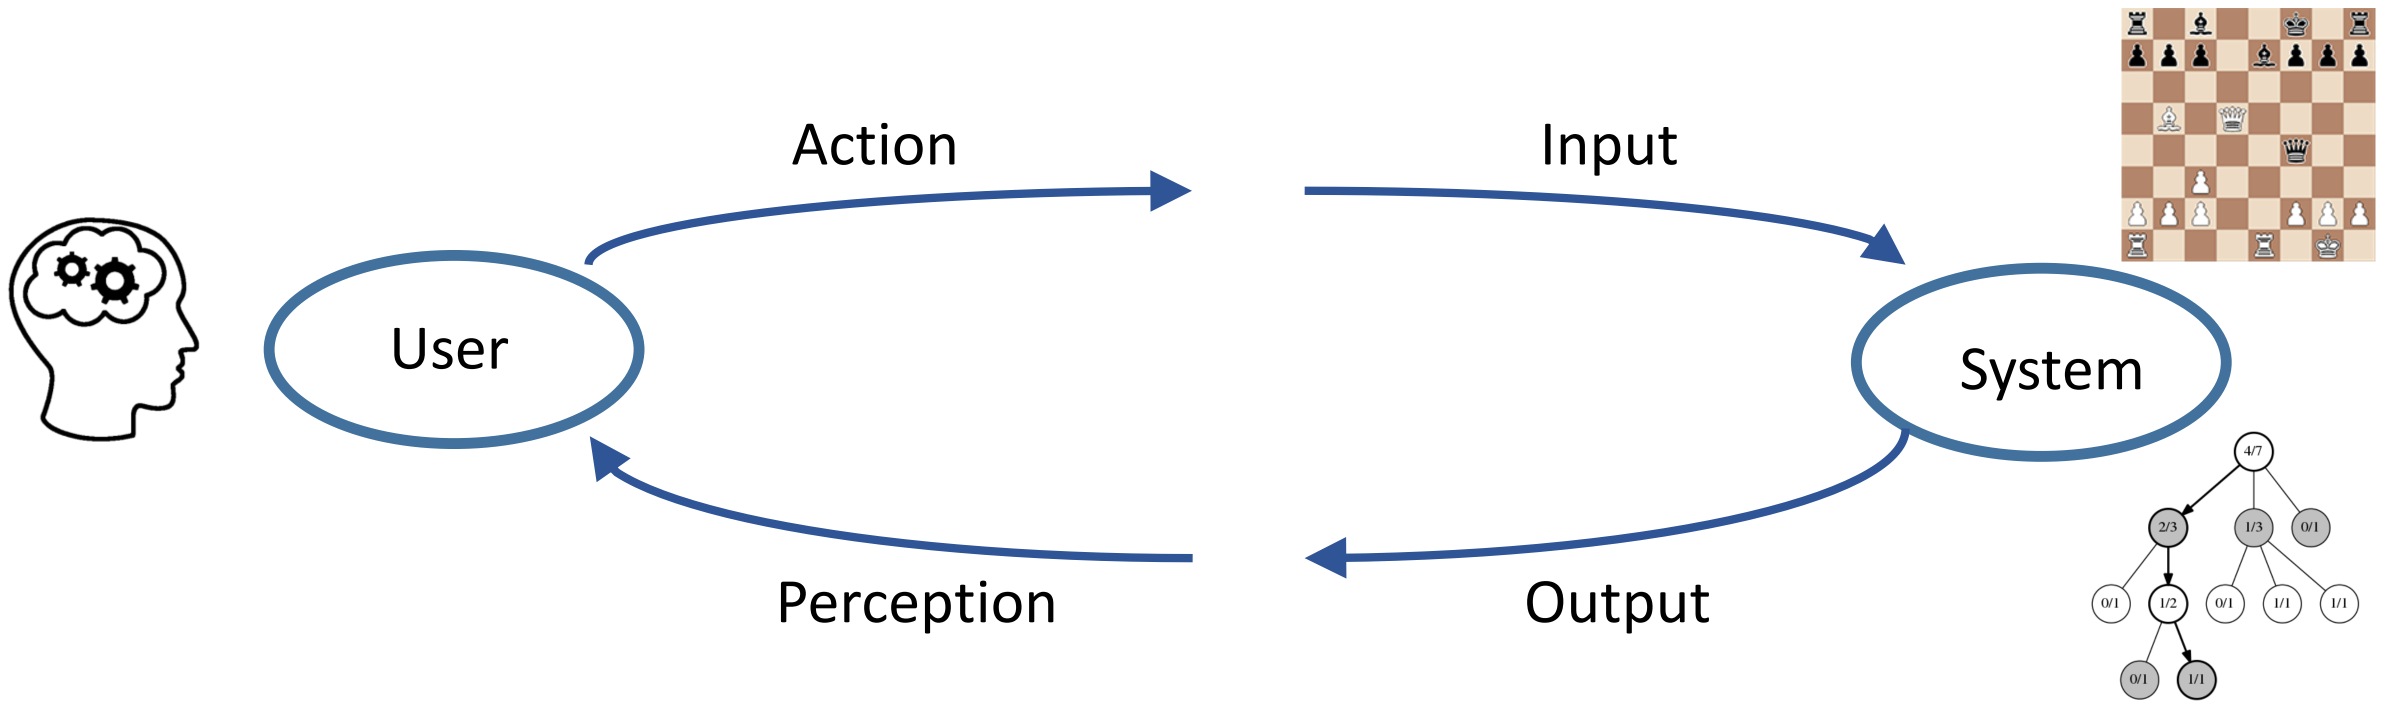
Action
Input
User
System
Perception
Output

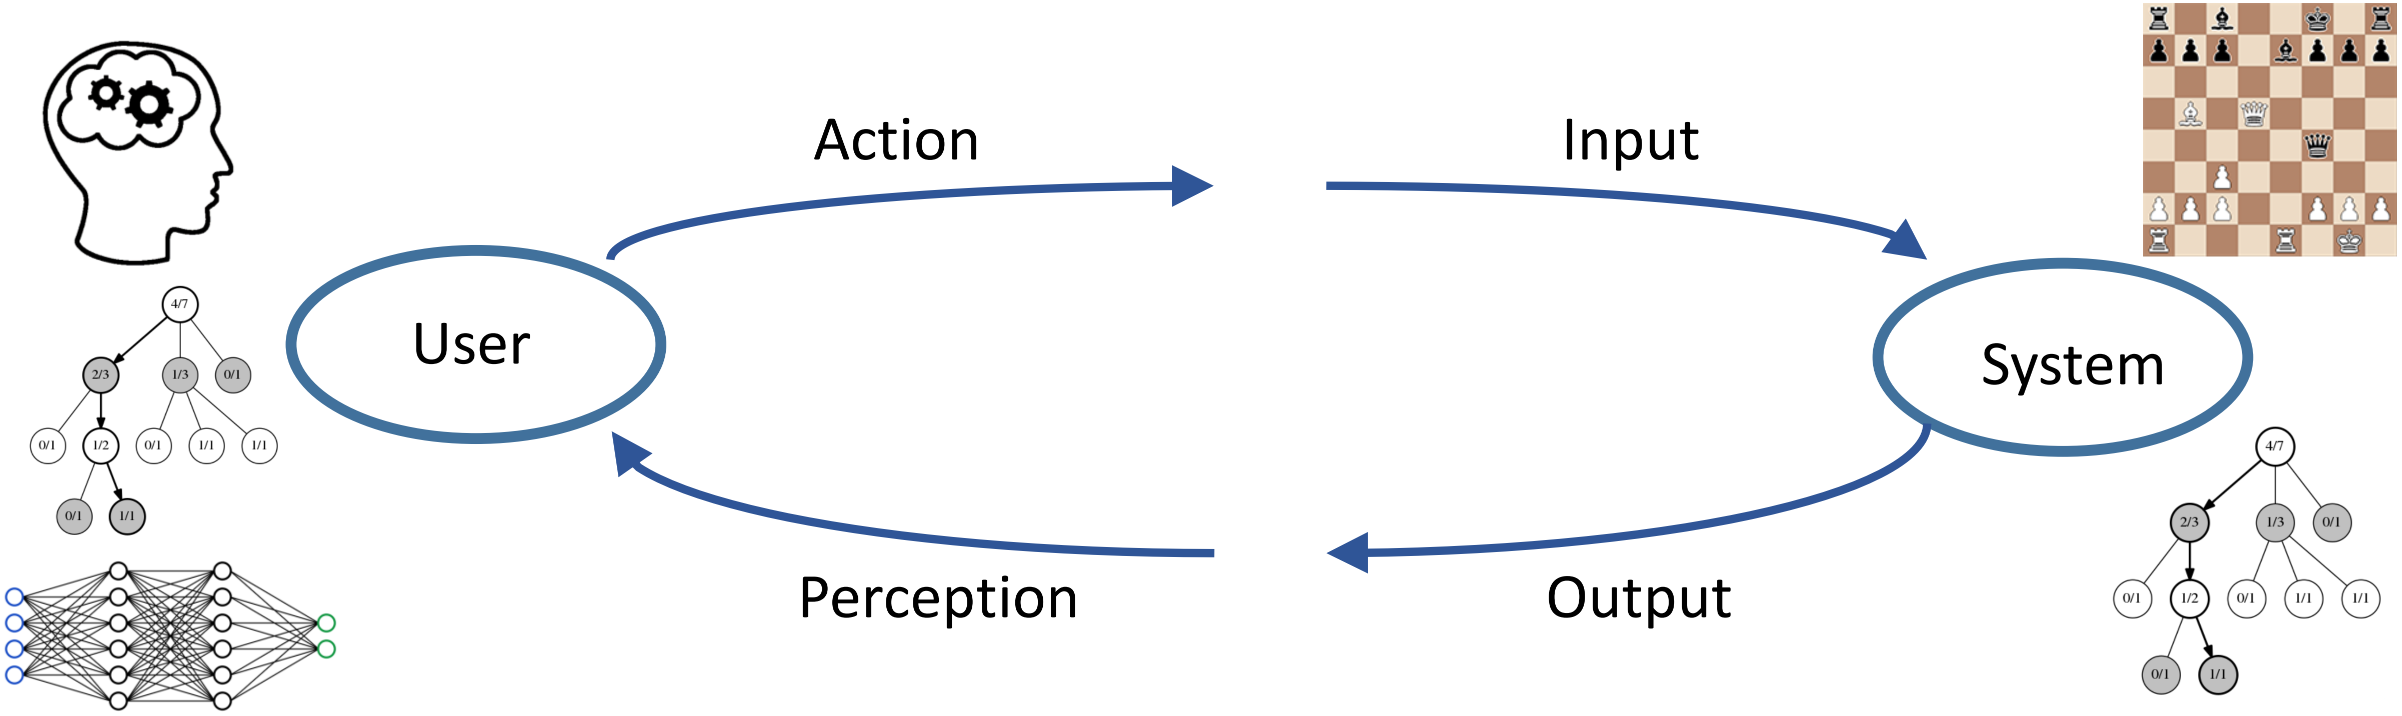
Action
Input
User
System
Perception
Output

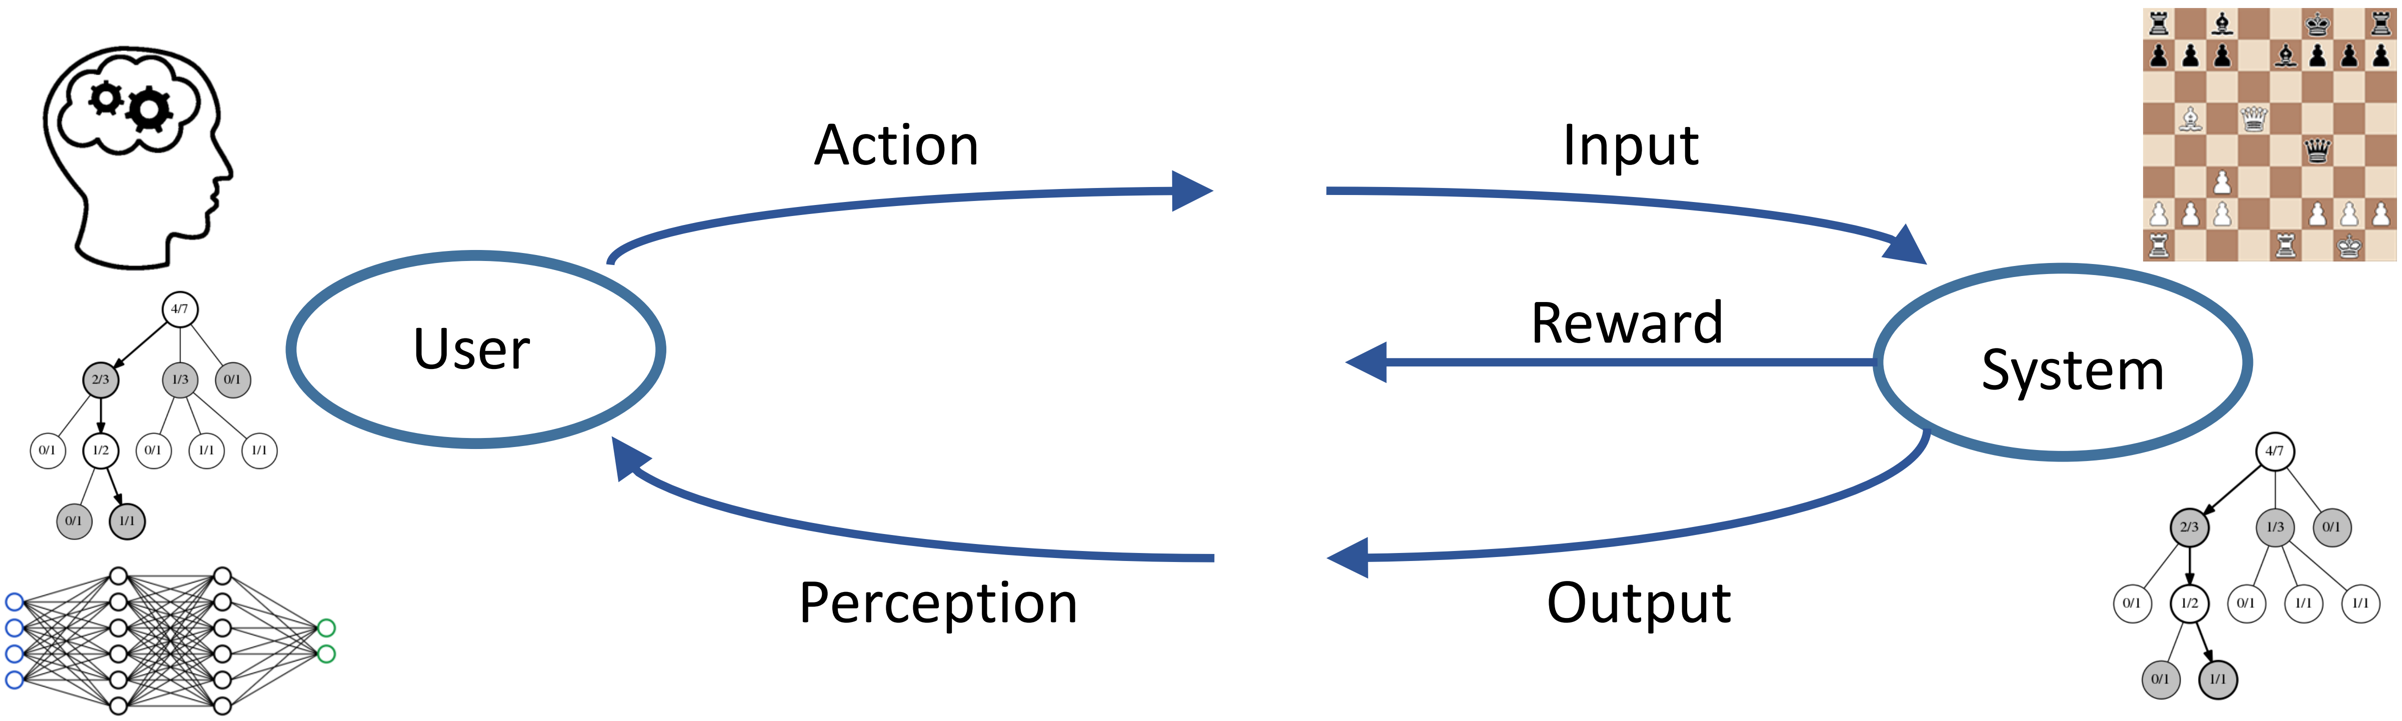
Action
Input
User
Reward
System
Perception
Output

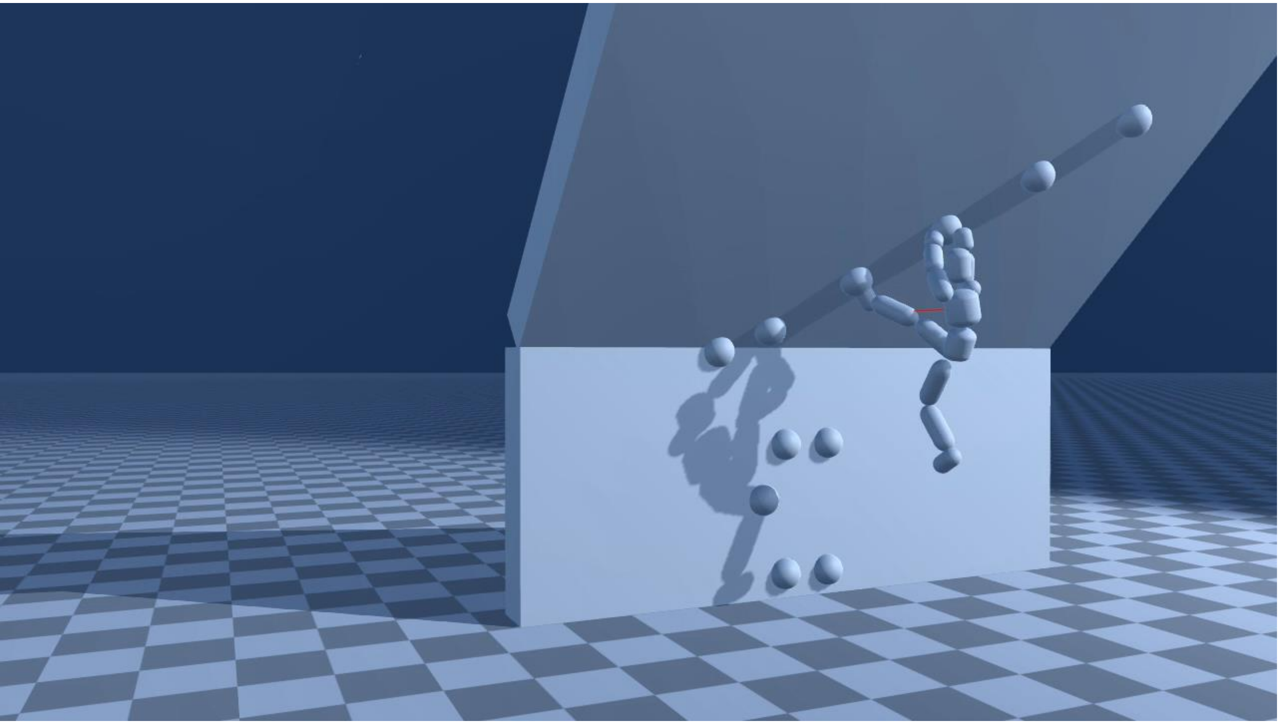

Computer Animation: Simulated motor control

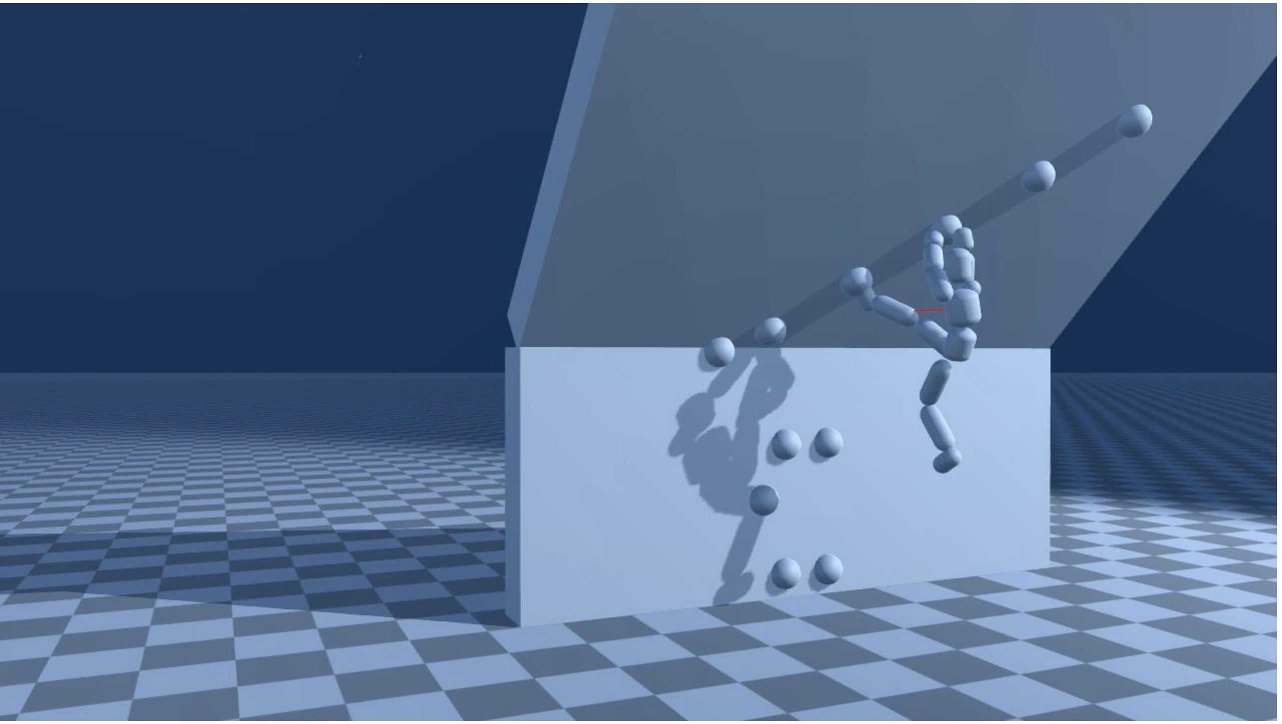

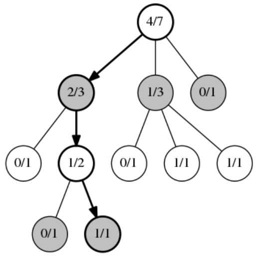

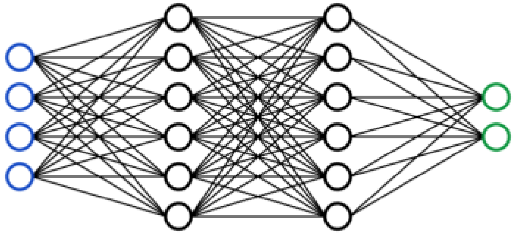

Computer Animation:
Simulated motor control

Computer vision:
Simulated perception

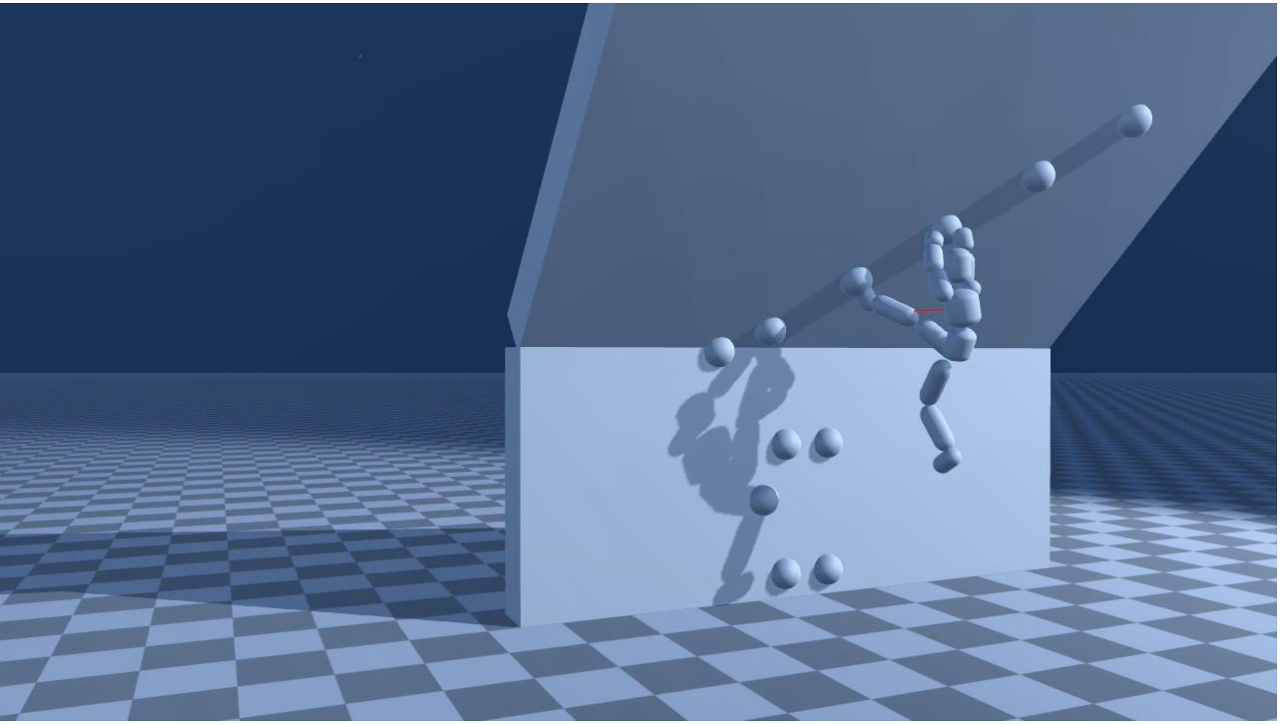

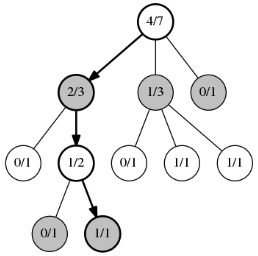

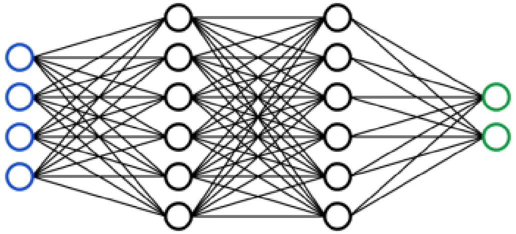

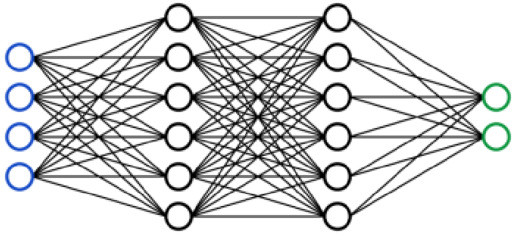

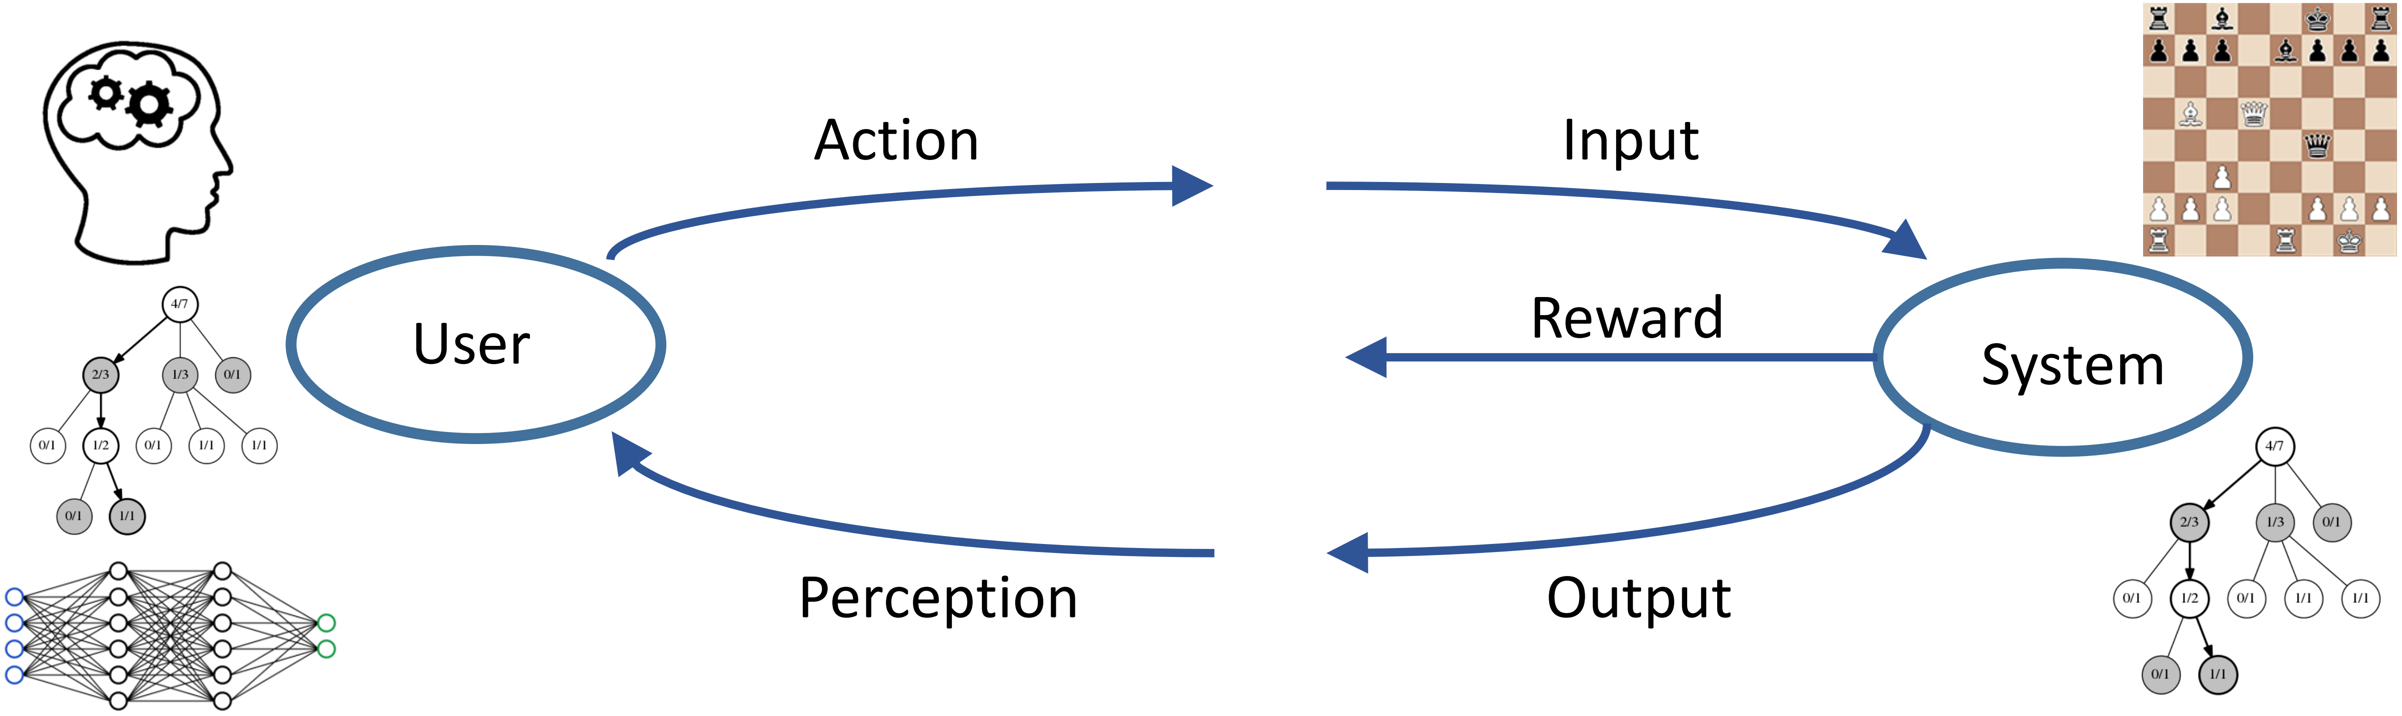
Action
Input
User
Reward
System
Perception
Output

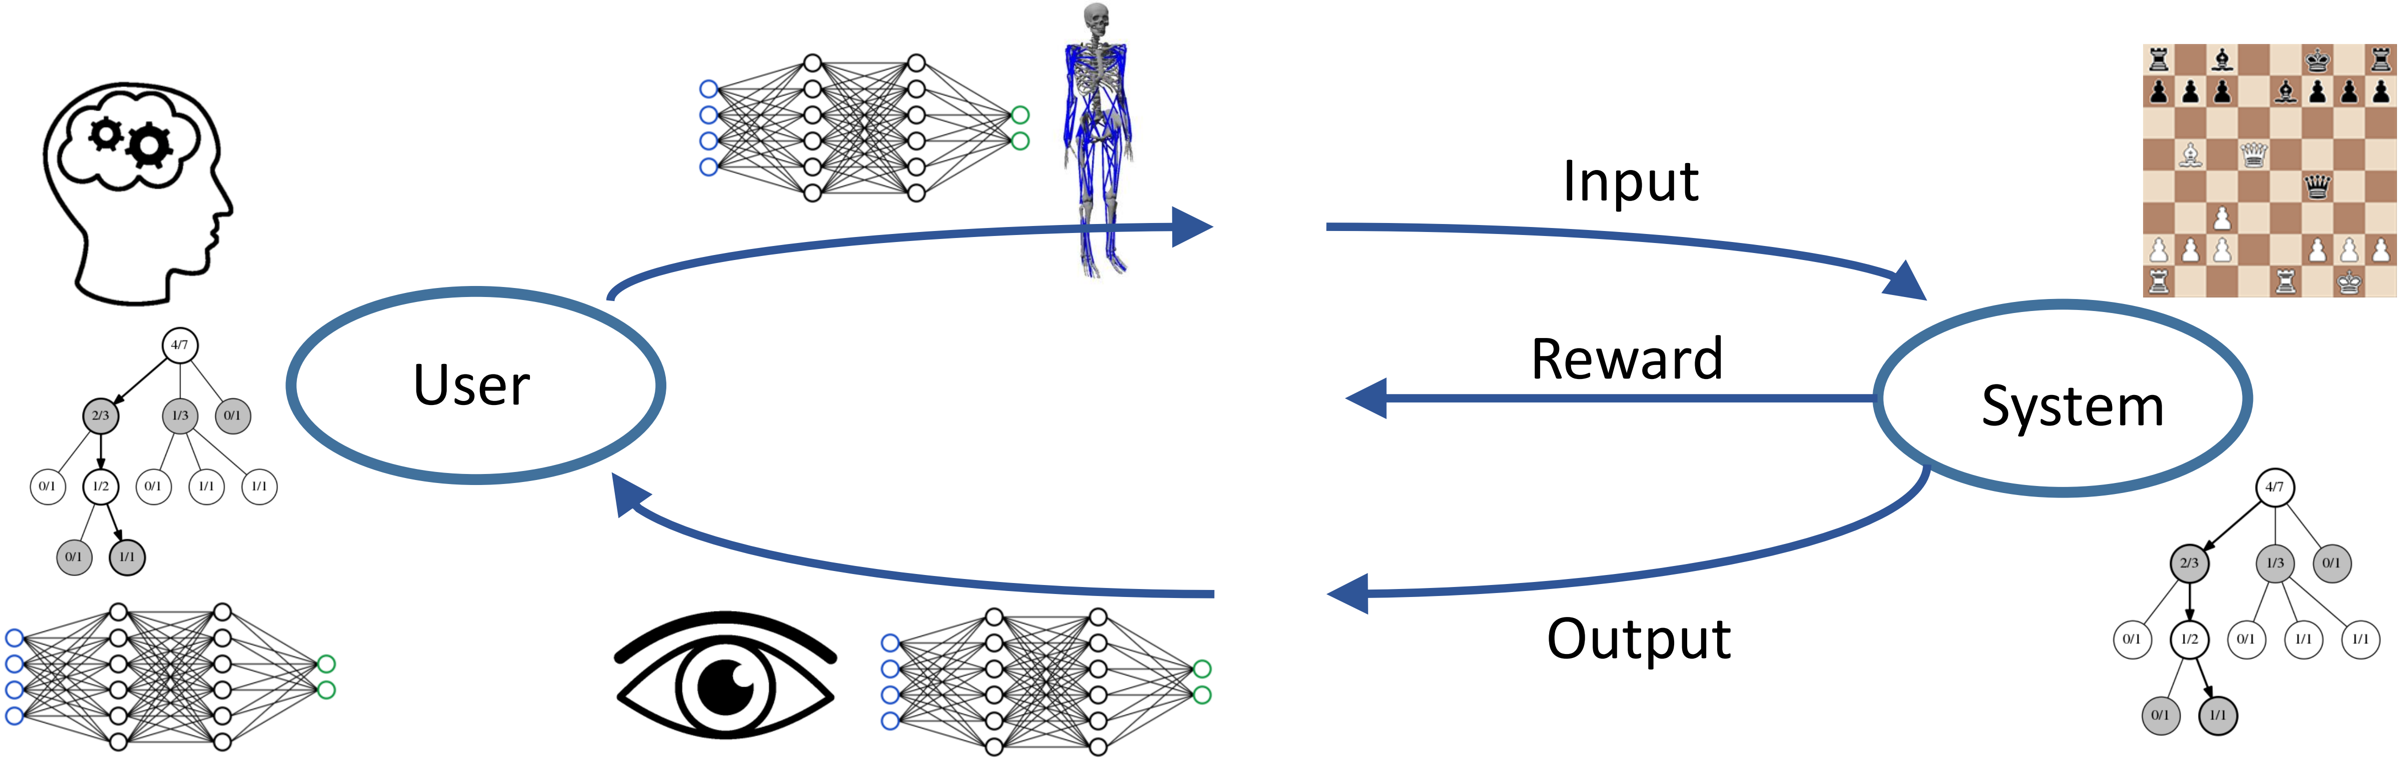
Input
Reward
Output
User
System

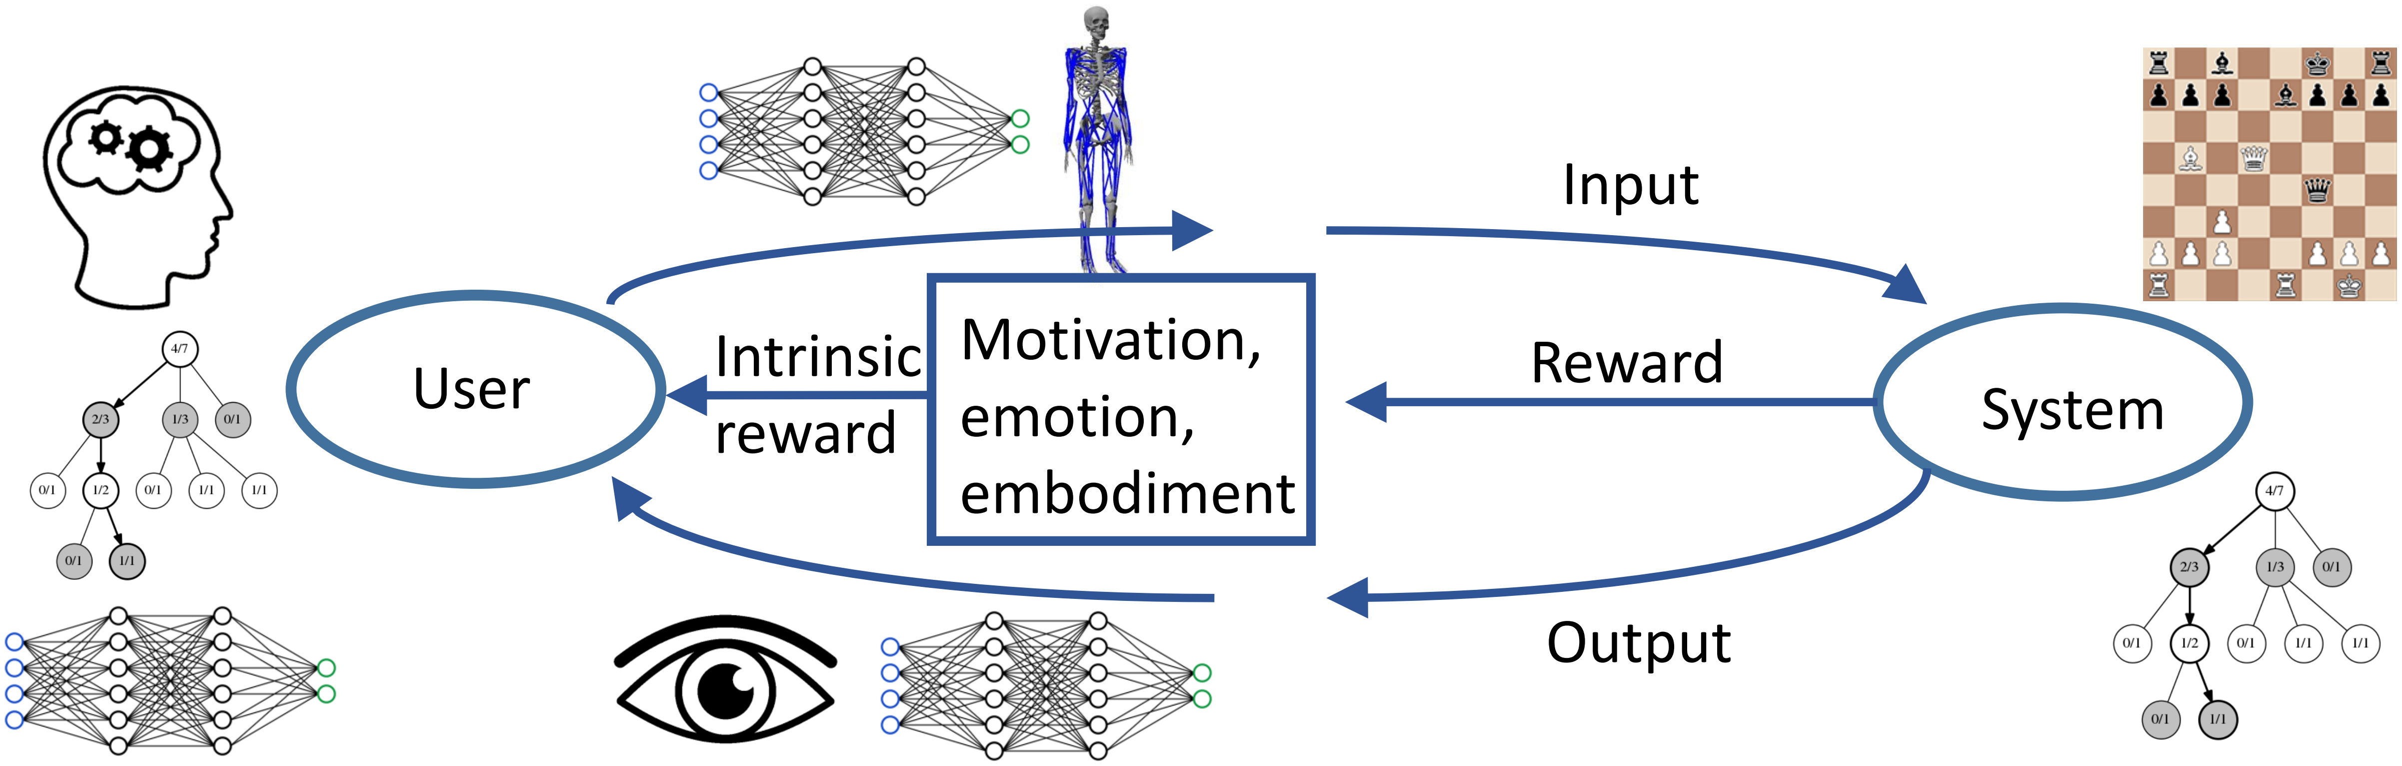
Input
Intrinsic reward
Motivation, emotion, embodiment
Reward
User
System
Output
4/7
2/3
1/3
0/1
1/2
1/1

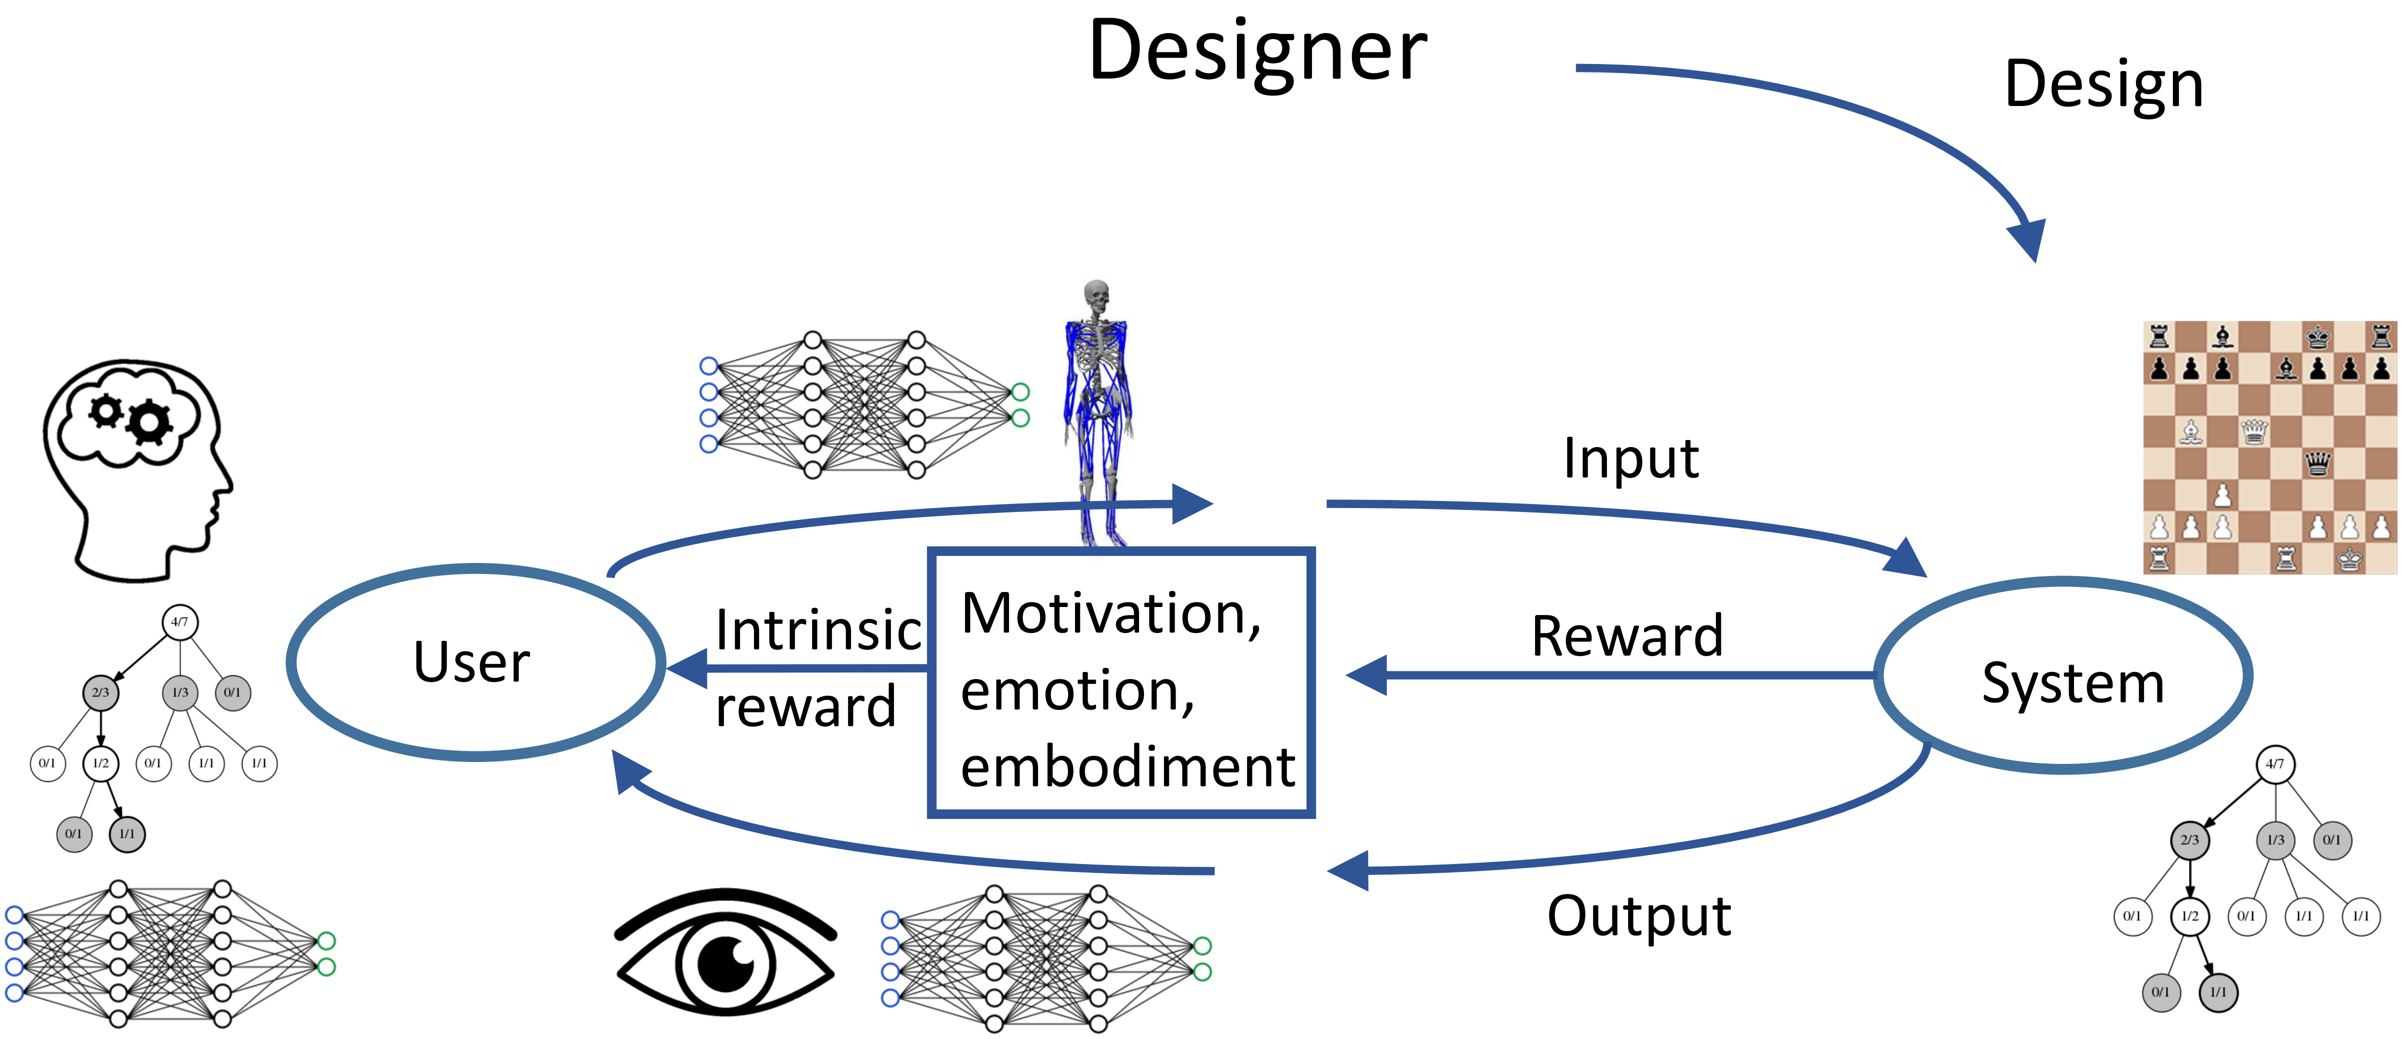
Designer
Design
Input
Motivation,
emotion,
embodiment
Intrinsic
reward
Reward
User
System
Output

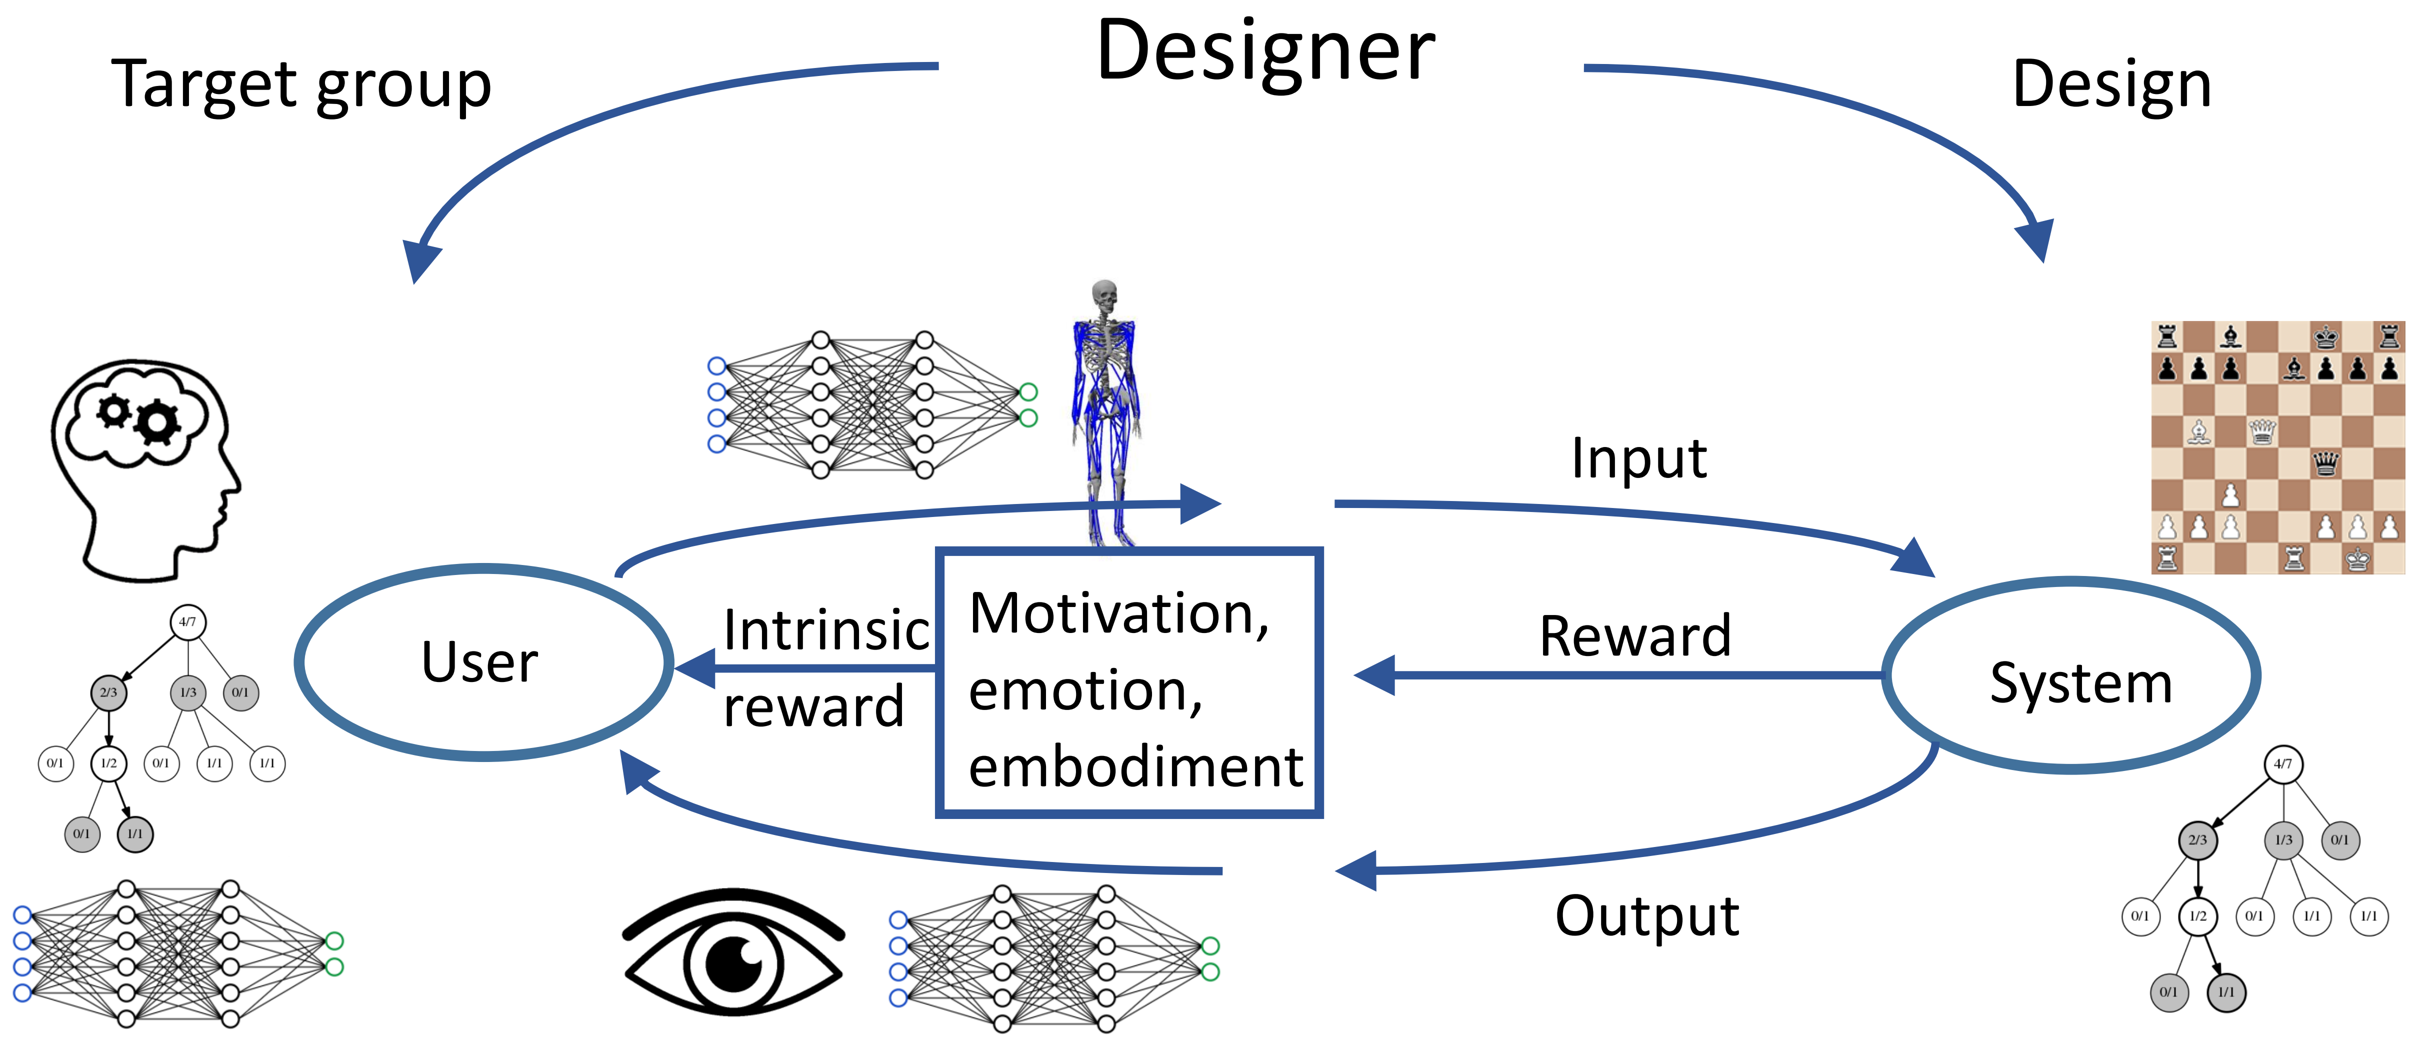

Designer
Target group
Design
Input
User
Intrinsic
reward
Motivation,
emotion,
embodiment
Reward
System
Output
4/7
2/3
1/3
0/1
0/1
1/2
0/1
1/1
1/1
0/1
1/1

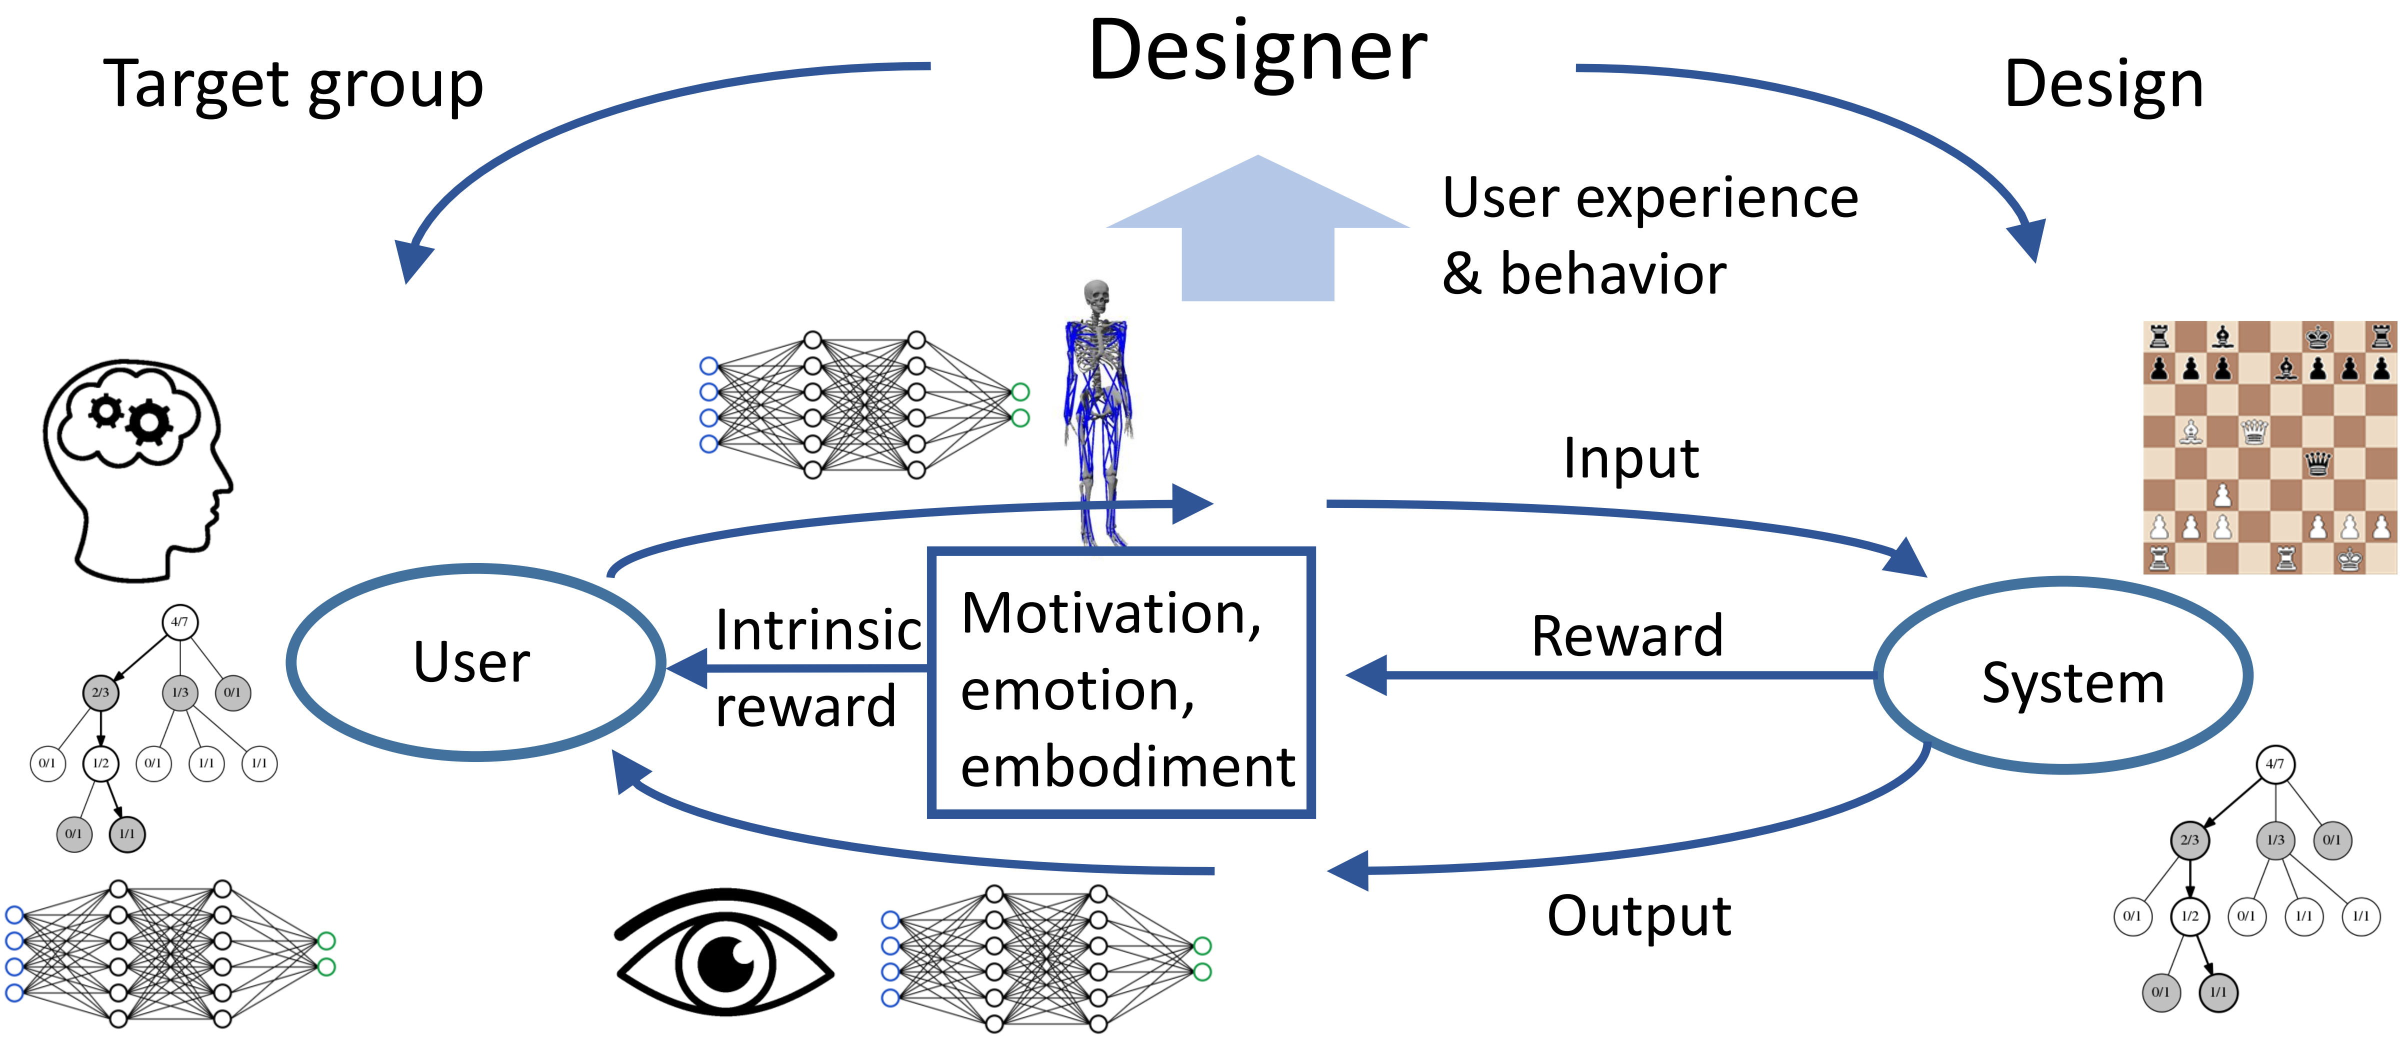
Designer
Target group
Design
User experience
& behavior
Input
Intrinsic
reward
Motivation,
emotion,
embodiment
Reward
User
System
Output

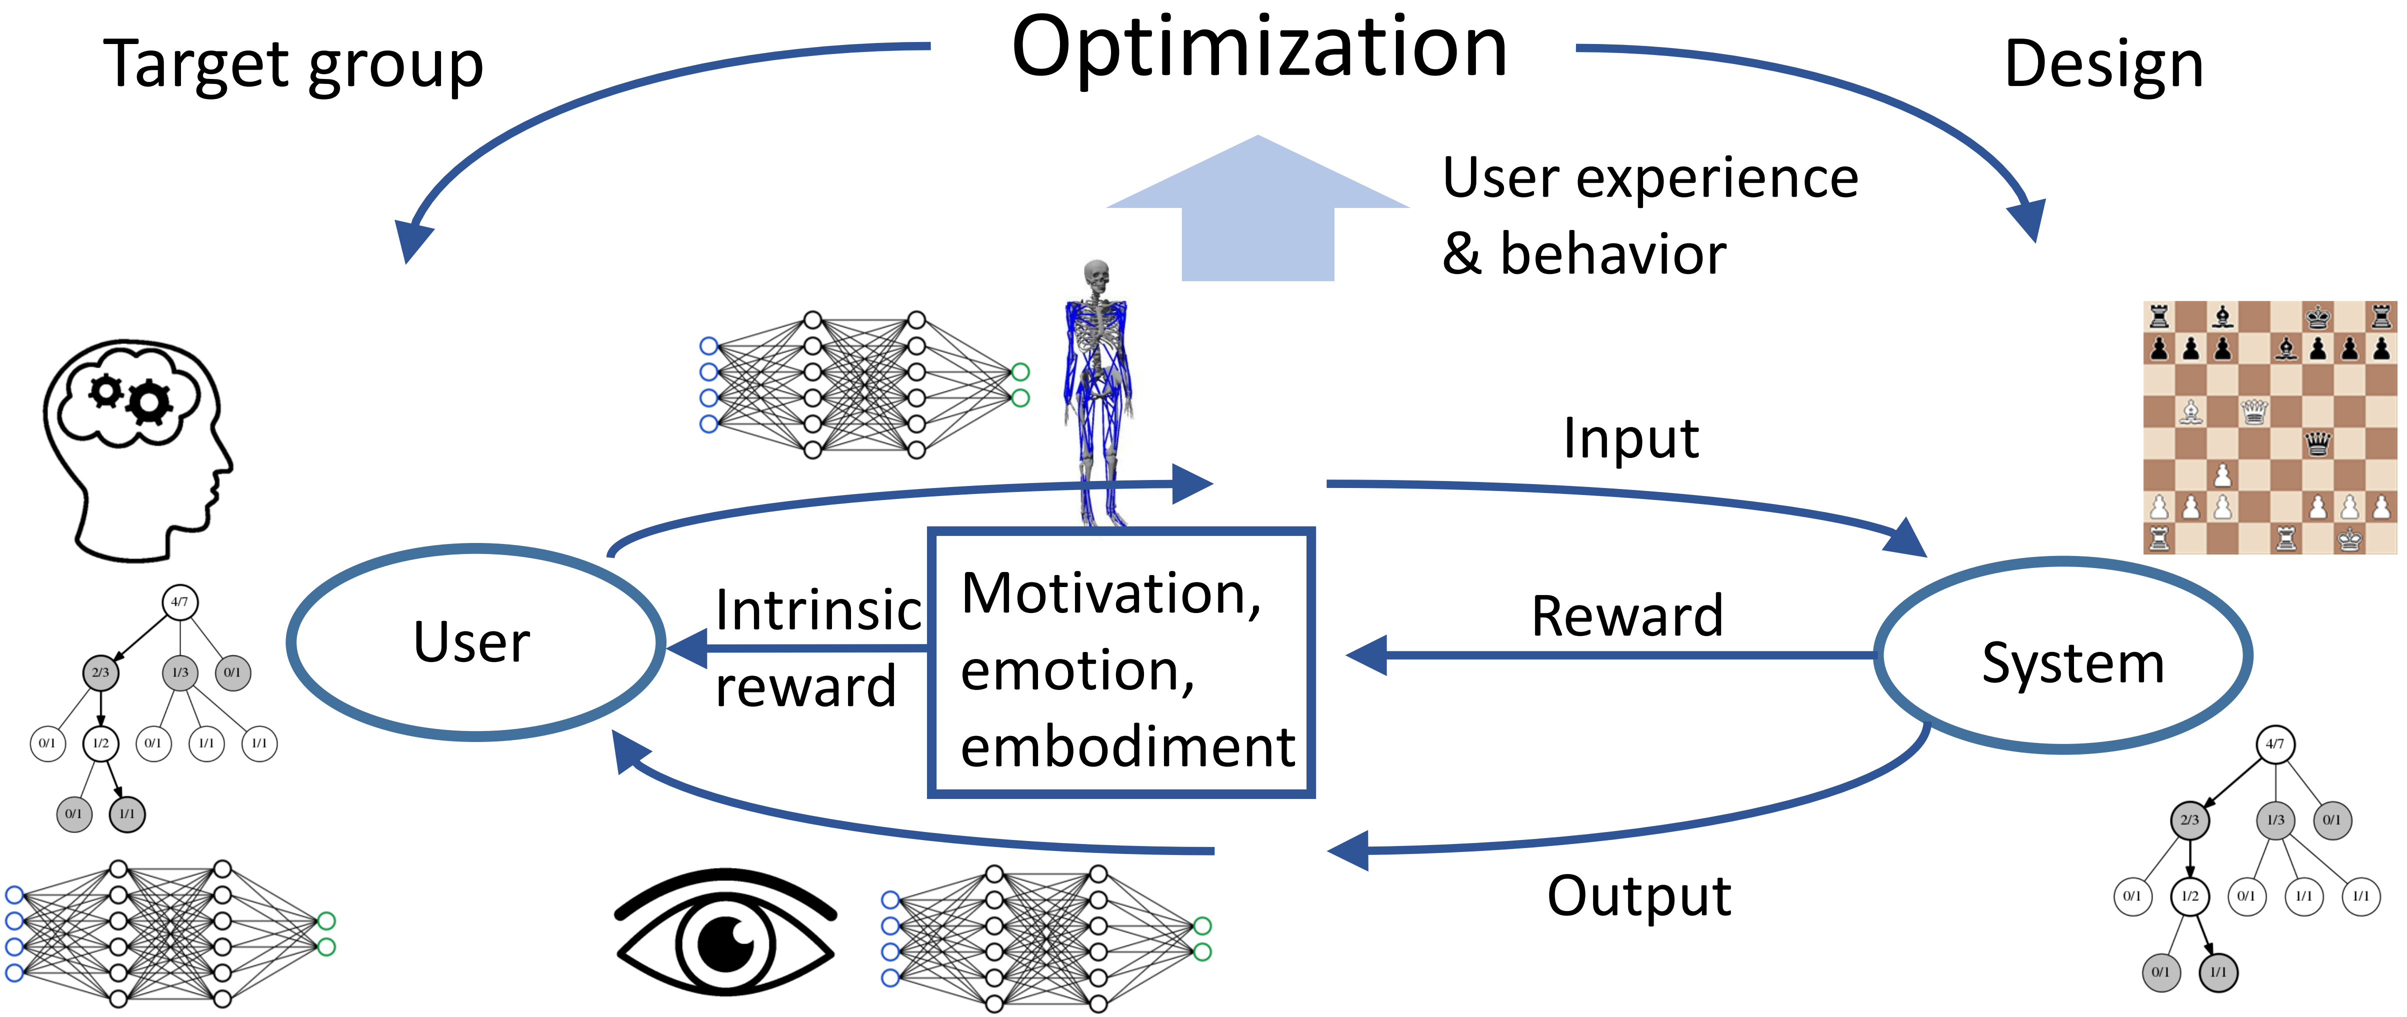
Optimization
Target group
Design
User experience
& behavior
Input
Intrinsic
reward
Motivation,
emotion,
embodiment
Reward
User
System
Output
4/7
2/3
1/3
0/1
1/2
1/1

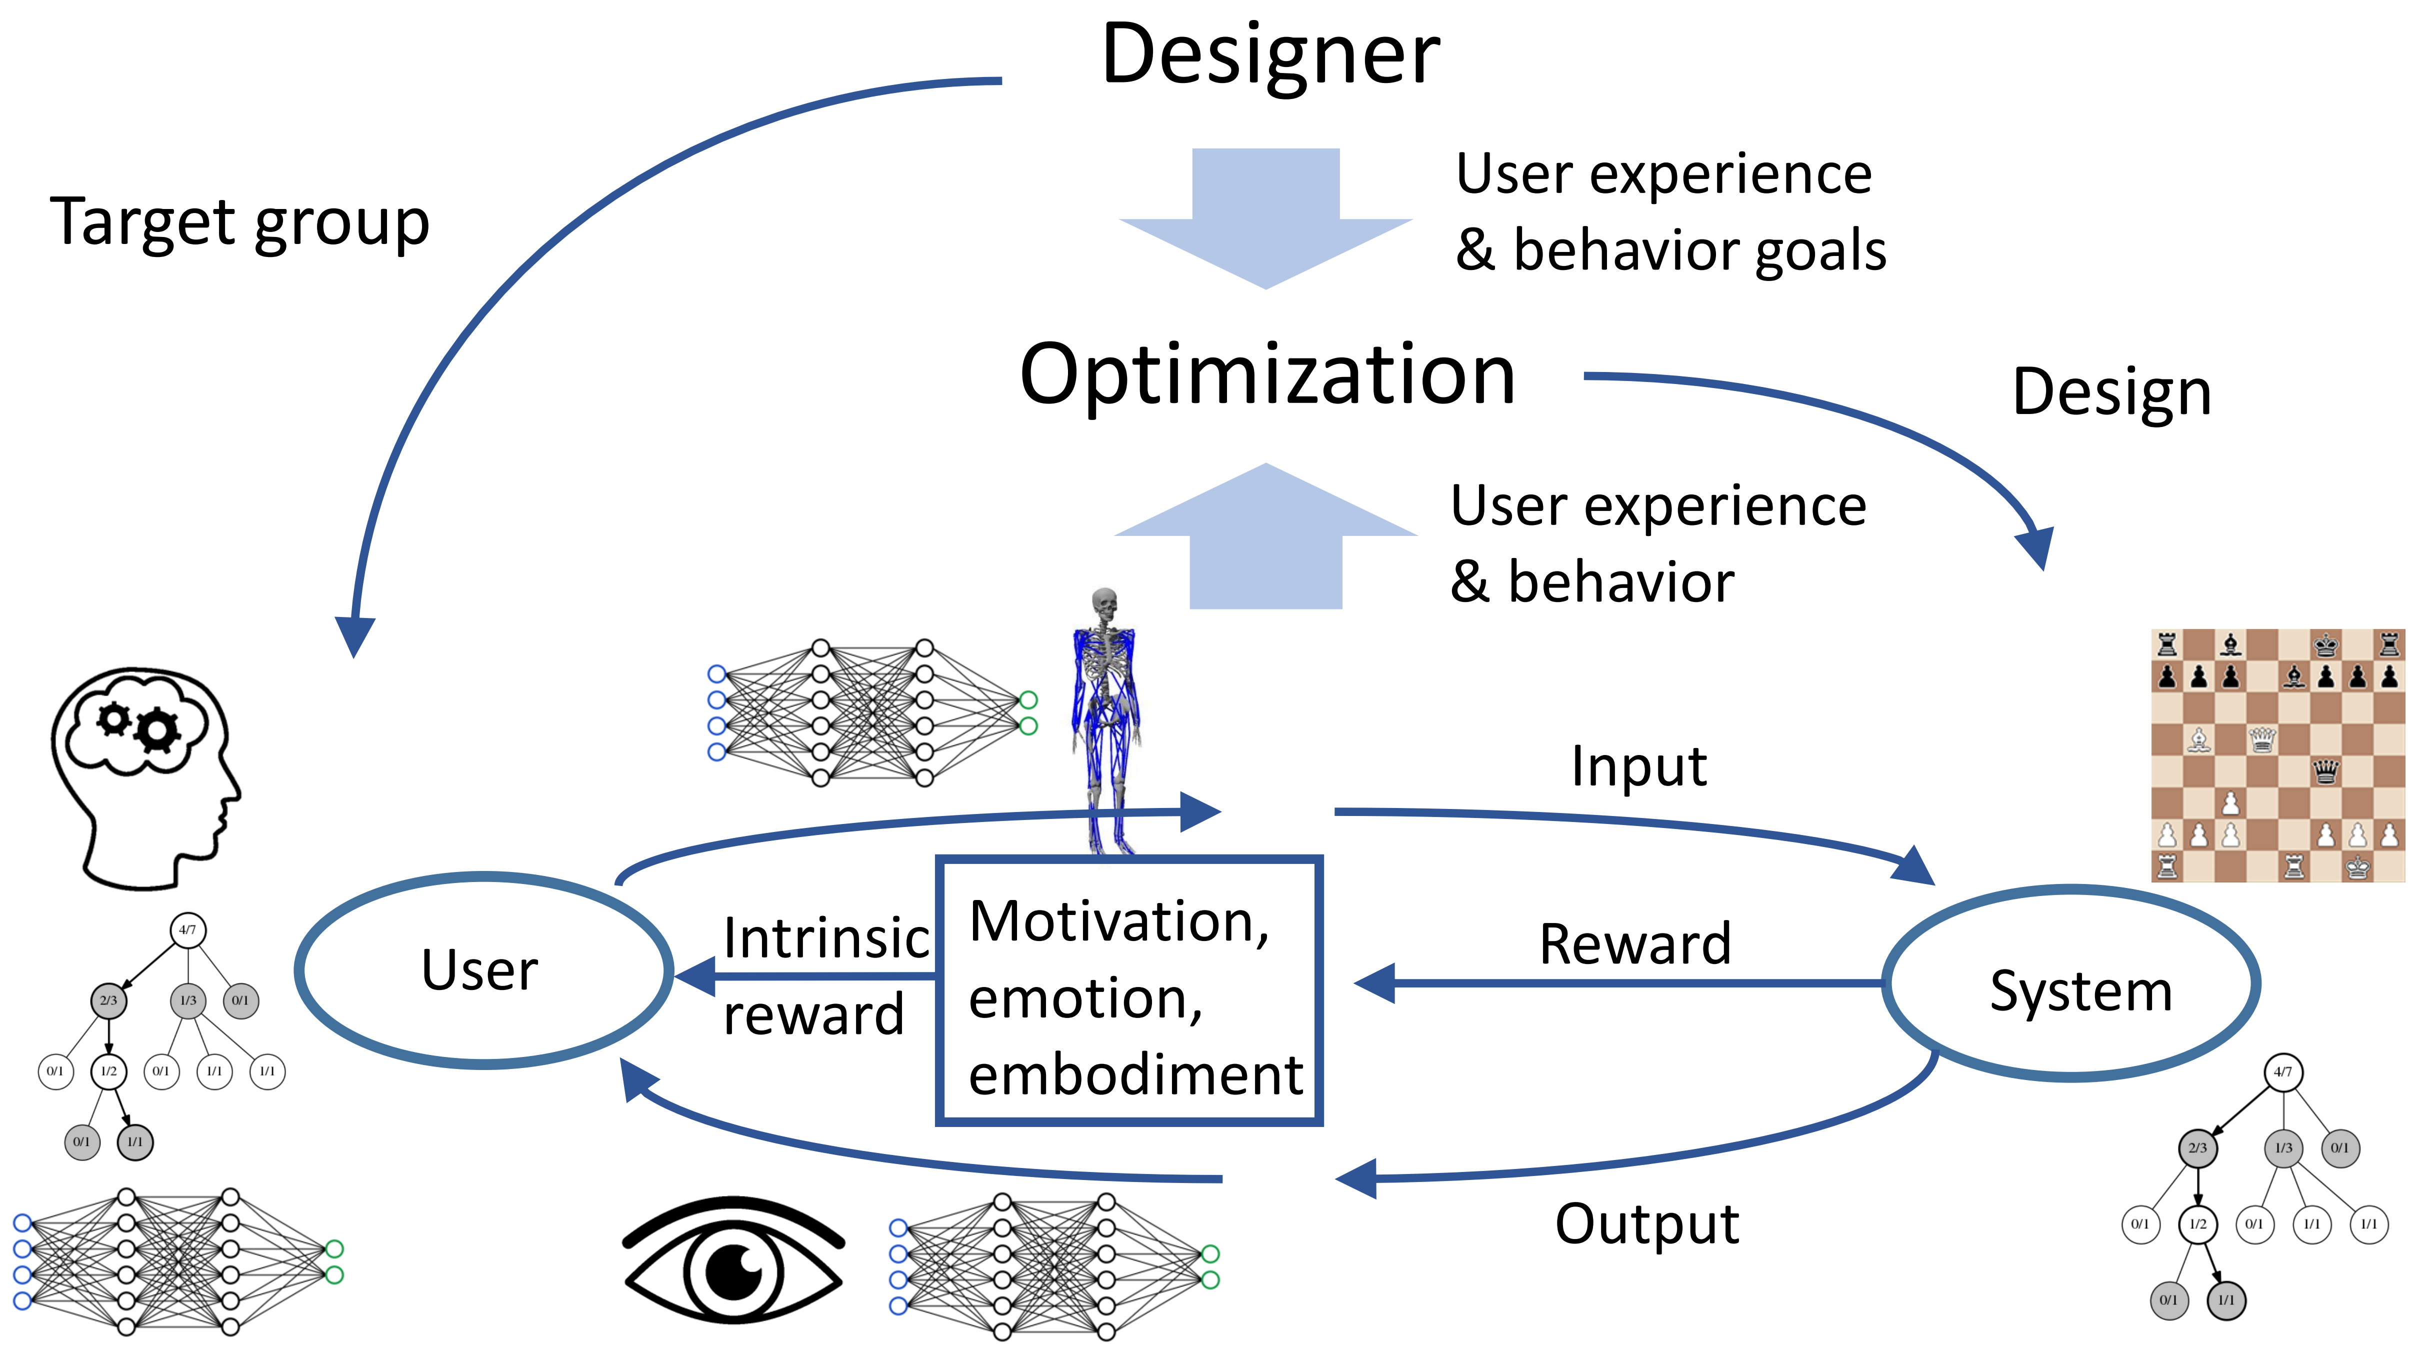
Designer
Target group
User experience
& behavior goals
Optimization
Design
User experience
& behavior
Input
Intrinsic
reward
Motivation,
emotion,
embodiment
Reward
User
System
Output

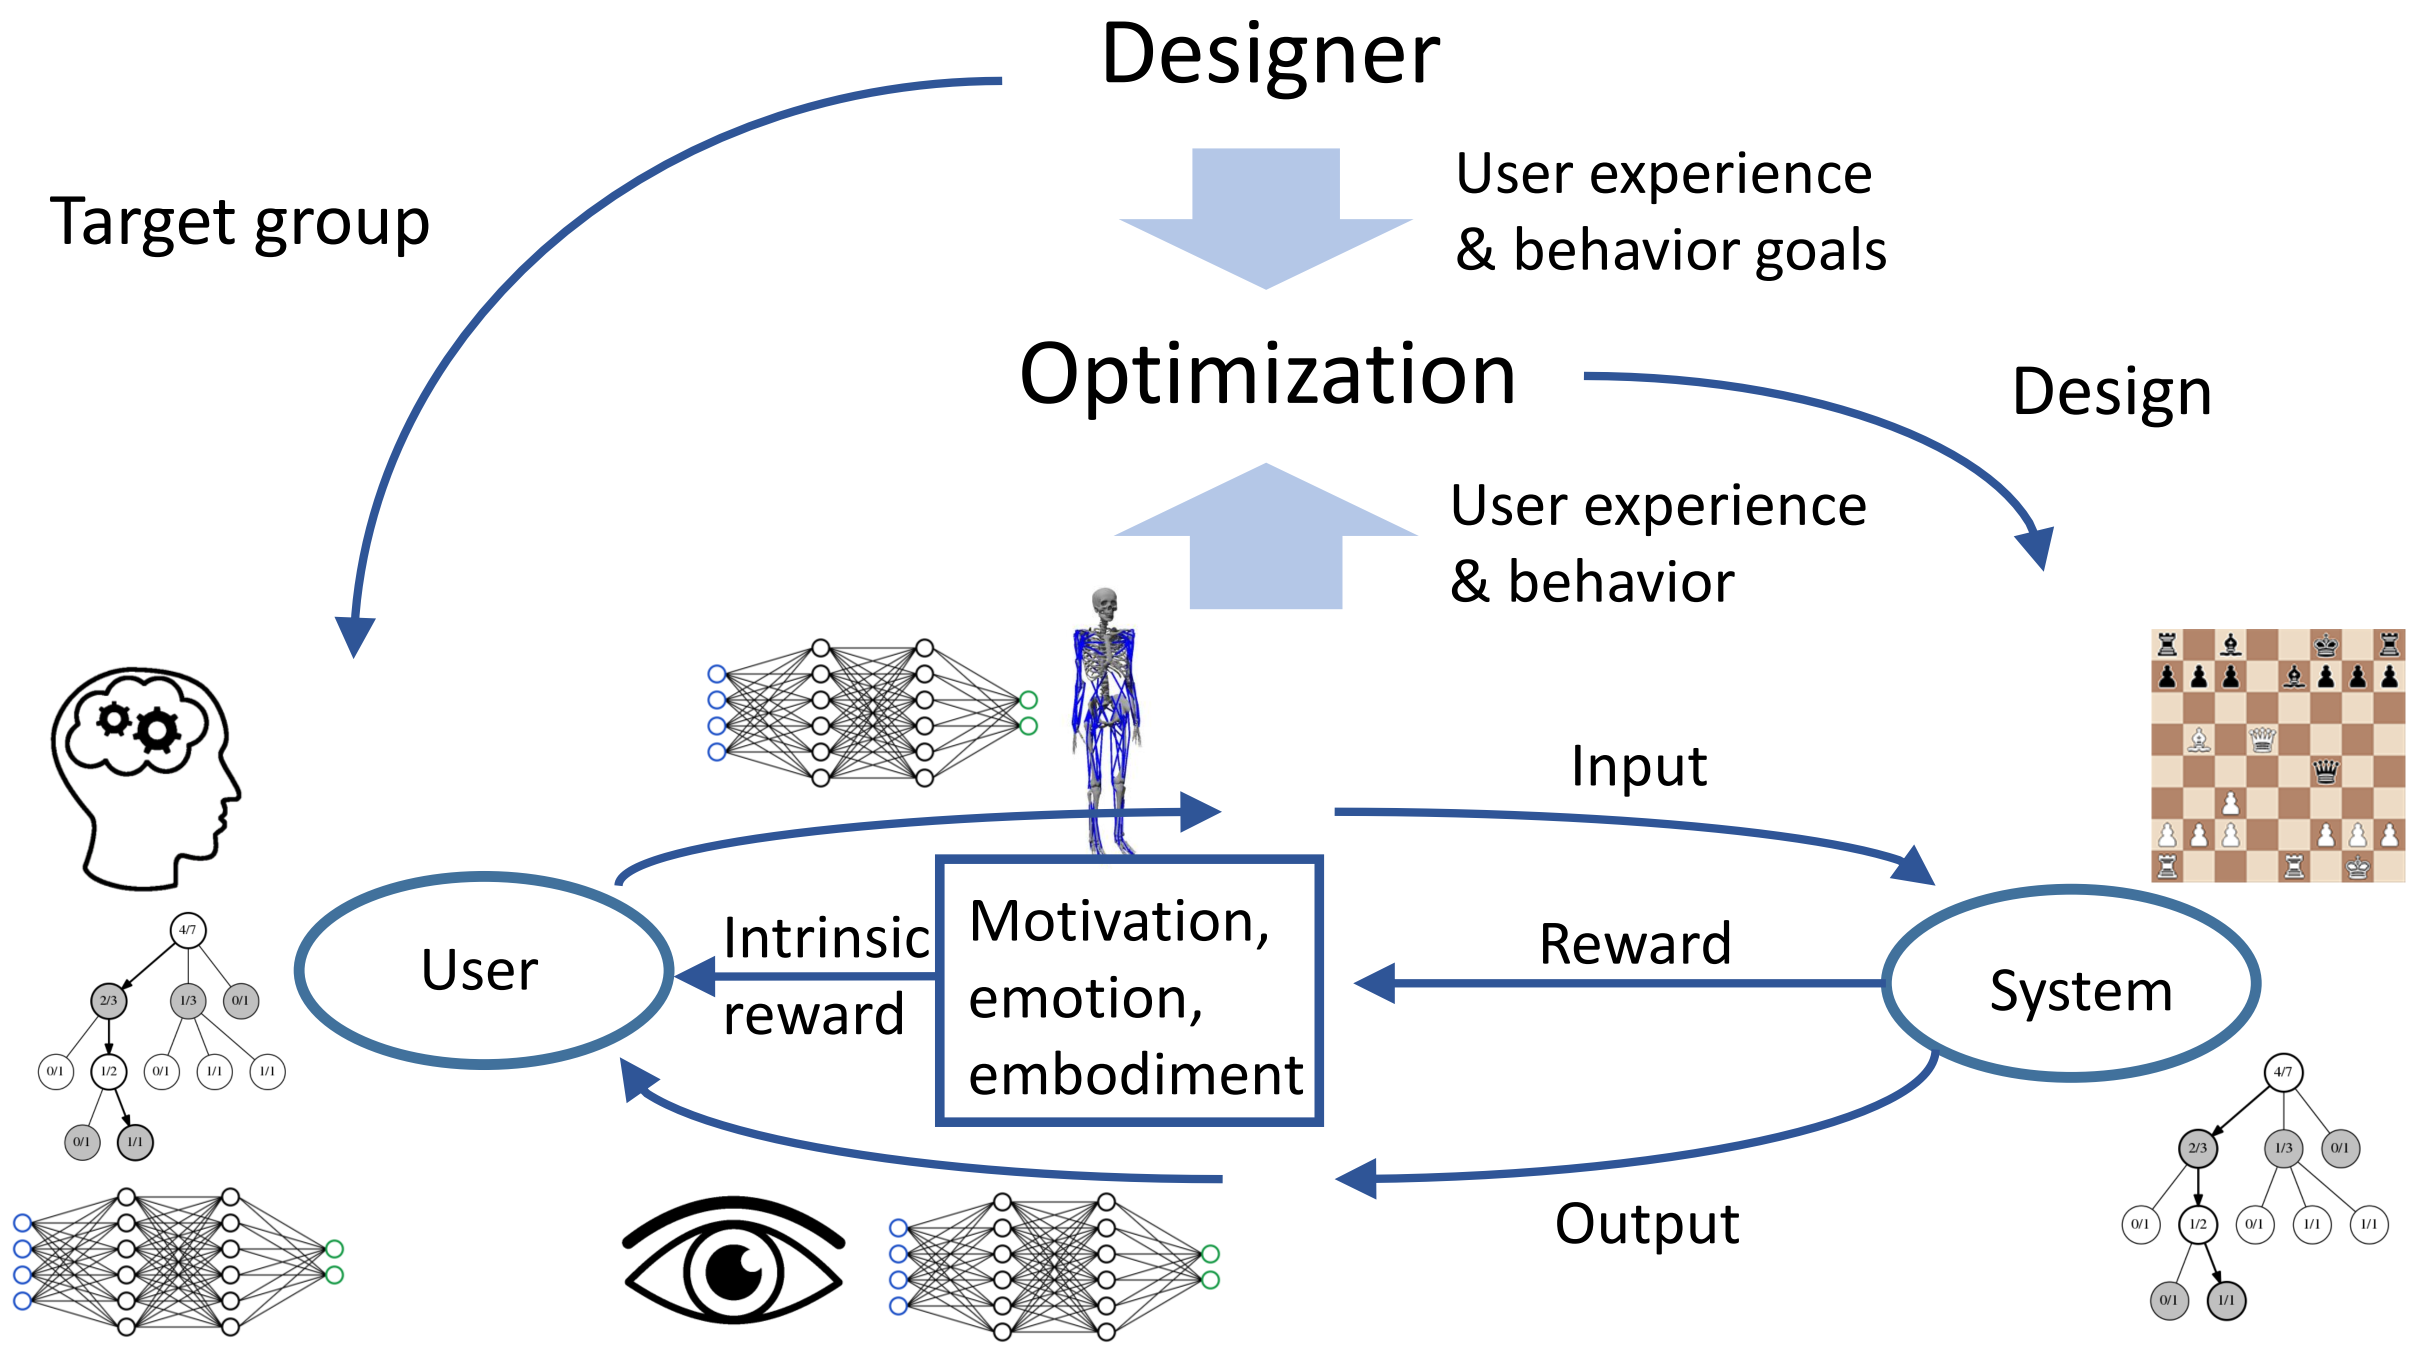
Designer
Target group
User experience
& behavior goals
Optimization
Design
User experience
& behavior
Input
Intrinsic
reward
Motivation,
emotion,
embodiment
Reward
User
System
Output

Examples

Monte-Carlo Tree Search for Persona Based Player Modeling

Christoffer Holmgård[1], Antonios Liapis[2], Julian Togelius[1,3], Georgios N. Yannakakis[1,2]
1: Center for Computer Games Research, IT University of Copenhagen, Copenhagen, Denmark
2: Institute of Digital Games, University of Malta, Msida, Malta
3: Department of Computer Science and Engineering, New York University, New York, USA

Abstract

Is it possible to conduct player modeling without any players? In this paper we use Monte-Carlo Tree Search-controlled procedural personas to simulate a range of decision making styles in the puzzle game *MiniDungeons 2*. The purpose is to provide a method for synthetic play testing of game levels with synthetic players based on designer intuition and experience. Five personas are constructed, representing five different decision making styles archetypal for the game. The personas vary solely in the weights of decision-making utilities that describe their valuation of a set affordances in MiniDungeons 2. By configuring these weights using designer expert knowledge, and passing the configurations directly to the MCTS algorithm, we make the personas exhibit a number of distinct decision making and play styles.

to inform their design decisions or to integrate in their content creation systems. When a designer defines such an agent for a particular game we call them *procedural personas*. As player types akin to those described by Bartle (1996) and the play personas described by Tychsen and Canossa (2008) and Canossa and Drachen (2009), they describe archetypal ways of interacting with the game. They are formal representations of the game designer's assumptions about her players.

Each persona may be used interactively or automatically in the level design process. Interactively, a level designer can inspect different interaction patterns (e.g. play-traces or completion statistics) in the level and iteratively adapt either the level or the persona behavior (Yannakakis, Liapis, and Alexopoulos 2014). Automatically, a procedural content

Human-Like Playtesting with Deep Learning

Tech at King Follow

Jan 11 · 14 min read

By Alex Nodet—Artificial Intelligence Engineer

King, like many other game companies, follows a free-to-play business model. This trend has increased within the gaming industry in recent years. Its efficiency is driven by frequent releases of new in-game content.

As of October 2018, Candy Crush Saga offers more than 3,700 levels to its players, and 15 new ones are released every week. Offering high-quality content is one of our core values at King. Therefore it is important to make sure that every level we release is correctly balanced. One traditional way is to ask playtesters for feedback. However, it comes with limitations we will discuss later in this article.

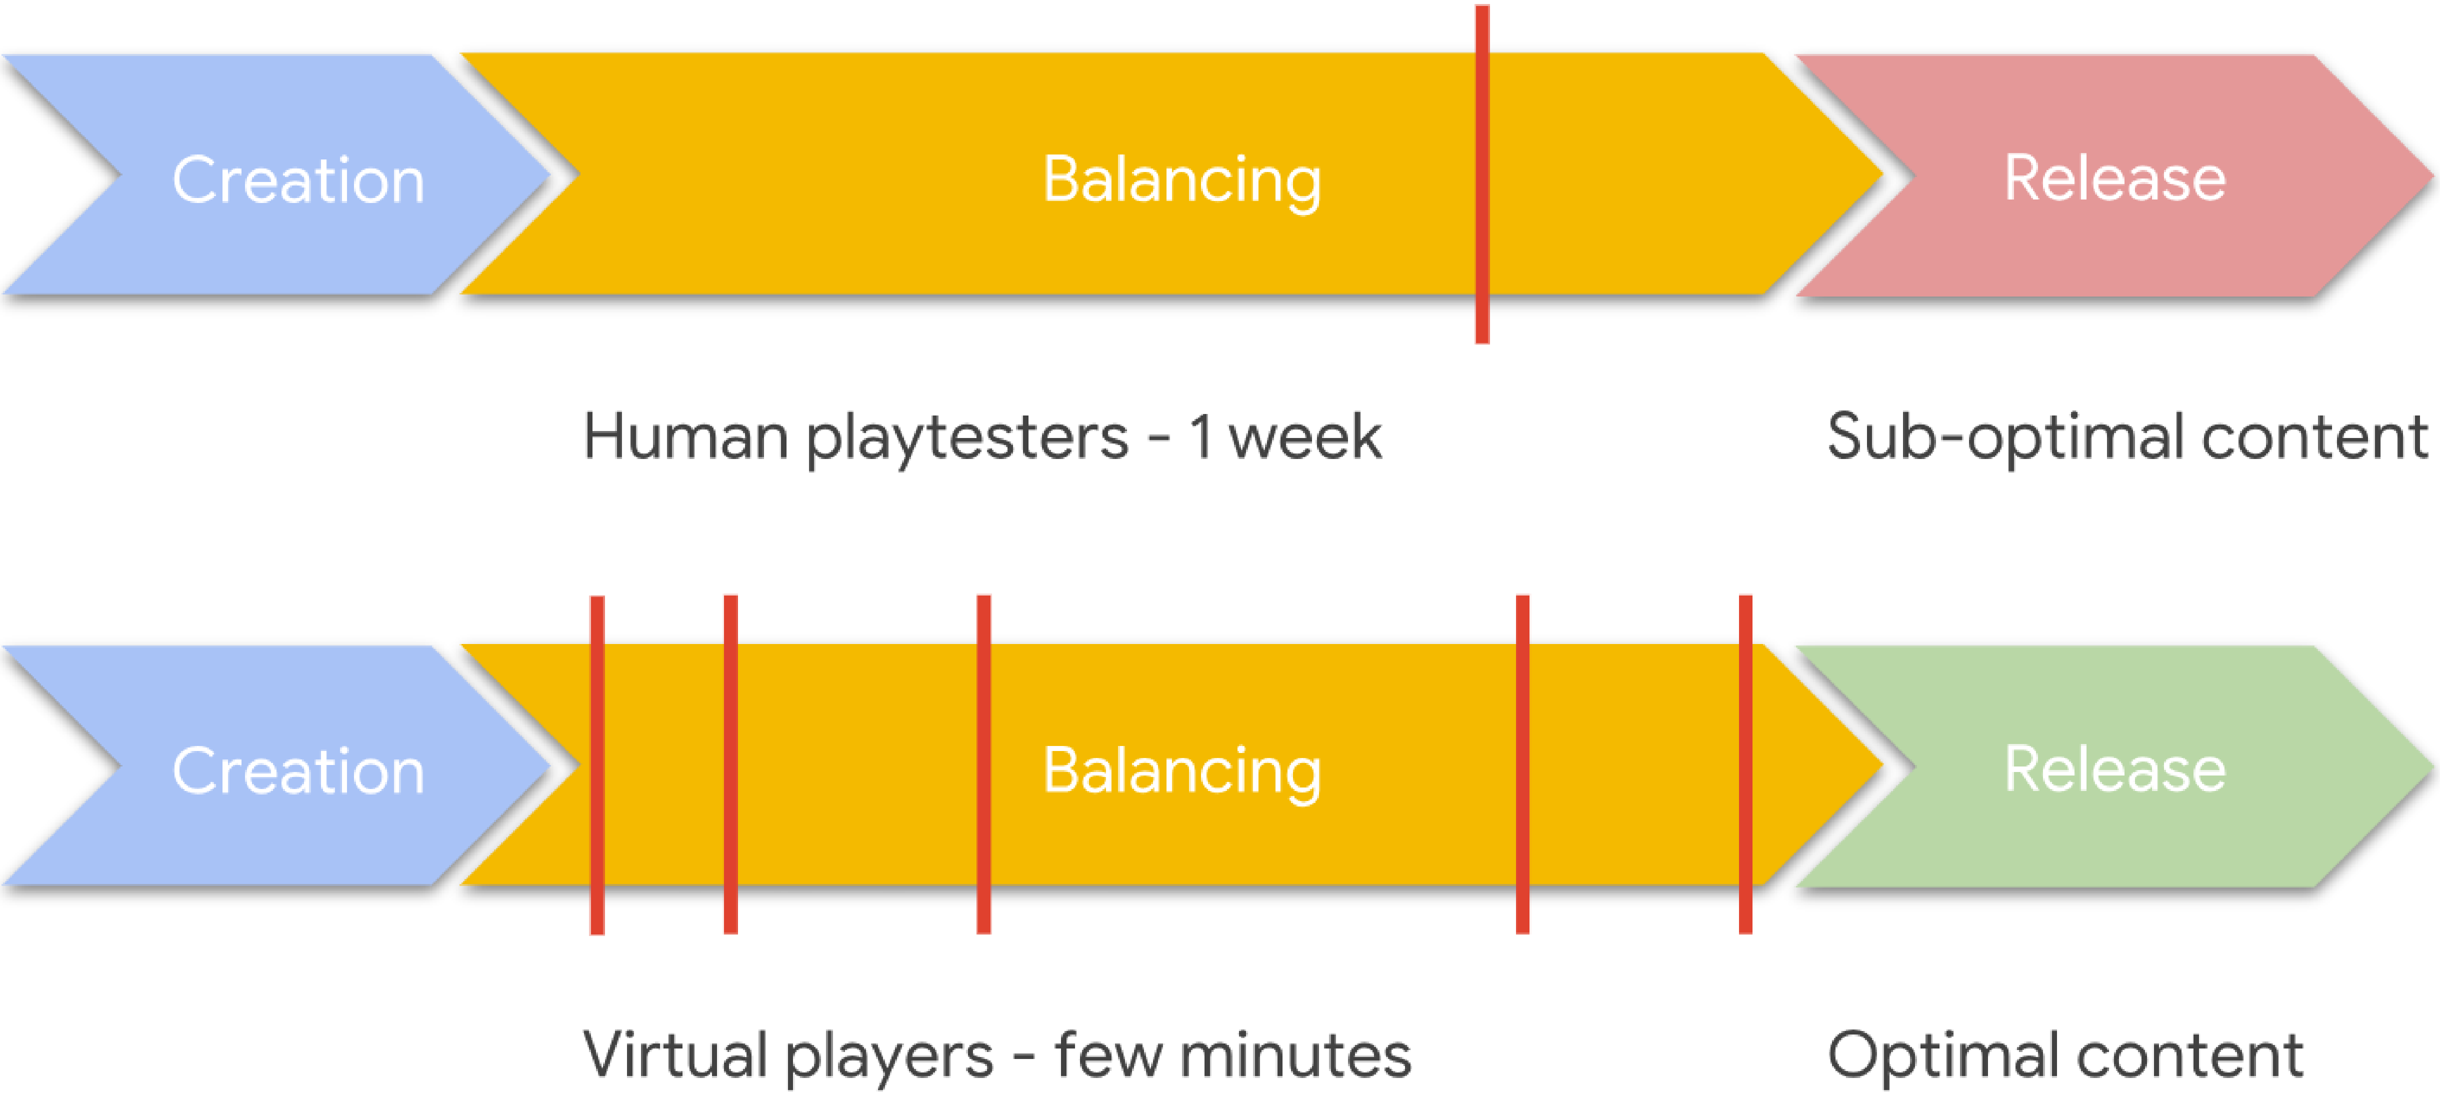
Creation
Balancing
Release
Human playtesters - 1 week
Sub-optimal content
Creation
Balancing
Release
Virtual players - few minutes
Optimal content

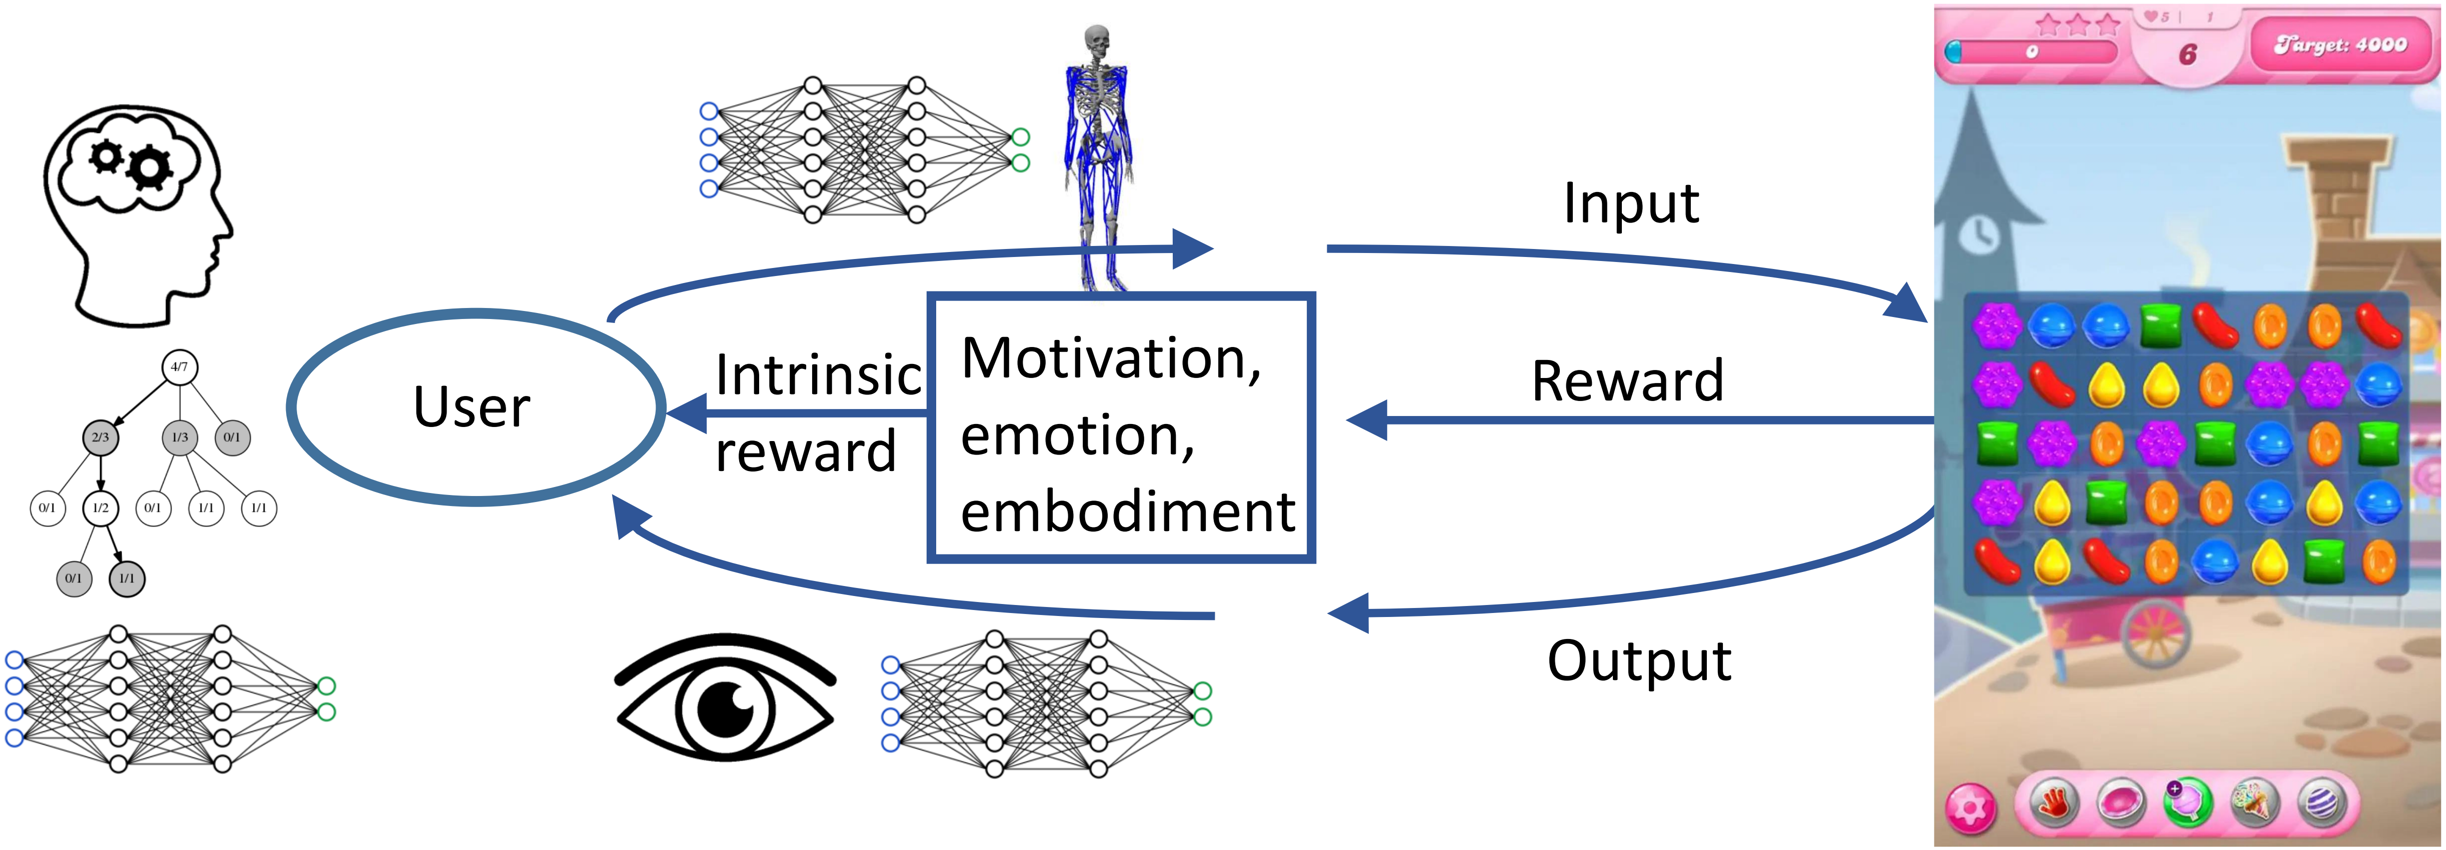
User
Intrinsic reward
Motivation, emotion, embodiment
Input
Reward
Output
Target: 4000
6

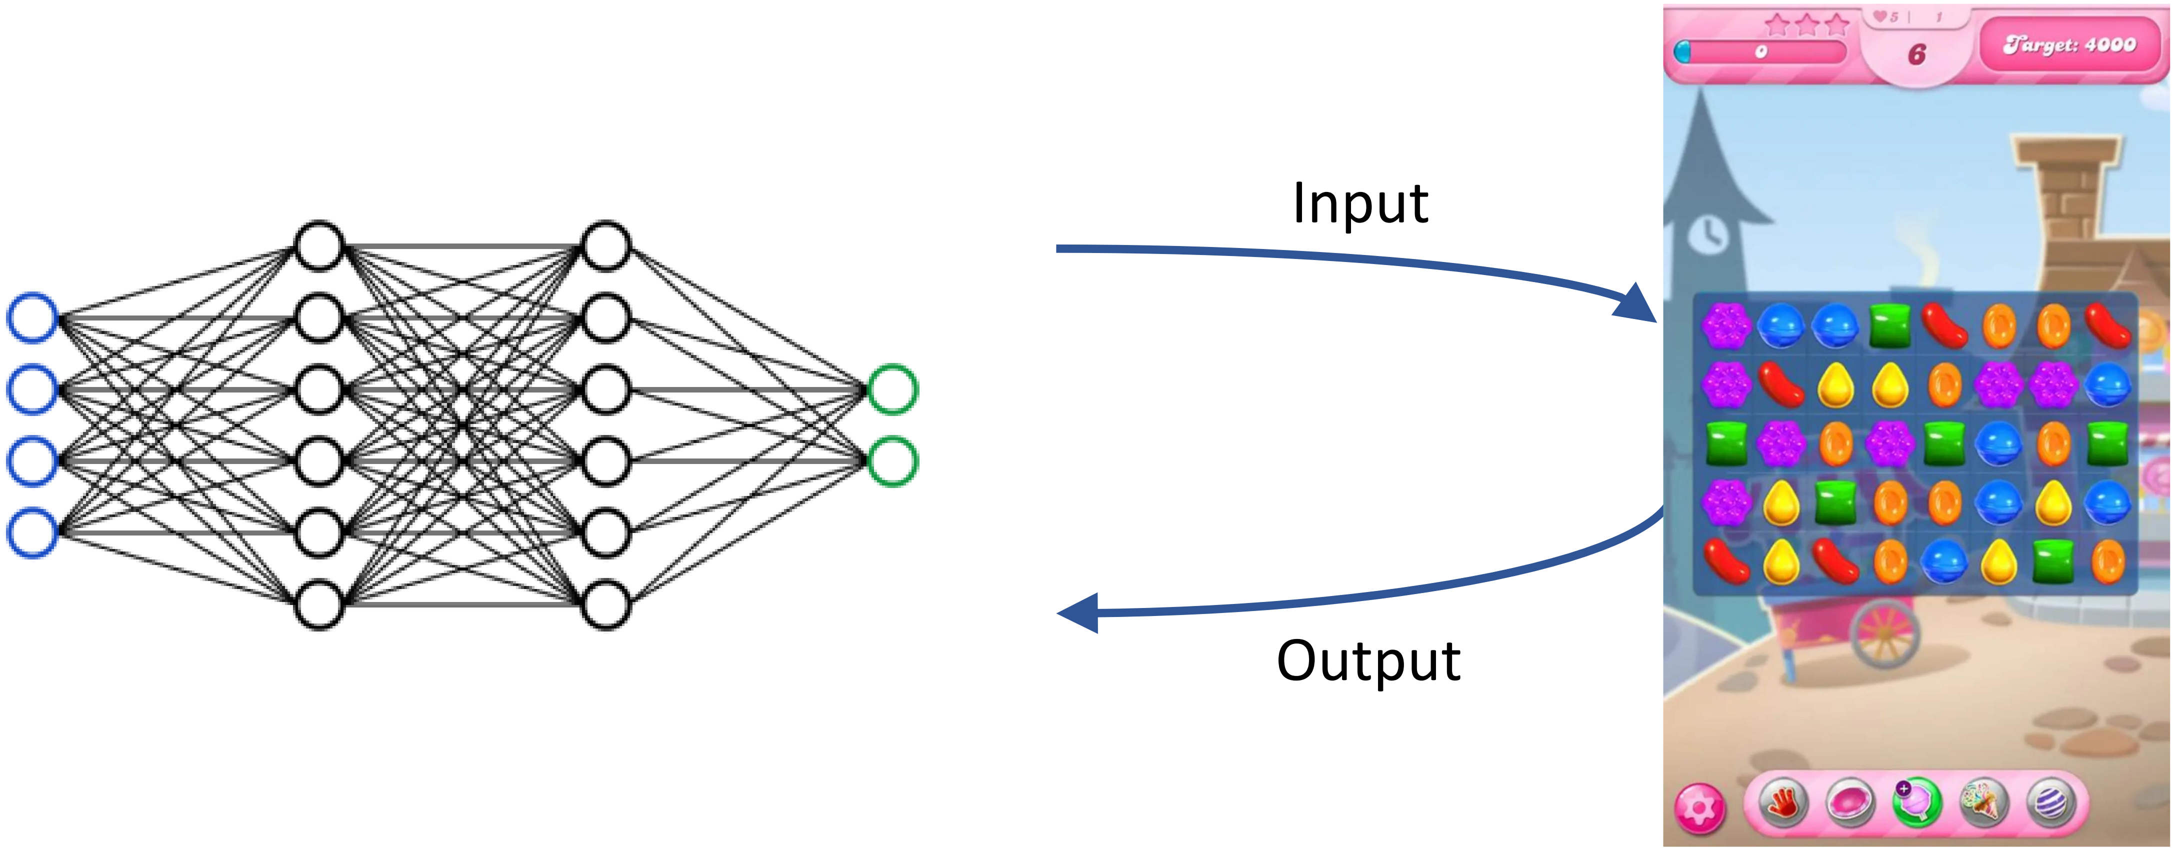
Input
Output
0
6
Target: 4000

Predicting Player Experience Without the Player
An Exploratory Study

Christian Guckelsberger[1,*], Christoph Salge[2,3], Jeremy Gow[1], and Paul Cairns[4]

[1]Computational Creativity Group, Goldsmiths, University of London, London, UK
[2]Adaptive Systems Research Group, University of Hertfordshire, Hatfield, UK
[3]Game Innovation Lab, New York University, New York, US
[4]Department of Computer Science, University of York, York, UK
[*]Corresponding author. Email: c.guckelsberger@gold.ac.uk

ABSTRACT

A key challenge of procedural content generation (PCG) is to evoke a certain player experience (PX), when we have no direct control over the content which gives rise to that experience. We argue that neither the rigorous methods to assess PX in HCI, nor specialised methods in PCG are sufficient, because they rely on a human in the loop. We propose to address this shortcoming by means of computational models of intrinsic motivation and AI game-playing agents. We hypothesise that our approach could be used to automatically predict PX across games and content types without relying on a human player or designer. We conduct an exploratory study in level generation based on empowerment, a specific model of intrinsic motivation. Based on a thematic analysis, we find that empowerment can be used to create levels with qualitatively different PX. We relate the identified experiences to established theories of PX in HCI and game design, and discuss next steps.

PCG algorithms require formal guidelines about the desired content characteristics. A procedurally generated level should without doubt be *playable*, i.e. there must be a way for the player to succeed or fail, or to experience the whole content instance and not just a small part of it. Content should also be *novel* and *typical* (cf. [38]): a generated quest for instance should be different from existing quests, but still fit the game under consideration. However, nobody would care about a level, character or as a consequence even the overall game, if the content in question did not lead to a desired experience. *Player experience* (PX) substantially determines the *value* of a content instance, and consequently its acceptance and replayability (cf. [46]). The key challenge for PCG arising from this is to evoke a certain PX, when we have *no direct control over the content* which gives rise to that experience.

We argue that established methods in human computer interaction (HCI) and PCG do not address the challenge of predicting

ACM CHI PLAY 2017

Coupled Empowerment Maximisation (CEM)

- Observation: progress in games often **aligned with increase in options** and influence.
- Quantify this influence with **empowerment** [1]:

An agent's perceivable influence on the world, including other agents.

- **Coupled Empowerment Maximisation**:
  NPC selects actions to maximise its own-, while minimising / maximising player's empowerment.

(NPC = non-player character)

[1] Salge, Glackin & Polani. "Empowerment – An Introduction." Guided Self-Organization: Inception. Springer, 2014. 67-114.

[2] Guckelsberger et al. "Predicting Player Experience without the Player: An Exploratory Study." In Proc. ACM CHI'Play, 2017.

Coupled Empowerment Maximisation (CEM)

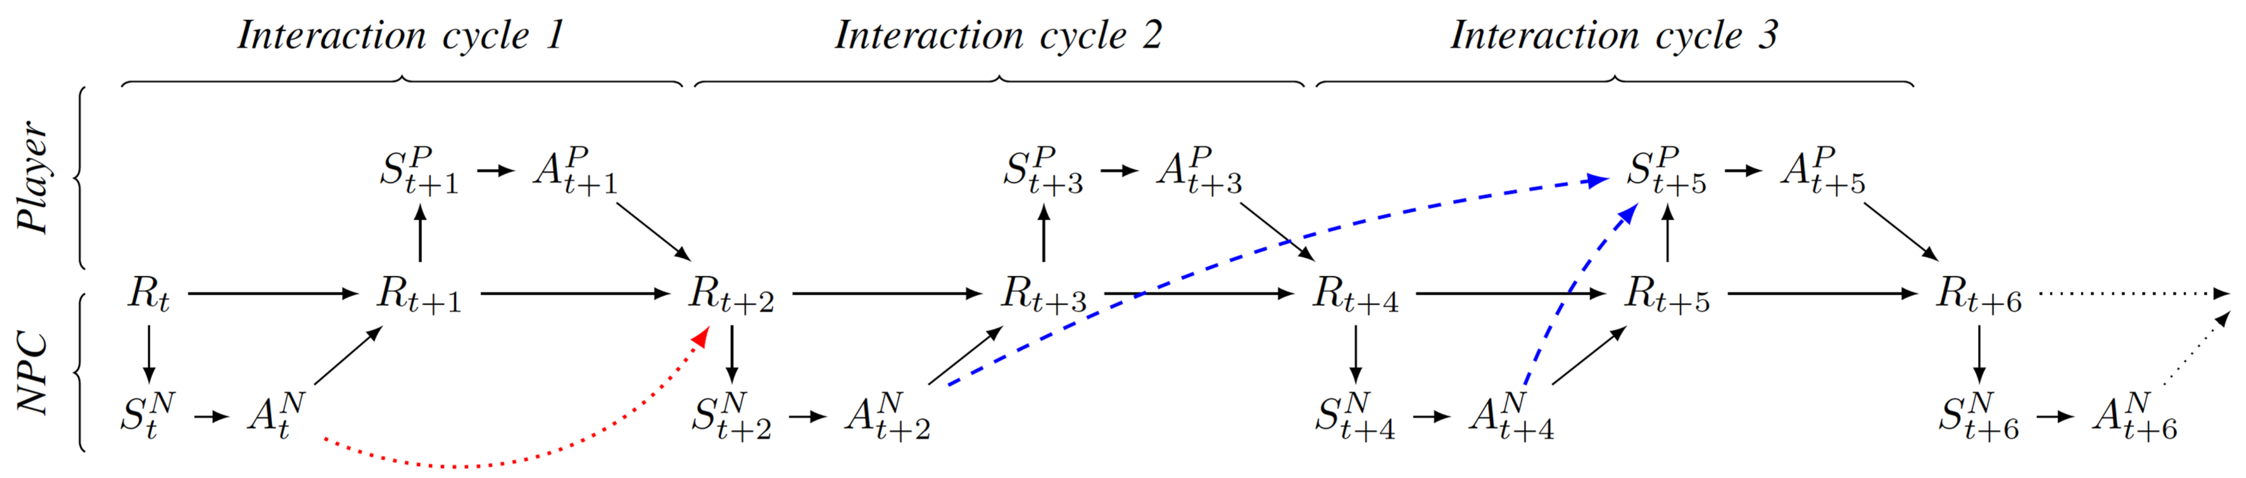

- Combining 3 types of empowerment (R: world state, S: sensor, A: actuator):
  1. **NPC empowerment**: NPC's influence on its own sensor n steps ahead - $\mathfrak{E}^N(s_{t+2})$
  2. **Player empowerment**: Player's influence on its sensor n steps ahead - $\mathfrak{E}^P(s_{t+1})$
  3. **NPC-Player Transfer empowerment**: NPC's influence on player's future sensor - $\mathfrak{E}^T(s_{t+2})$
- **CEM Policy:** $\pi^{cem}(a_t|s_t) = \arg\max_{a_t} \left(\alpha_N\, \mathbb{E}[\mathfrak{E}^N]_{a_t} + \alpha_P\, \mathbb{E}[\mathfrak{E}^P]_{a_t} + \alpha_T\, \mathbb{E}[\mathfrak{E}^T]_{a_t}\right)$

CEM-Induced Support (Companion NPCs)

Operational proximity　Evade　Protect

Cf. Guckelsberger, Salge and Colton. "Intrinsically Motivated General Companion NPCs via Coupled Empowerment Maximisation ". In Proc. IEEE Conf. Computational Intelligence and Games, 2016.

CEM-Induced Adversarialness (Enemy NPCs)

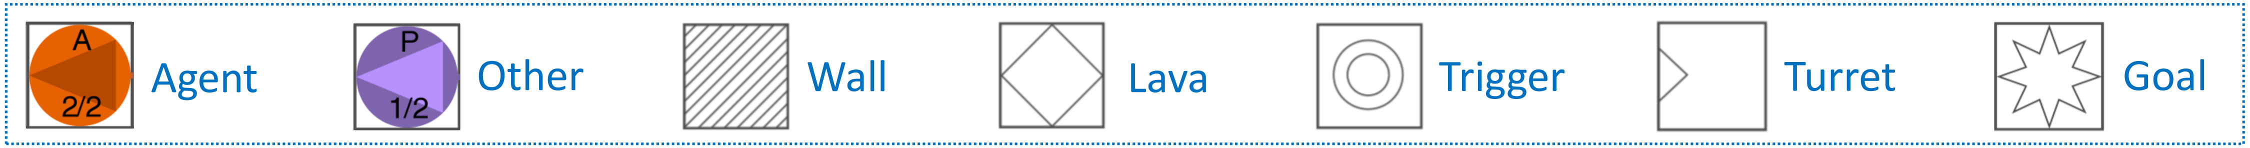

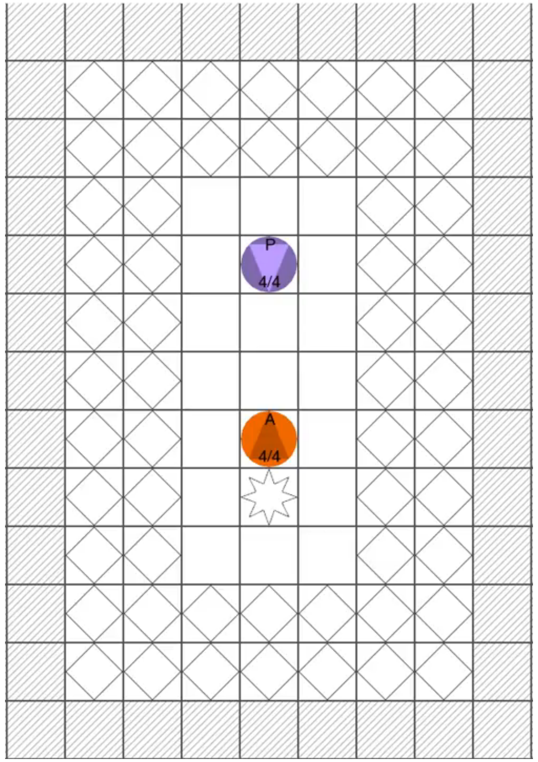

Push into lava

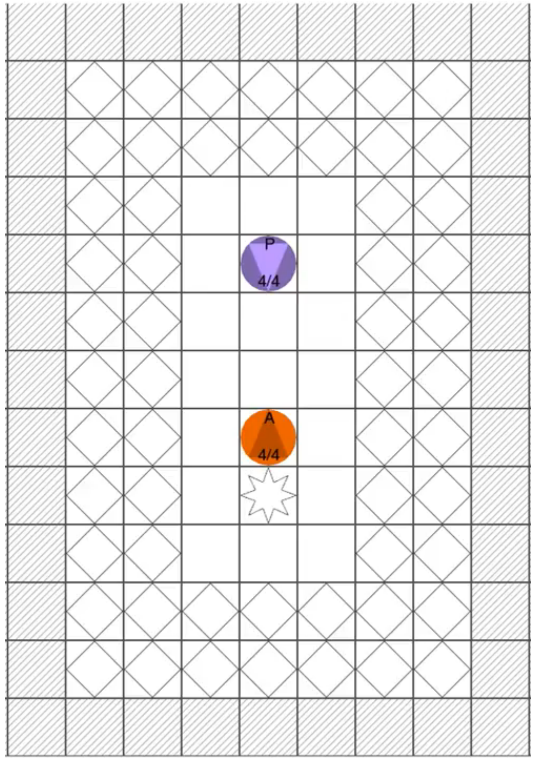

Escape over lava

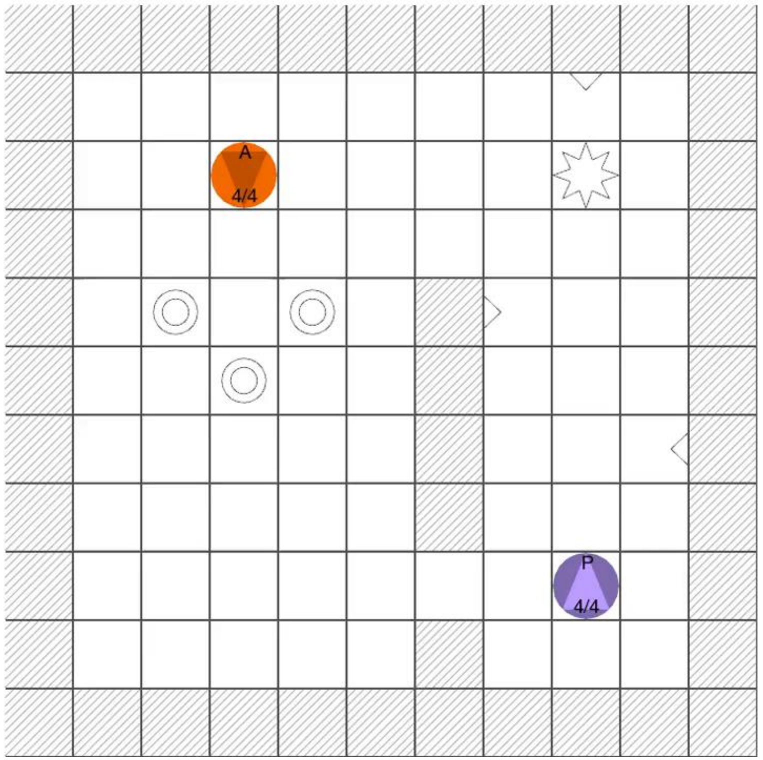

Strike from a distance

Cf. Guckelsberger, Salge, and Togelius. "New And Surprising Ways to Be Mean. Adversarial NPCs with Coupled Empowerment Minimisation". In Proc. IEEE Conf. Computational Intelligence and Games, 2018.

Simulating curiosity

Bellemare et al. 2016 (Google DeepMind)

Simulating curiosity

Bellemare et al. 2016 (Google DeepMind)

Intrinsic reward = satisfaction of psychological needs

COVER FEATURE **OUTLOOK**

User Interface Design with Combinatorial Optimization

Antti Oulasvirta, Aalto University

Optimization methods have revolutionized almost every field of engineering design, so why not user interface design? The author reviews progress and challenges in model-driven UI optimization, in which an optimizer utilizes predictive models of human perception, behavior, and experience to anticipate users' responses to computer-generated designs.


F1
F2
F3
F4
1 à §
2 é
3 è
4 &
A
Z £
E €
R ®
T
Y
D $
F ¤
G µ
H
B
N
V
C ç
S ß
Q
X
W
Tab
Caps Lock
Alt

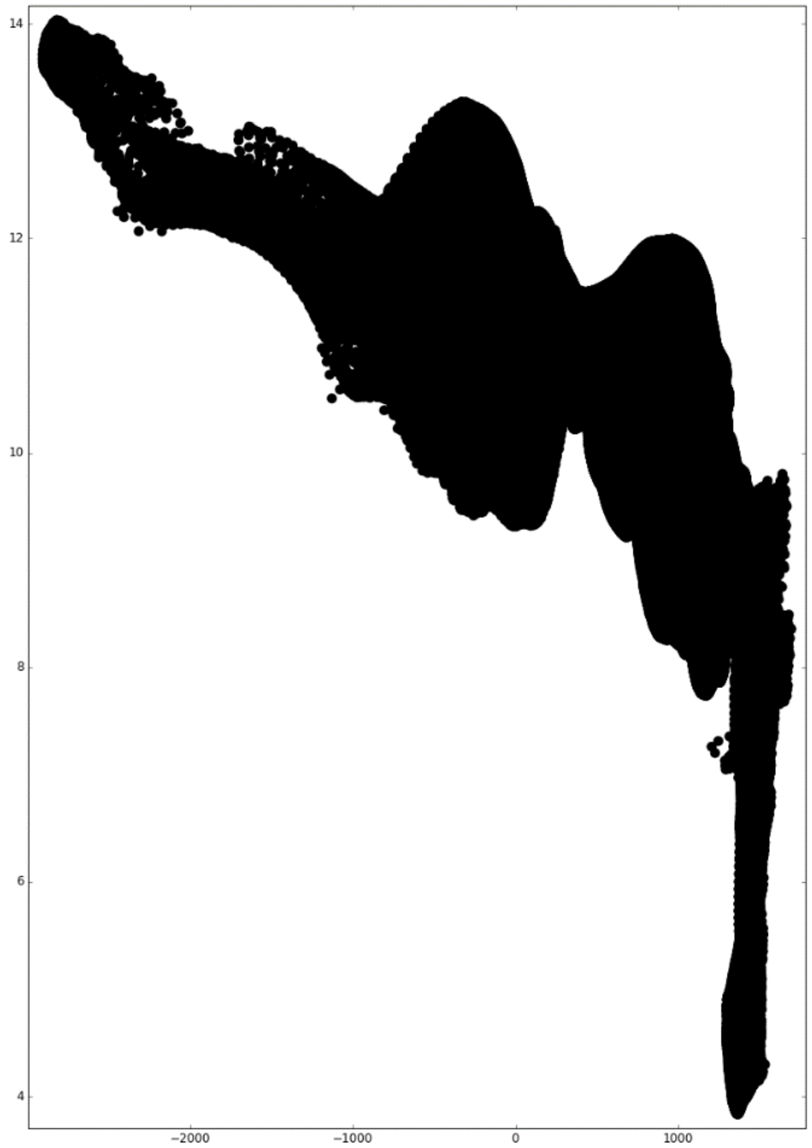

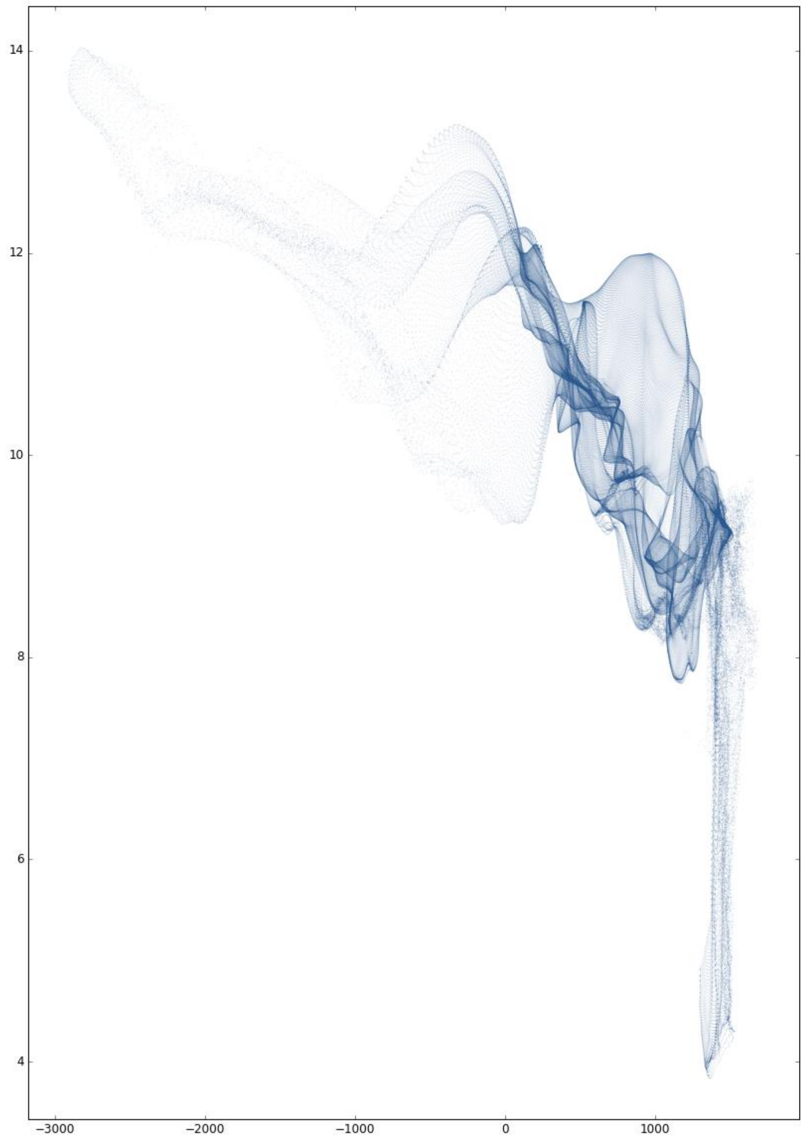

Micallef, Palmas, Oulasvirta, Weinkauf 2017: Scatterplot optimization

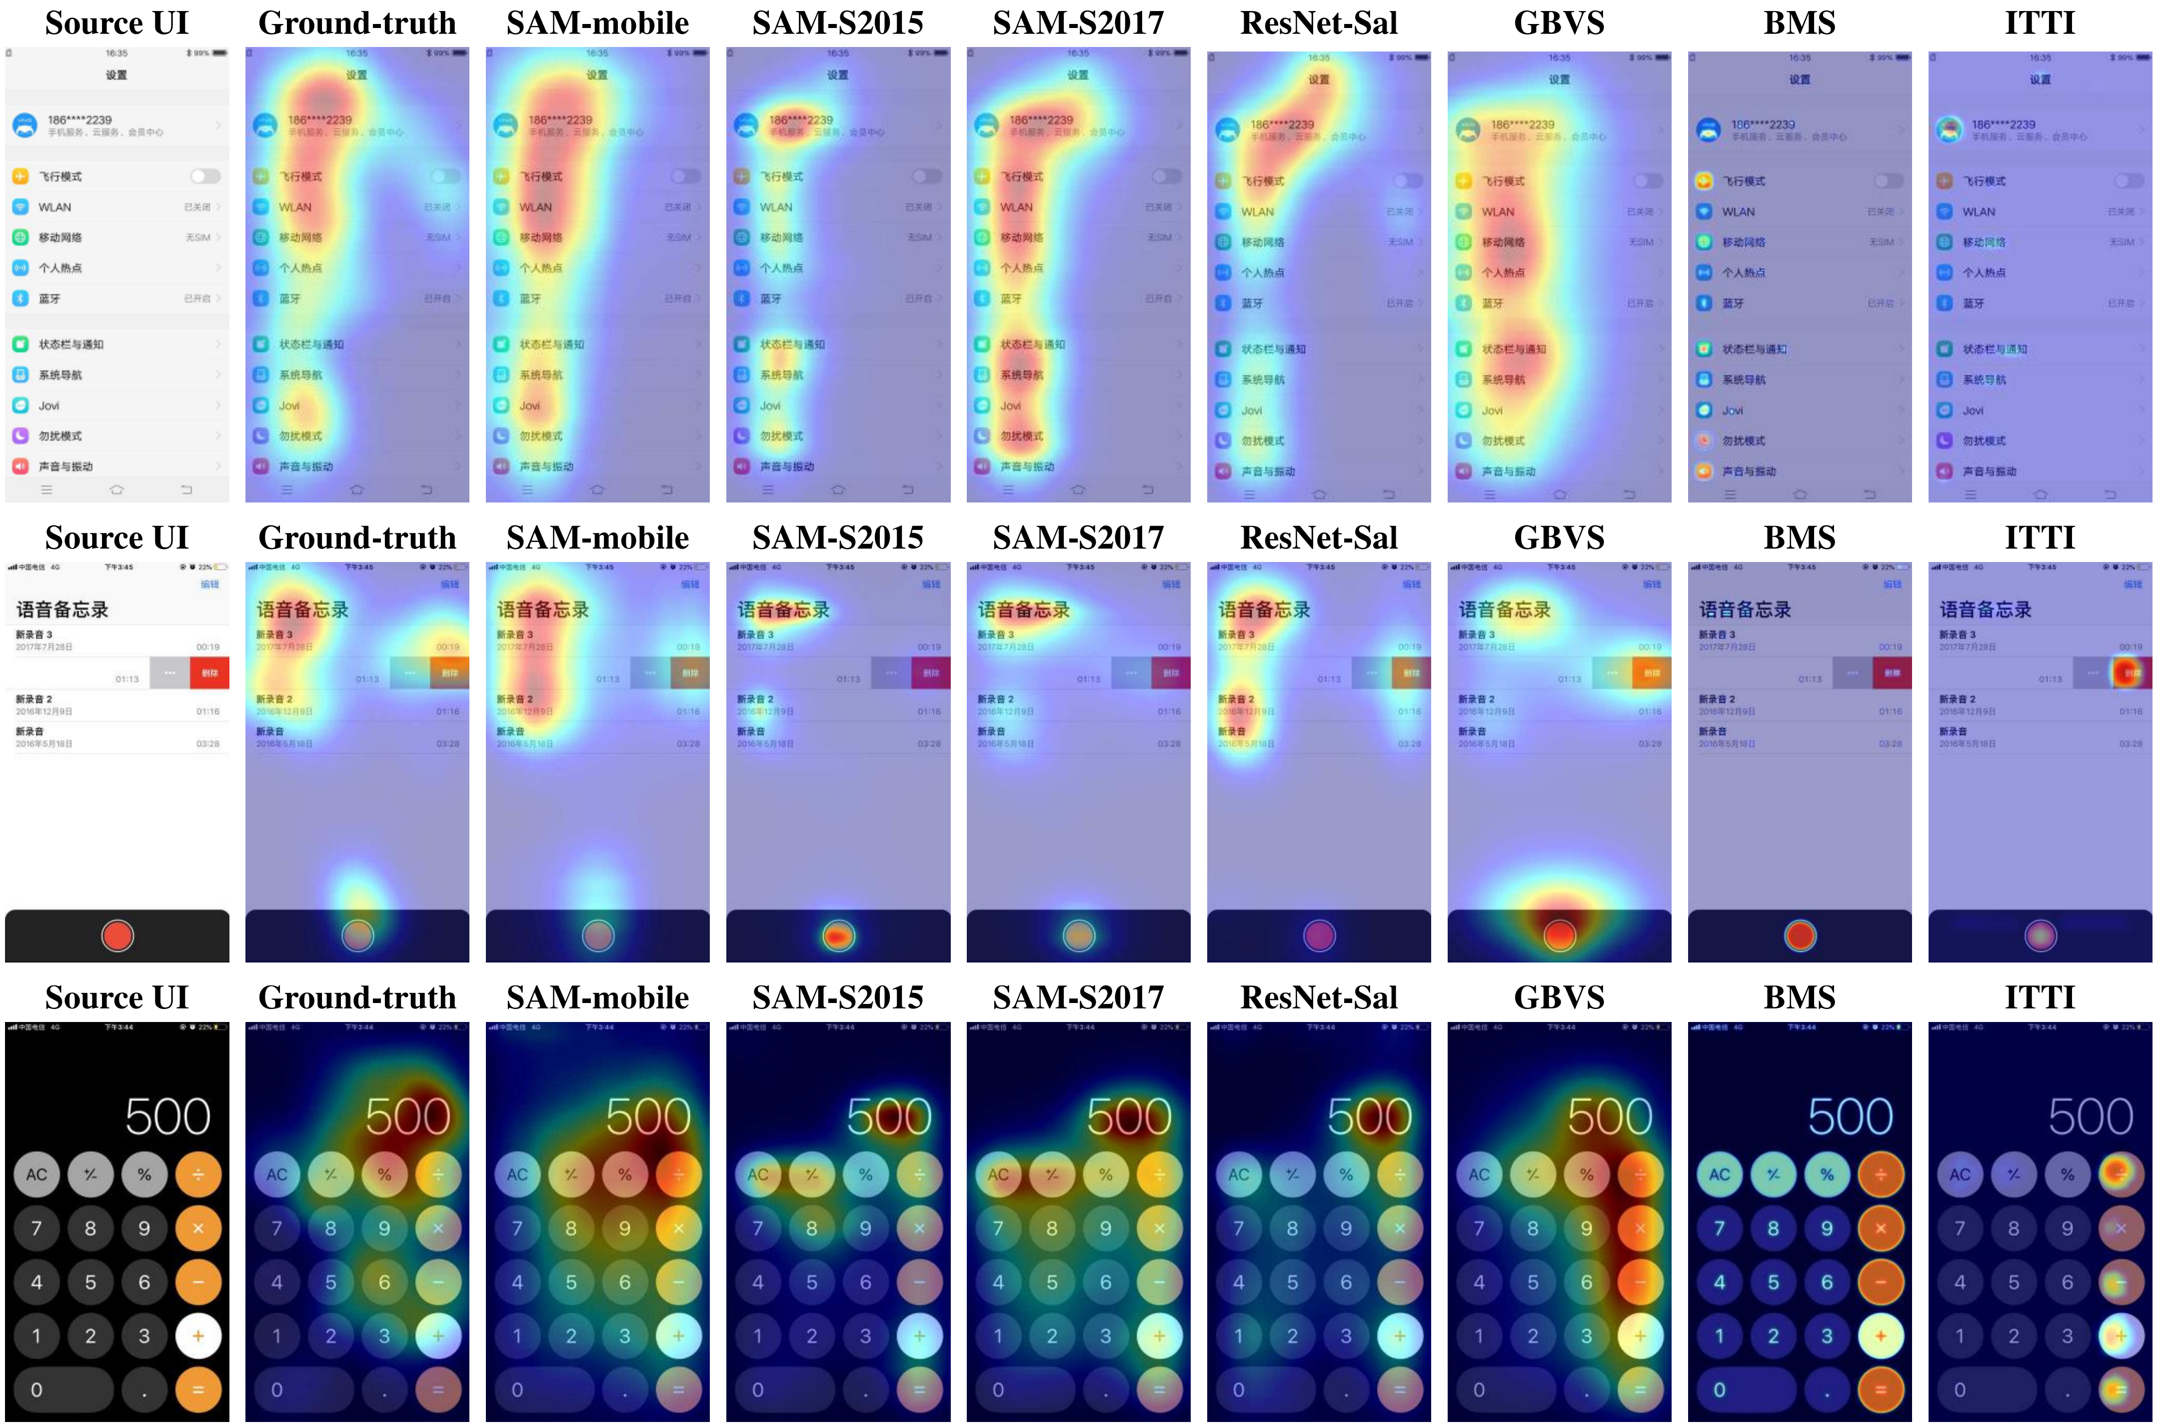
Source UI
Ground-truth
SAM-mobile
SAM-S2015
SAM-S2017
ResNet-Sal
GBVS
BMS
ITTI
Source UI
Ground-truth
SAM-mobile
SAM-S2015
SAM-S2017
ResNet-Sal
GBVS
BMS
ITTI
Source UI
Ground-truth
SAM-mobile
SAM-S2015
SAM-S2017
ResNet-Sal
GBVS
BMS
ITTI

My group: Simulating VR games & interfaces

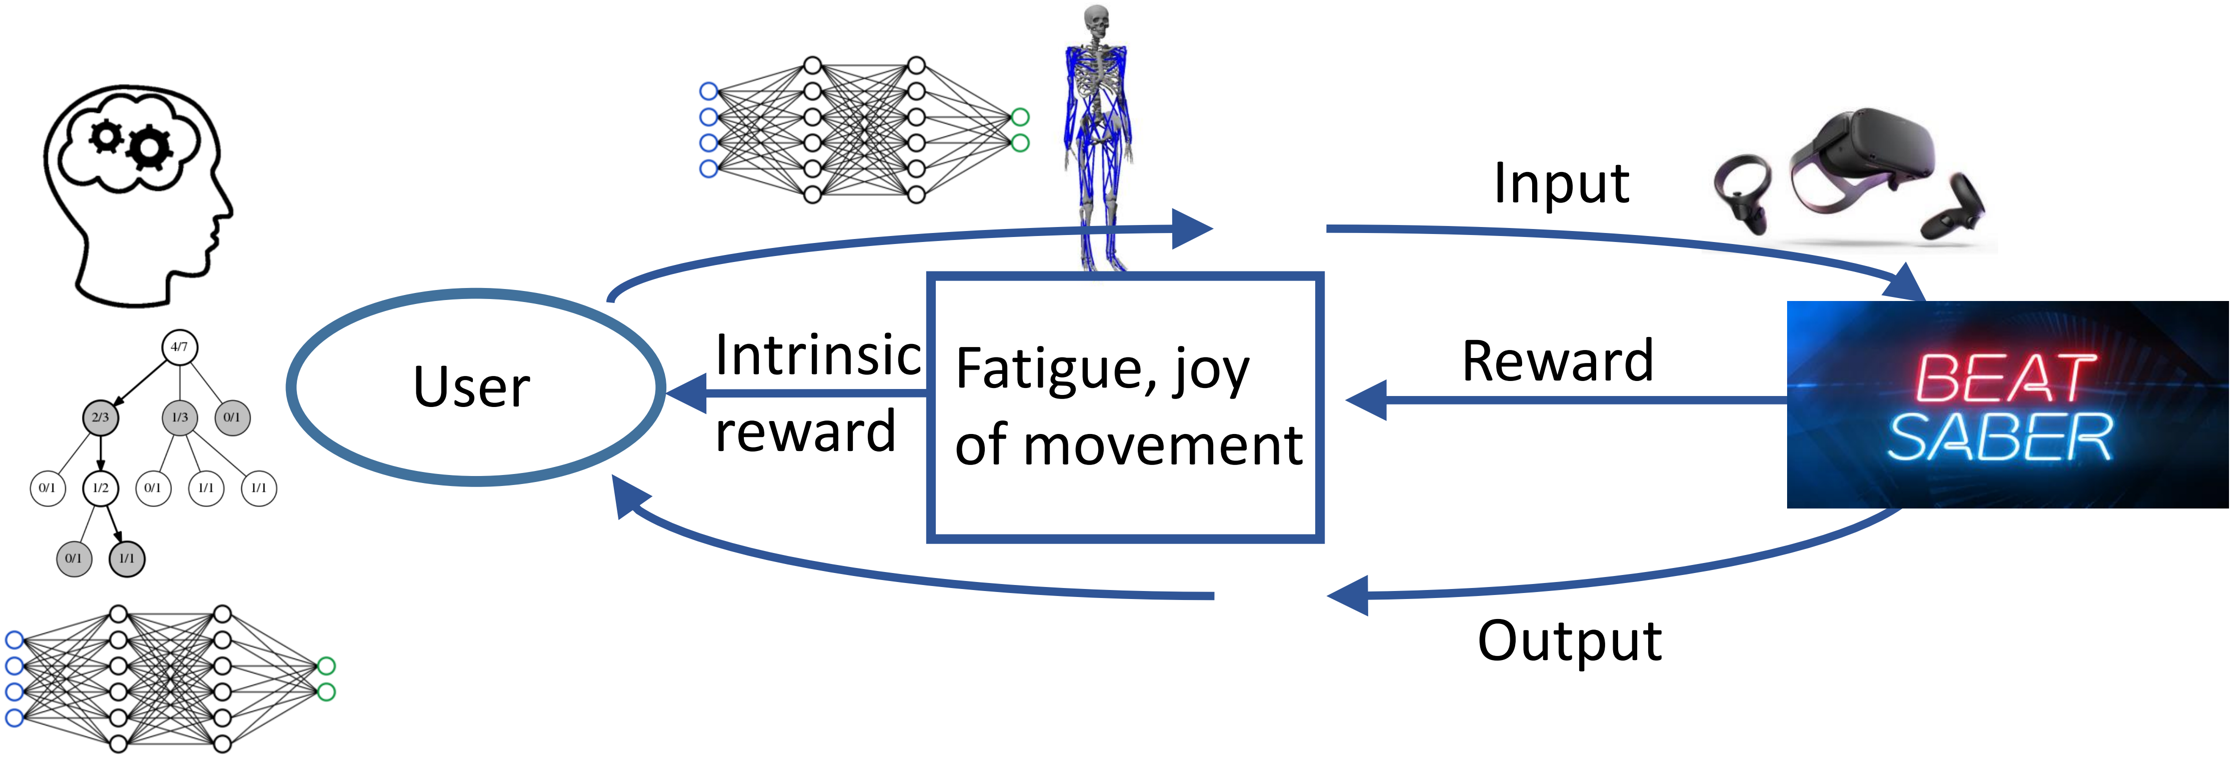

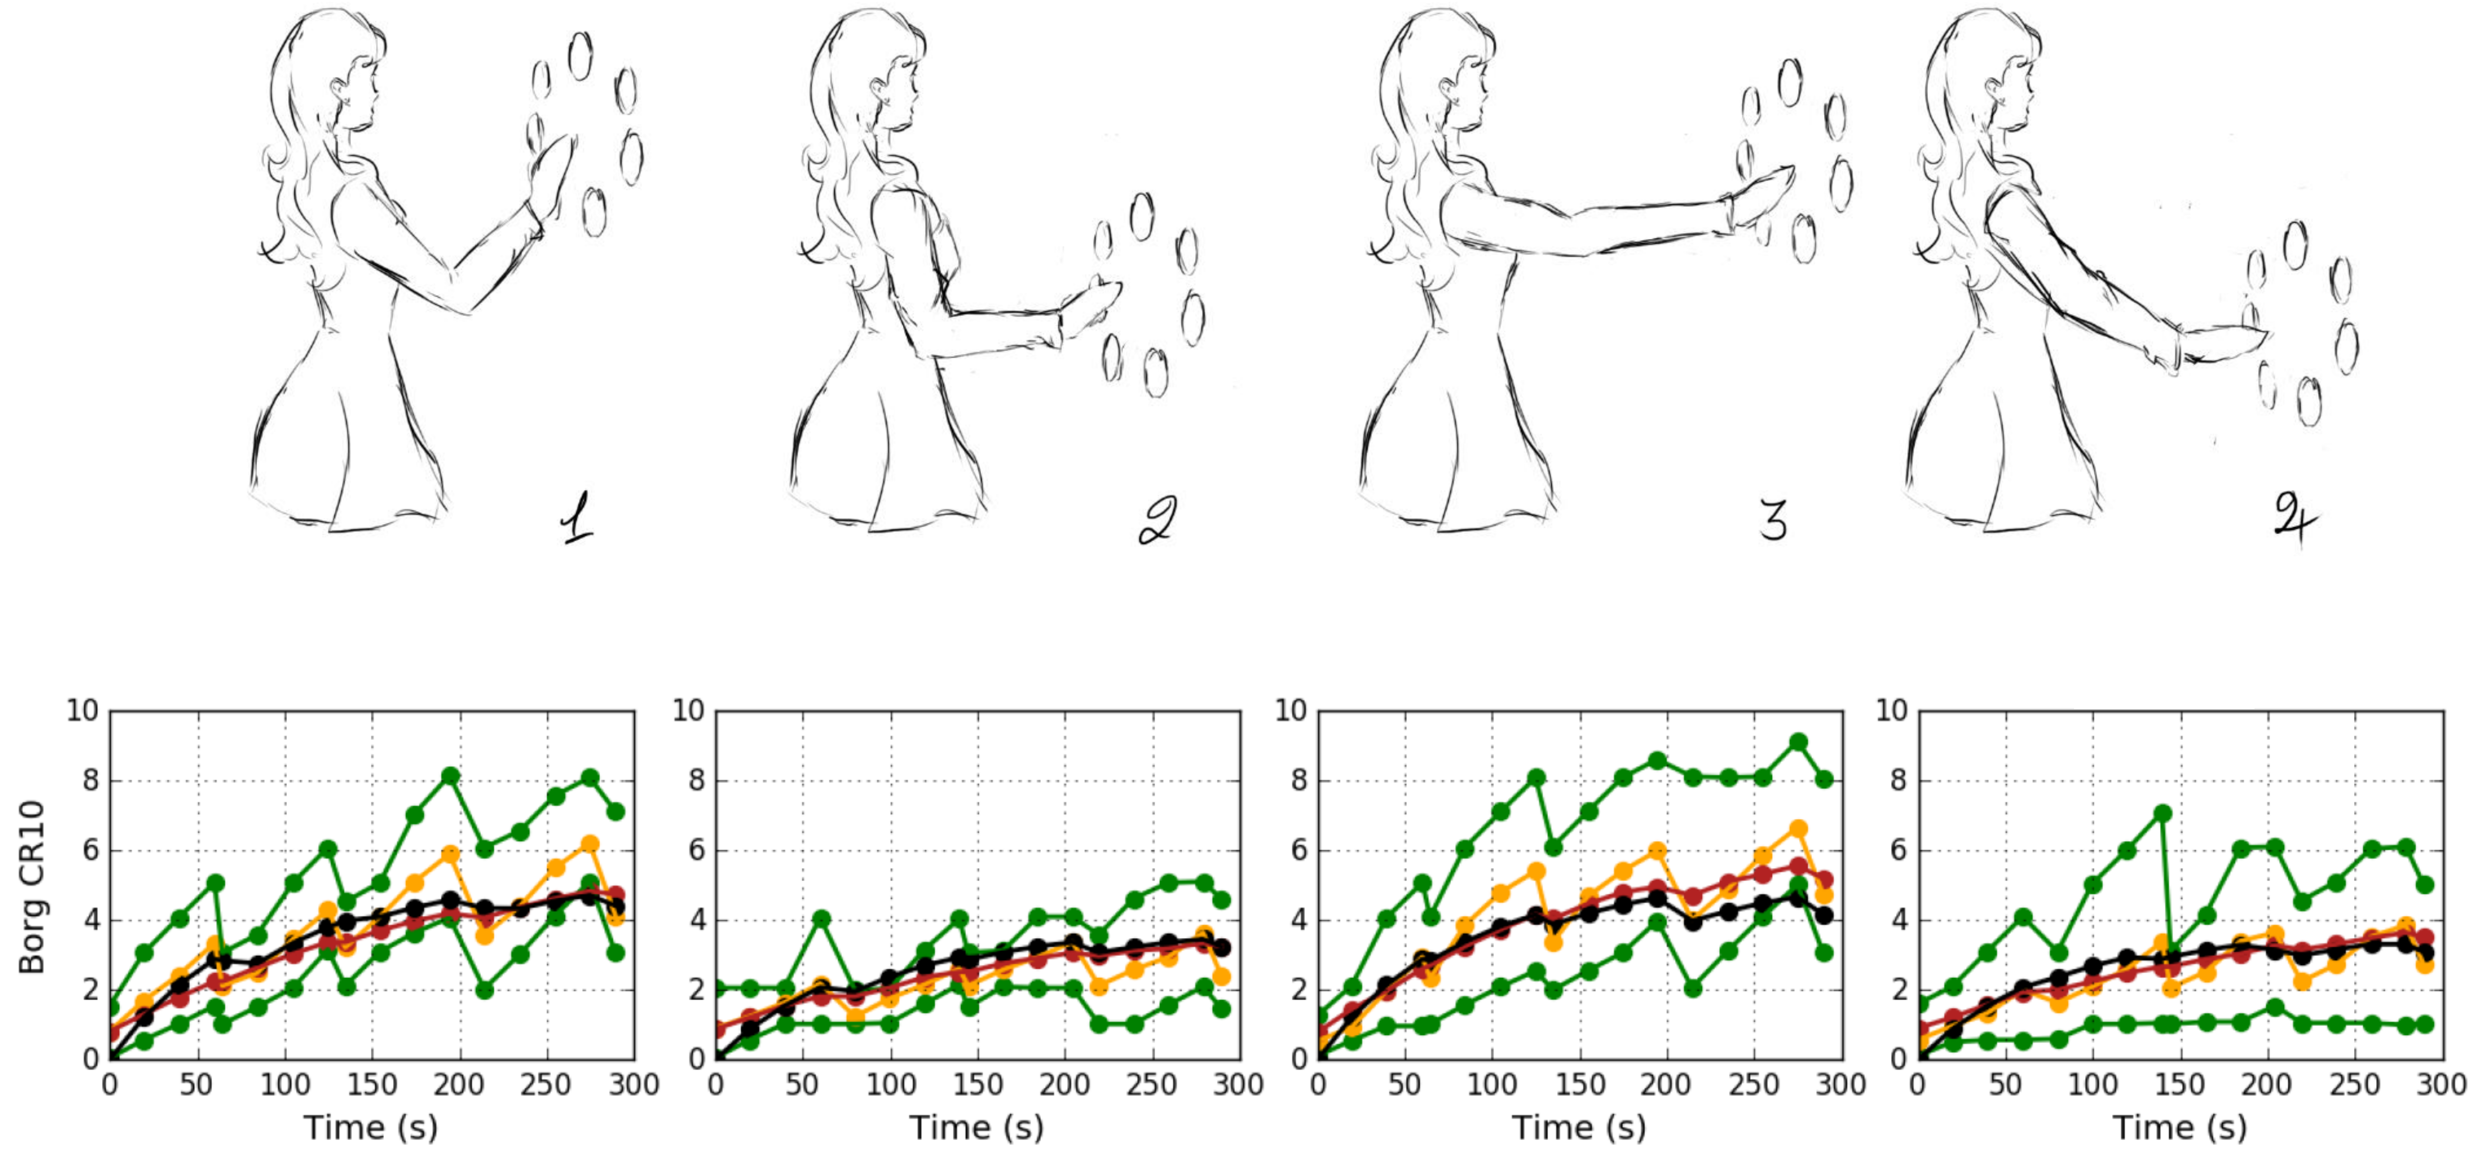

Figure 8. Results of predicting the Borg CR10 rating. Green: Upper/lower bound of ground truth. Yellow: Average of ground truth. Red: Average 3CC estimate of ground truth computed using motion capture data [33]. Black: Our simulation-based average 3CC-r estimate. Our simulation model yields similar modeling accuracy as [33], but does not require motion capture data.

Horizons, hard problems

WHAT IS THE VPH INSTITUTE

The Virtual Physiological Human Institute for Integrative Biomedical Research, in short VPH Institute, is an international non-profit organisation incorporated in Belgium, whose mission is to ensure that the Virtual Physiological Human is fully realised, universally adopted, and effectively used both in research and clinic.

WHAT IS THE VPH?

The Virtual Physiological Human (VPH), also identified with the word "in silico medicine" is the field that encompasses the use of individualised physiology based computer simulations in all aspects of the prevention, diagnosis, prognostic assessment, and treatment of a disease and development of a biomedical product.

Maximize aesthetic impression, media virality, wearability

Project Runway challenge: Maximize judge scores,
constrained by budget < *X*, all materials must be from a candy store

Maximize the feeling of tranquil awe.

Make the player feel grief, but still have a somehow satisfying experience

Though-provoking games: Maximize reflection.
(user experience is not limited to product use!)

“A Game that Makes You Question...” Exploring the Role of Reflection for the Player Experience

Elisa D. Mekler[1], **Ioanna Iacovides**[2], **Julia Ayumi Bopp**[1]
[1]Faculty of Psychology, University of Basel; [2]Department of Computer Science, University of York
elisa.mekler@unibas.ch, jo.iacovides@york.ac.uk, j.bopp@unibas.ch

ABSTRACT

Reflection is a core design outcome for HCI, and recent work has suggested that games are well suited for prompting and supporting reflection on a variety of matters. However, research about what sorts of reflection, if any, players experience, or what benefits they might derive from it, is scarce. We report on an interview study that explored when instances of reflection occurred, at what level players reflected on their gaming experience, as well as their reactions. Our findings revealed that many players considered reflection to be a worthwhile activity in itself, highlighting its significance for the player experience beyond moment-to-moment gameplay. However, while players engaged in reflective description and dialogic reflection, we observed little to no instances of higher-level transformative and critical reflection. We conclude with a discussion of the value and challenges inherent to evaluating reflection on games.

Yet there is much to gain from studying reflection in this context. First, it may add to our understanding of player experience, especially with regards to how games may sustain engagement beyond the moment-to-moment experience of gameplay. Second, reflection is a crucial component of learning [7, 14, 41], where already a substantial body of work considers how games and gaming practice support learning [24, 46, 48]. Third, recent work has explored games to promote thought-provoking ‘serious experiences’ [27, 37] to raise awareness or persuade. Some have pointed towards games as a means to facilitate transformative reflection [44], which could give way to attitudinal and behavioral change [19]. A better understanding of reflection in games may thus serve to inform the design and evaluation of such serious experiences. Finally, reflection may also make for richer aesthetic gaming experiences [10], as exemplified by design approaches such as ludic engagement [23] and slow technology [25].

ACM CHI PLAY 2018


Simultaneously maximize simplicity of geometry,
curious interest and engagement in movement play

Satisfying and meaningful gameplay in the absence of goals and rewards, in a malleable multiplayer world

Minecraft E3 trailer, https://www.youtube.com/watch?v=Nq_Q77bJ3H0

Motivation

- Designing games etc. is slow, expensive, and success not guaranteed
- Reason: need to prototype and test with users

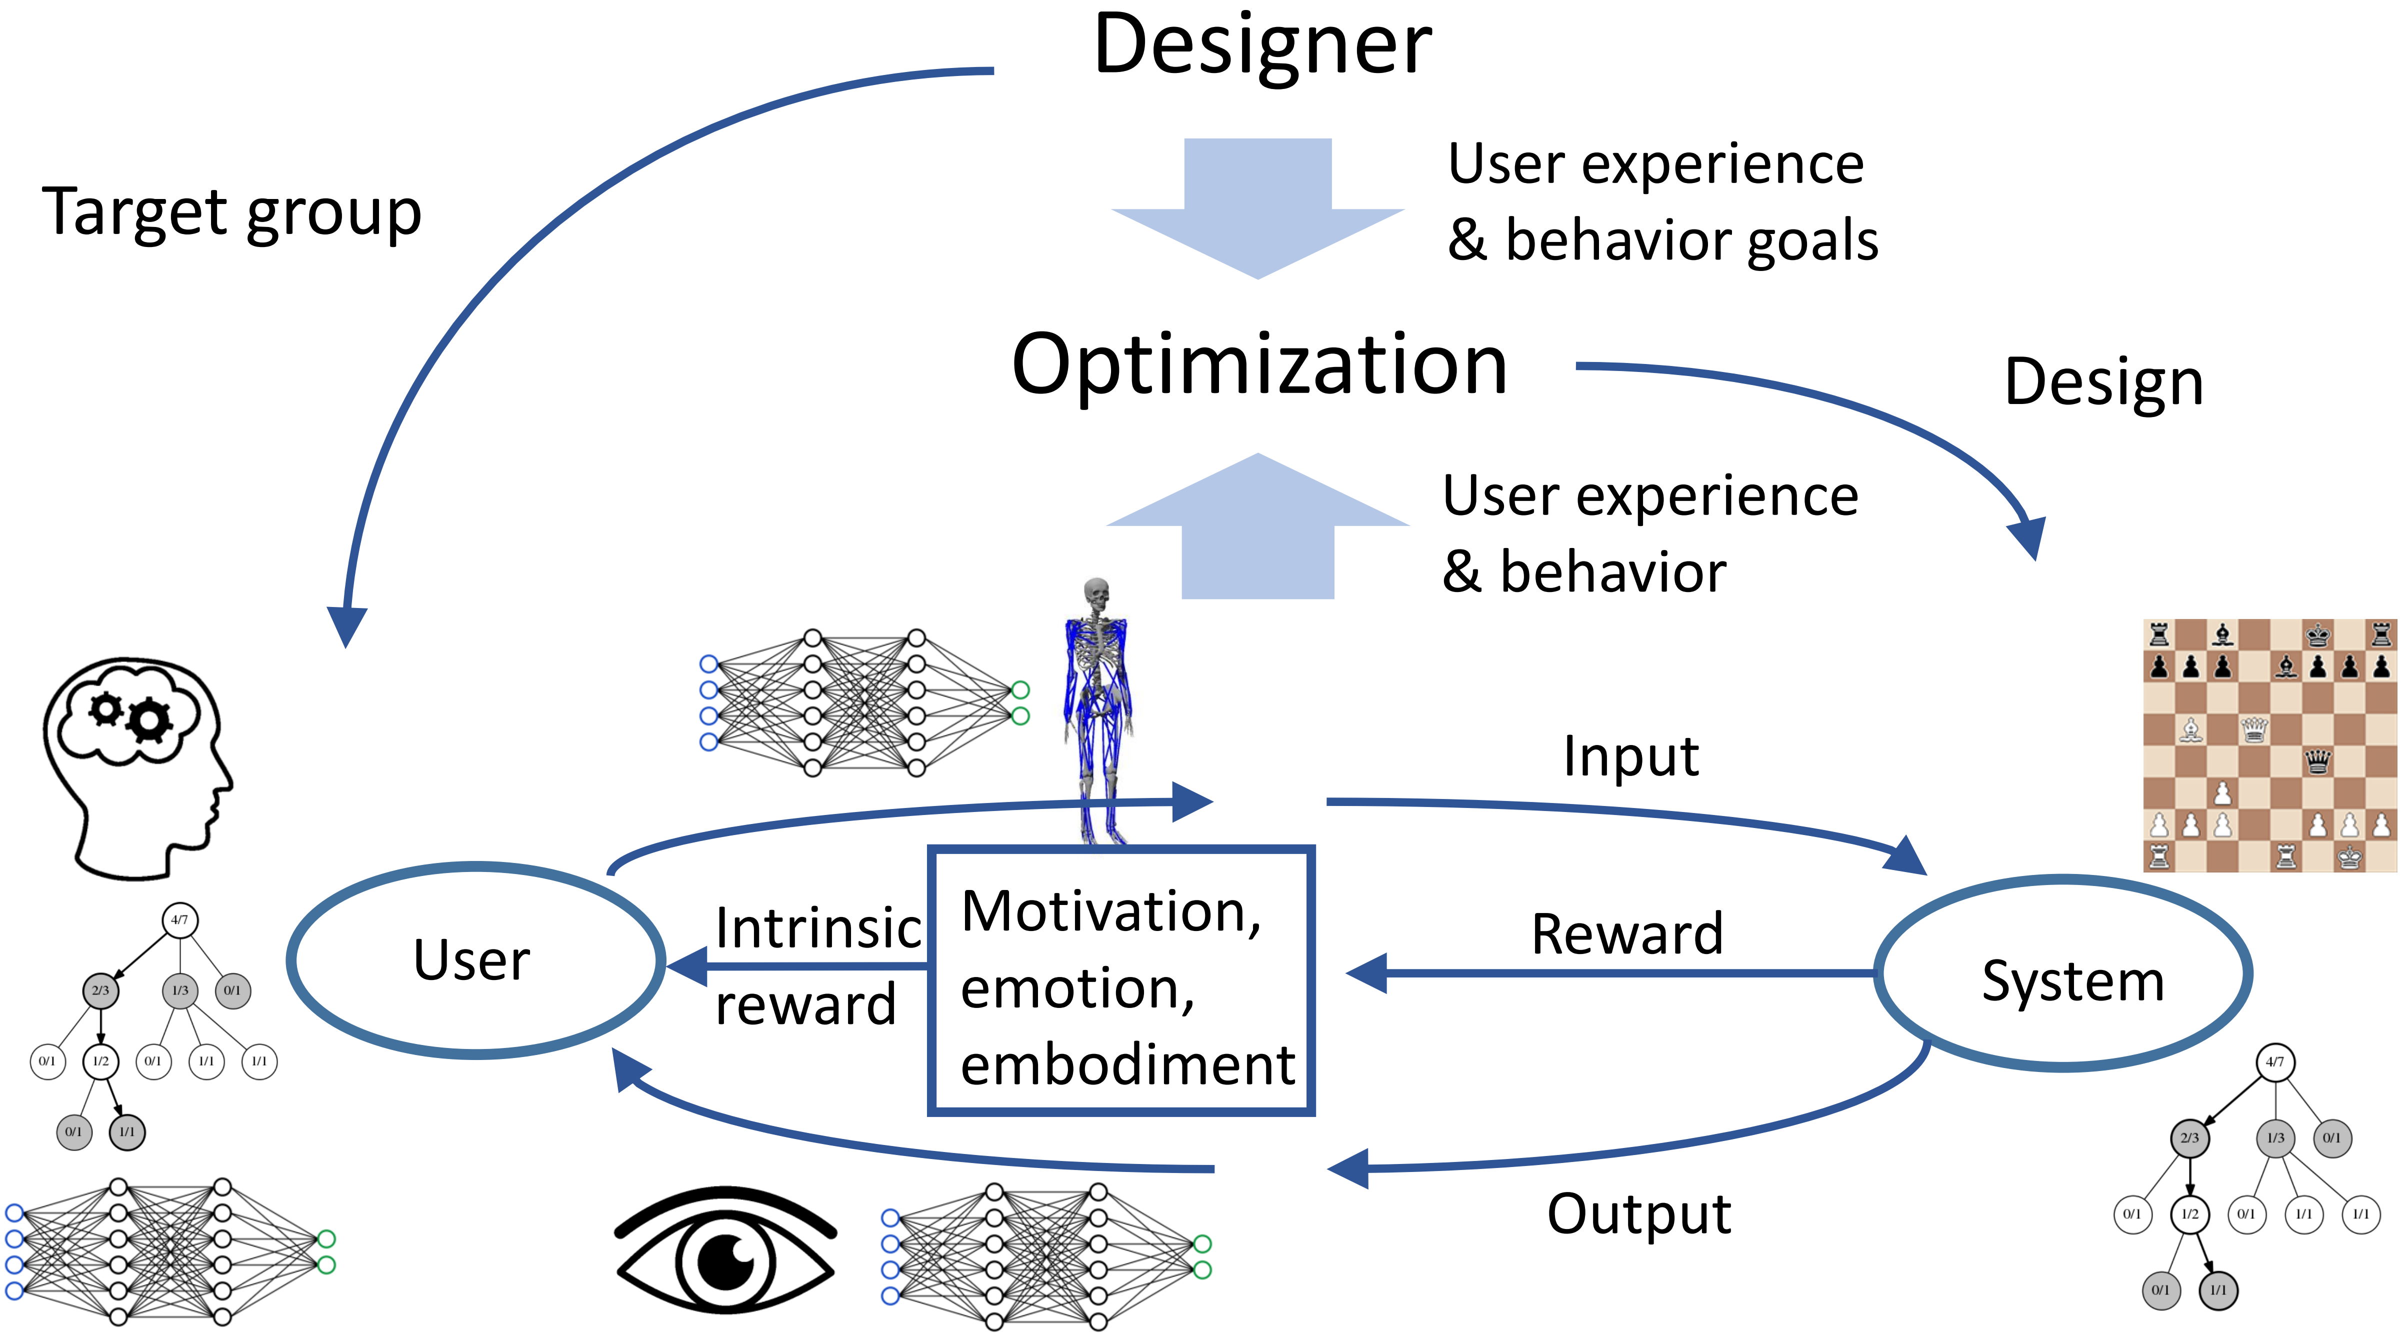
Designer
Target group
User experience
& behavior goals
Optimization
Design
User experience
& behavior
Input
Intrinsic
reward
Motivation,
emotion,
embodiment
Reward
User
System
Output

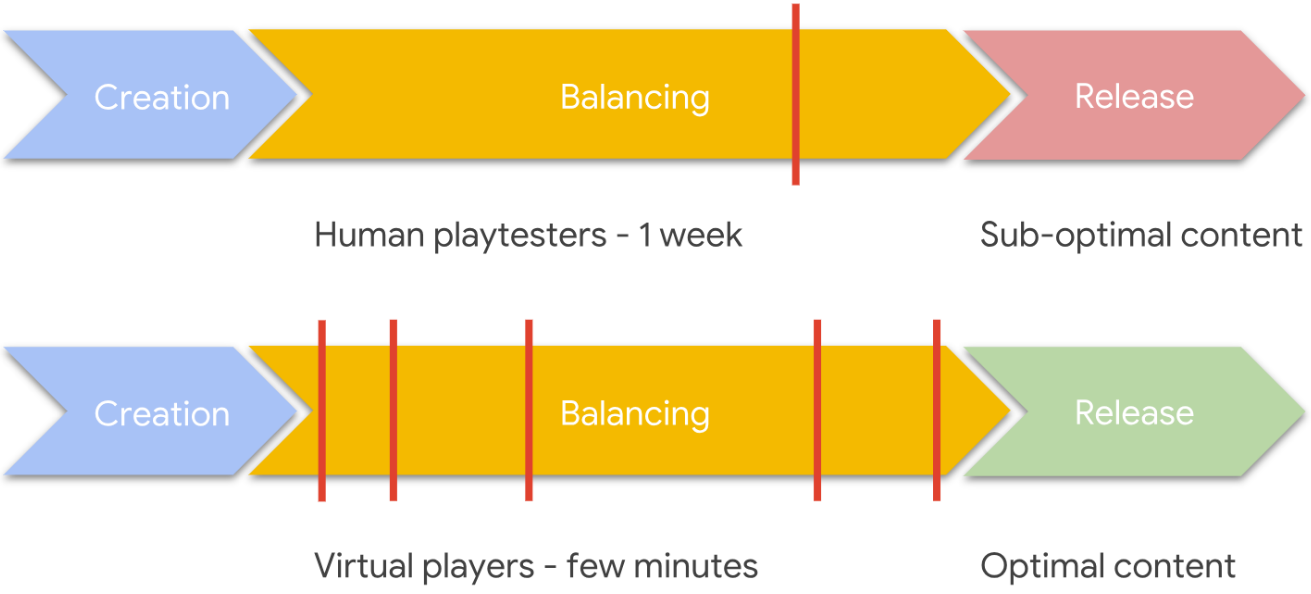
Creation
Balancing
Release
Human playtesters - 1 week
Sub-optimal content
Creation
Balancing
Release
Virtual players - few minutes
Optimal content


THANK YOU
FOR PLAYING
THAT DRAGON,
CANCER

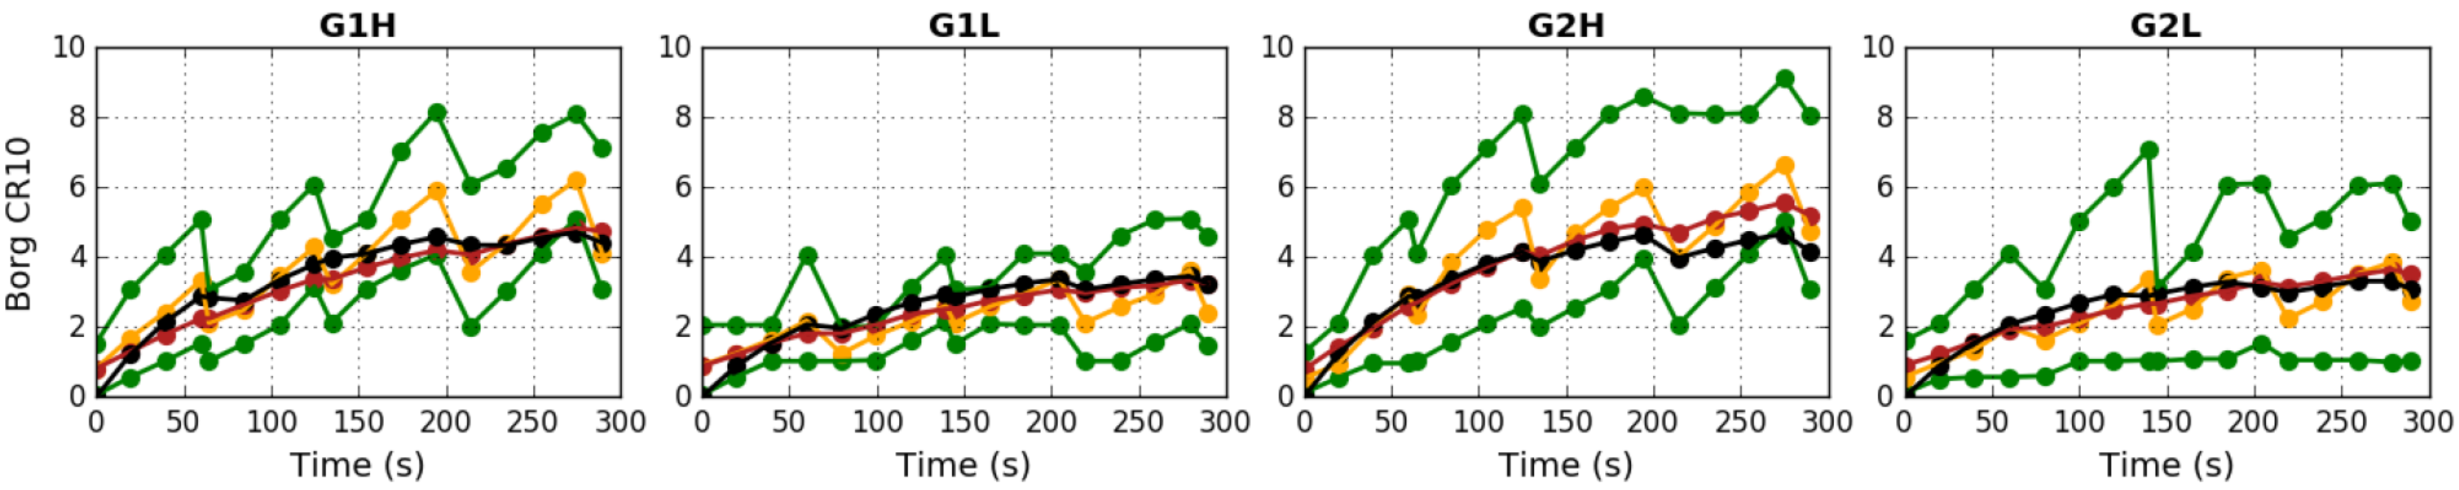
G1H
G1L
G2H
G2L
Borg CR10
Time (s)
0
2
4
6
8
10
0
50
100
150
200
250
300

Defining the target experience in the game Journey

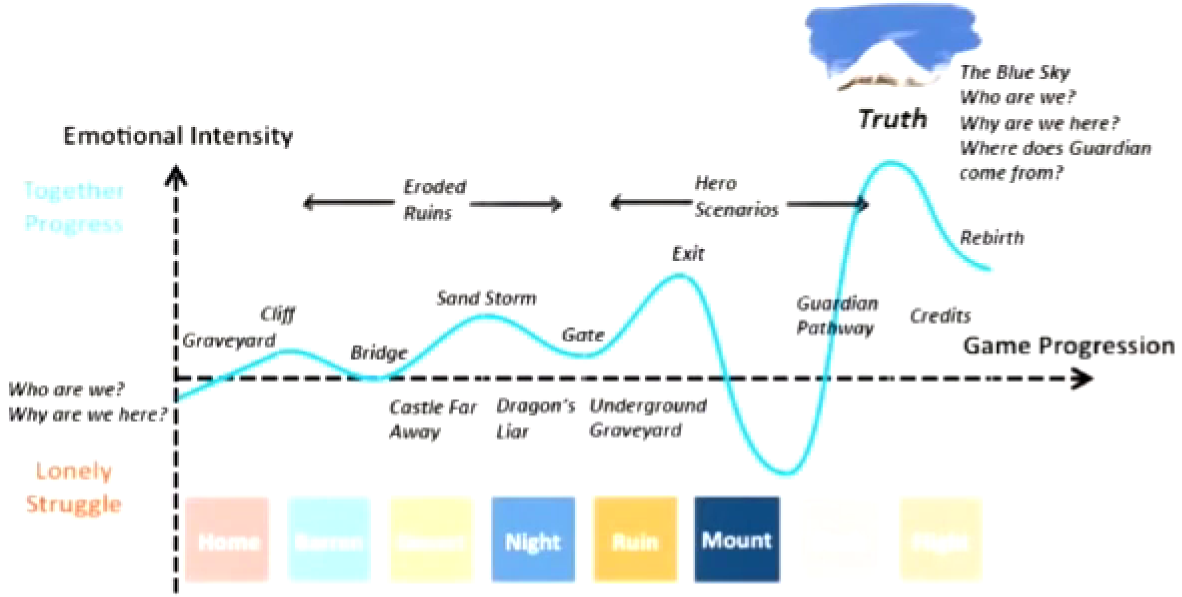

How to evaluate the objective function?

- Direct evaluation (theory-driven, data driven)
- Simulation-based evaluation (AI agents play the game)
- Interactive evaluation (human-in-the-loop, implicit or explicit)

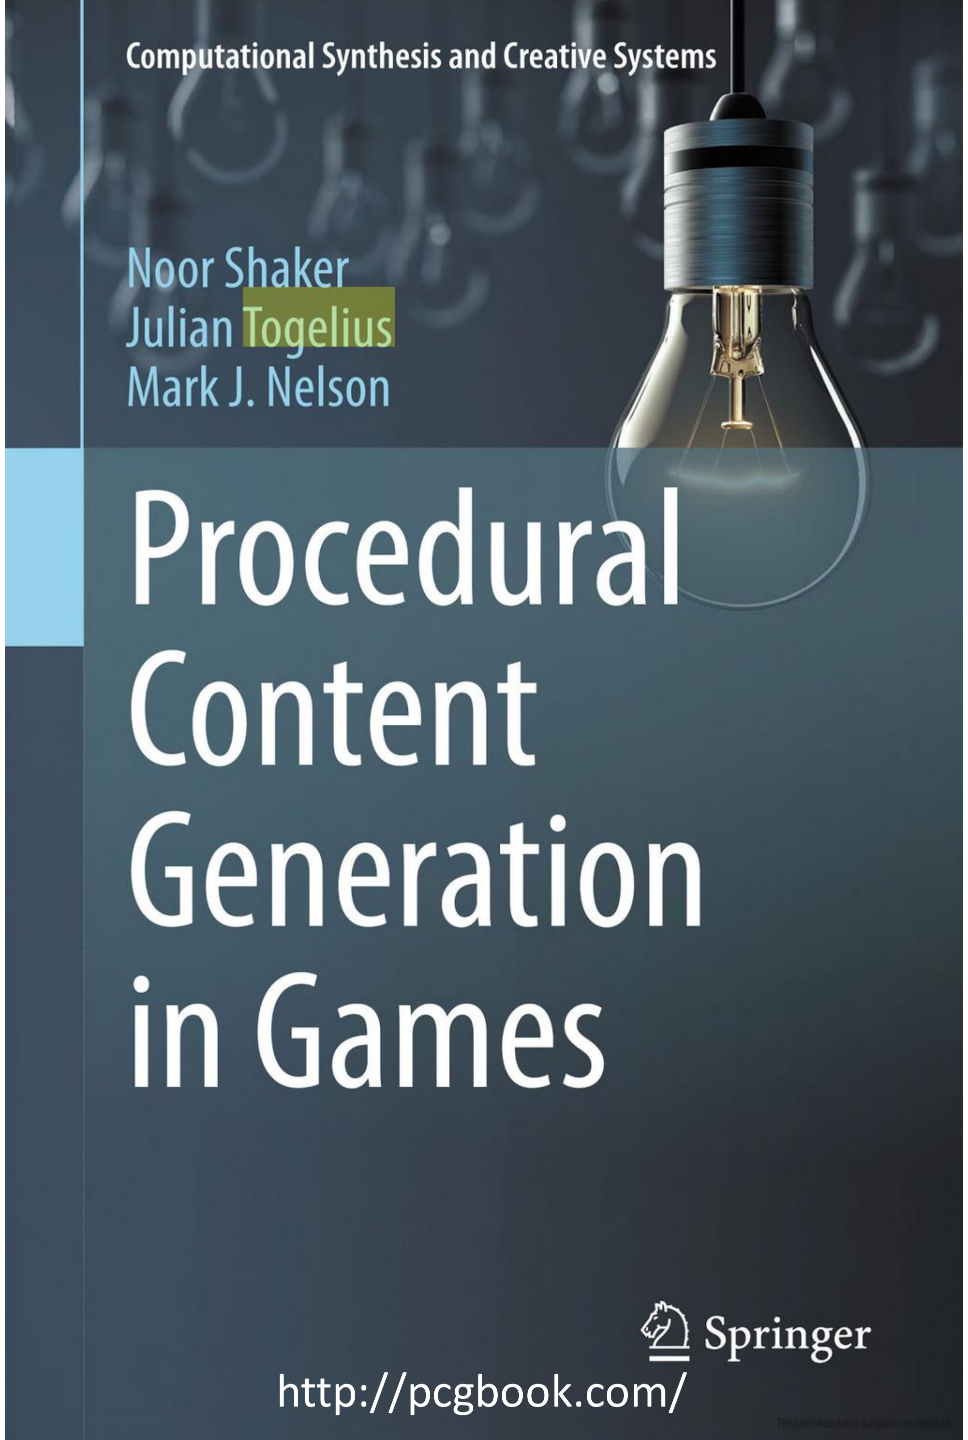

http://pcgbook.com/

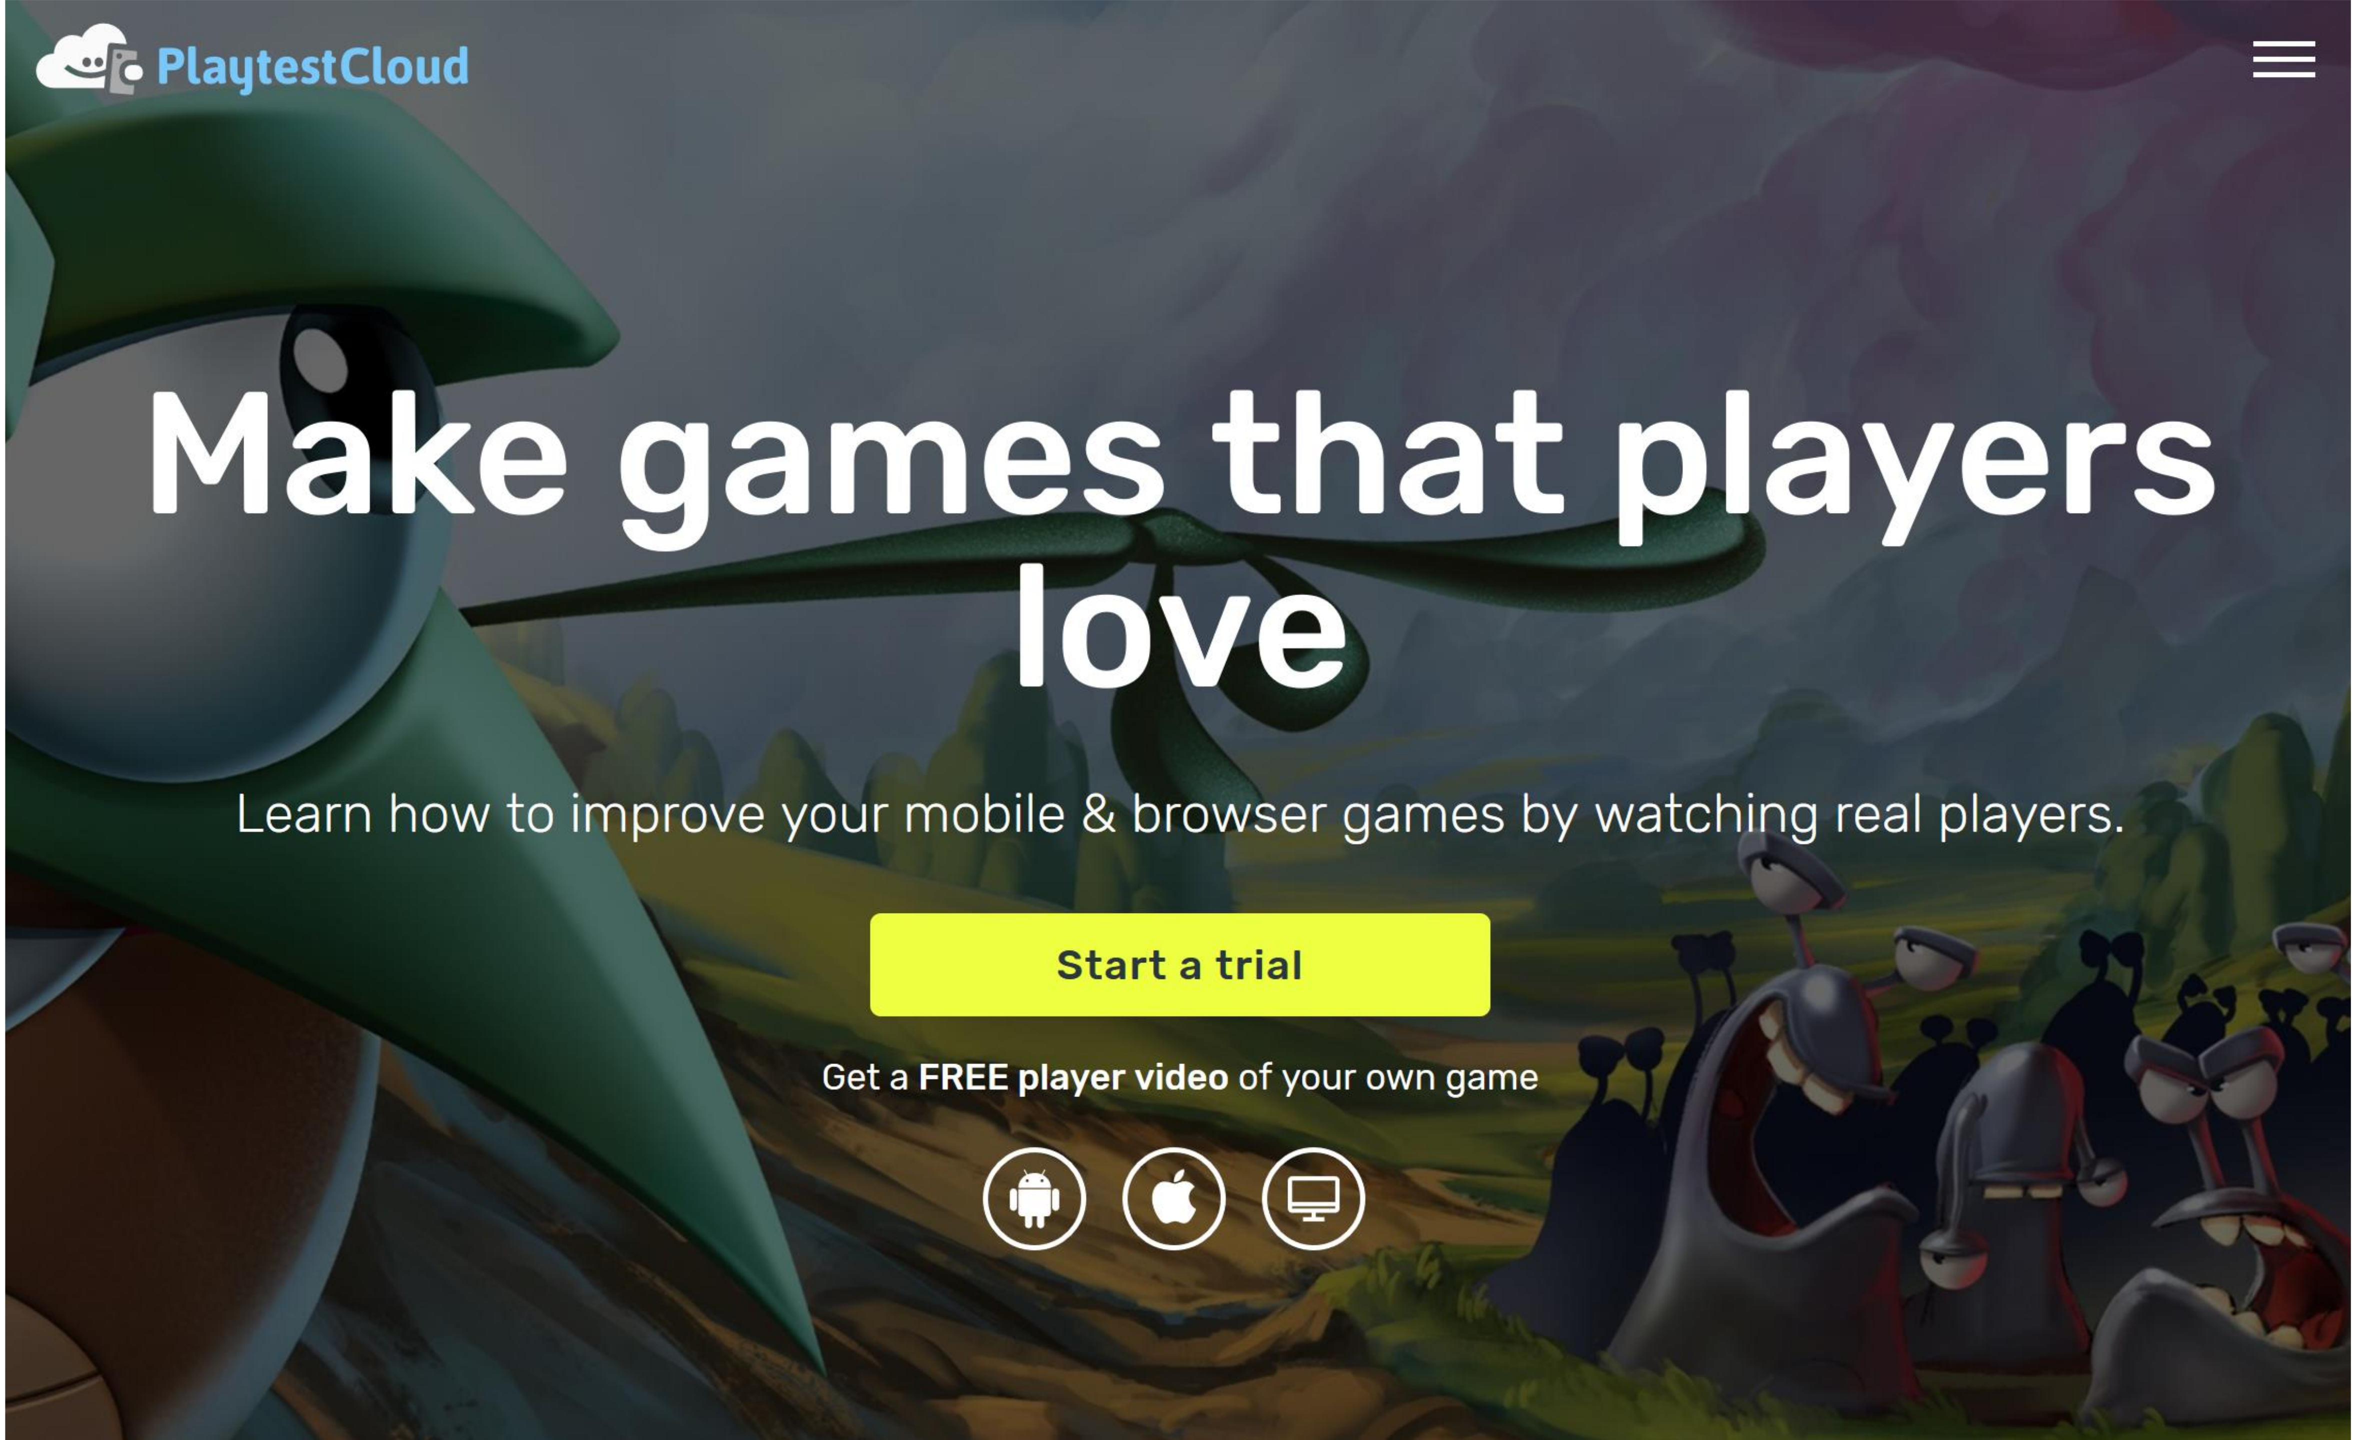
PlaytestCloud
Make games that players love
Learn how to improve your mobile & browser games by watching real players.
Start a trial
Get a FREE player video of your own game

PlaytestCloud

Products

Surveys

starting at **$9** /response

Send a survey to mobile gamers
from any target audience.

Playtesting

starting at **$49** /player

First-time player experience testing.
Players play 5 minutes or longer.

Multi-Session

starting at **$900** /playtest

Optimize the player experience
in the first 2-5 gameplay sessions.

Longitudinal

starting at **$1,020** /study

Test the first 3-10 days of gameplay.
Players play one or more sessions per day.