

CR-Eyes: A Computational Rational Model of Visual Sampling Behavior in Atari Games

Martin Lorenz

lorenz@cs.uni-leipzig.de
ScaDS.AI, Leipzig University
Leipzig, Germany

Niko Konzack

be66akax@studserv.uni-leipzig.de
Leipzig University
Leipzig, Germany

Alexander Lingler

alexander.lingler@it-u.at
Interdisciplinary Transformation
University
Linz, Austria

Philipp Wintersberger

philipp.wintersberger@it-u.at
Interdisciplinary Transformation
University
Linz, Austria

Patrick Ebel

ebel@uni-leipzig.de
ScaDS.AI, Leipzig University
Leipzig, Germany

Abstract

Designing mobile and interactive technologies requires understanding how users sample dynamic environments to acquire information and make decisions under time pressure. However, existing computational user models either rely on hand-crafted task representations or are limited to static or non-interactive visual inputs, restricting their applicability to realistic, pixel-based environments. We present CR-Eyes, a computationally rational model that simulates visual sampling and gameplay behavior in Atari games. Trained via reinforcement learning, CR-Eyes operates under perceptual and cognitive constraints and jointly learns where to look and how to act in a time-sensitive setting. By explicitly closing the perception–action loop, the model treats eye movements as goal-directed actions rather than as isolated saliency predictions. Our evaluation shows strong alignment with human data in task performance and aggregate saliency patterns, while also revealing systematic differences in scanpaths. CR-Eyes is a step toward scalable, theory-grounded user models that support design and evaluation of interactive systems.

CCS Concepts

• **Human-centered computing** → **User models**; • **Computing methodologies** → **Simulation types and techniques**.

Keywords

Visual Sampling, Atari, Reinforcement Learning, Computational Modeling

ACM Reference Format:

Martin Lorenz, Niko Konzack, Alexander Lingler, Philipp Wintersberger, and Patrick Ebel. 2026. CR-Eyes: A Computational Rational Model of Visual Sampling Behavior in Atari Games. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3772363.3798498>



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI EA '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2281-3/2026/04
<https://doi.org/10.1145/3772363.3798498>

1 Introduction

Understanding human perception during visually complex, dynamic tasks is essential for designing technologies that people can use efficiently in everyday life. Whether glancing at a smartwatch while navigating a crowded street or making rapid decisions in fast-paced video games, users must continuously determine where to direct attention and how to act accordingly. These decisions are made under perceptual and cognitive constraints with the goal to maximize task performance or expected utility [9]. This fundamental need for efficient information acquisition under time-pressure links gaming to high-stakes real-world domains, despite differing costs of failure. The ability to predict scanpaths can inform the design of intelligent user interfaces, adaptive notification systems, and assistive technologies [8, 16]. Models that close the perception-action loop by treating gaze as a goal-directed action that determines what users perceive next, and how this affects their behavior can serve as “simulated users” complementing the design and evaluation process of interactive systems while reducing reliance on costly experimental testing [22].

Still, before models of visual sampling are able to exploit their capabilities to support research and design, four key challenges need to be addressed. They must be able to (1) Predict full eye movement patterns, enabling the analysis of gaze behavior; (2) Operate directly on raw pixel input, removing the need for hand-crafted task representations; (3) Produce accurate predictions in fast-changing, dynamic environments where attentional demands shift rapidly; (4) Be grounded in theoretical models of perception and cognition, ensuring interpretability and supporting the modeling of individual differences among users. To date, most existing models fall short on at least one of these key challenges. Many focus primarily on visual attention allocation in static settings, such as data visualizations [26], single-screen app layouts [4, 31], web pages [5], or natural scenes [3, 32]. Others rely on task-specific abstractions to manage the complexity of dynamic environments [33], limiting their generalizability. Moreover, many of these models are black-box systems based on deep neural networks [17], which lack interpretability and offer limited insight into the cognitive mechanisms that drive human behavior. While they may predict where users are likely to look, they fail to explain why those decisions are made or how behavior might adapt to individual differences or

changes in task demands. As a result, such models are of limited usage to answering “what if” questions that drive the design of interactive technology.

To fill this gap, we present CR-Eyes, a model based on the theory of Computational Rationality (CR) [9, 22] that simulates human sampling behavior in fast-paced, visually complex Atari gameplay. In contrast to related approaches [33] CR-Eyes works directly on raw pixel input and does not require hand-crafted task environments. We evaluate CR-Eyes on three Atari games using the Atari-HEAD dataset [34].

Our results show that the model learns to sample the environment in a way that achieves game scores comparable to non-expert humans. While saliency maps derived from the agent’s behavior closely align with data generated by humans, the underlying eye-movement patterns still deviate substantially from human scanpaths.

2 Related Work

Traditionally, cognitive architectures such as EPIC [14] and ACT-R [24] have been used to model human attention allocation. While successful in HCI tasks such as goal-directed search [11, 13] and menu selection [2], these approaches rely on handcrafted environments and production rules, rendering them impractical for real-world applications.

A different line of work comes from robotics-inspired active vision systems, which learn visual sampling strategies to optimize task performance using sensory–motor reinforcement learning [29]. For example, Shang and Ryoo [25] propose an active vision agent for Atari games in the Arcade Learning Environment (ALE) [1] with separate sensory and motor action heads. However, such systems do not incorporate human perceptual or cognitive constraints and therefore do not produce human-like gaze behavior or strategies.

One approach for modeling human-like behavior, combining recent developments in reinforcement learning and cognitive theory, is Computational Rationality (CR) [22]. CR frames human behavior as rational optimization under constraints imposed by the task design, the user’s perceptual, motor, and cognitive capabilities, and competing goals such as speed and accuracy. By explicitly modeling such bounds, these models are more interpretable than deep neural networks, because they learn an optimization under explicitly defined, human-like constraints. Furthermore, CR models allow to answer “what if” questions through adjustment of these constraints. Recent examples of CR-based models that simulate human vision are Chartist that models top-down eye movements over static scenes [26], or Typoist that captures hand–eye coordination during touchscreen typing [27]. While Chartist [26] relies on a Large Language Model for high-level information acquisition, memory, and planning, to inform oculomotor control via RL on a pixel basis, Typoist [27] on the other hand, relies on screen-based raw pixel input for the RL agent alone. These models assume gaze-contingent perception in largely static or user-controlled environments. Yoon et al. [33] extend this paradigm to a dynamic first-person shooter by modeling saccadic latency and peripheral uncertainty, but rely on structured state representations rather than raw pixels. In contrast, CR-Eyes operates directly on pixel input from a continuous video game and accounts for environment changing actions.

To evaluate our approach we employ Atari-HEAD [34], a large-scale dataset that captures over 117 hours of gameplay across 20 Atari 2600 games. In contrast to regular Atari games, the games in Atari-HEAD run at a lower frame rate of 20 Hz and could be completely paused, resulting in near-optimal information acquisition behavior and thus game performance, which we mirror to match the human action space and timing conditions during data collection.

3 CR-Eyes

The goal of CR-Eyes is to reproduce human eye movements and gameplay behavior in Atari games while directly interacting with the ALE [1] based on raw pixel inputs. To achieve these goals, we make the following design choices: (1) We leverage the theory of Computational Rationality (CR), modeling humans’ internal environment (e.g. perception, cognition, motor control) and their external environment (e.g., game design); (2) In line with evidence from cognitive science [7] and related work [27, 28], we design distinct components to model perception, motor control and memory; (3) Inspired by findings on human sensorimotor learning [30], we use Sensorimotor Understanding Guided Active Reinforcement Learning (SUGARL) as proposed by Shang and Ryoo [25] to learn joint motor and sensory policies. The whole architecture is shown in Fig. 1.

3.1 Problem Formulation

Following related work on CR and visual sampling [26, 27], we formulate human visual sampling and gameplay as a sequential decision-making problem under uncertainty via a constrained, computationally rational Partially Observable Markov Decision Process (POMDP) (S, A, O, r) :

- S is the state space, where each state s_t is defined by the pixel representation of the environment at time t .
- A is the action space, in which the agent can execute a motor and a sensory action simultaneously. *Motor actions* describe game actions, such as *up*, *down*, *left*, *right*, or *fire* or to pause the game at the current frame. *Sensory actions* move the fovea, meaning they determine which portions of the screen to focus on.
- O denotes the observation space, in which o_t is represented by a stacked pixel representation of the last n observed image patches.
- r is the reward provided after every time step t . We formulate it as the sum of the achieved game score in a specific time step, e.g., the points dropped by the game for collecting objects and negative penalties for pausing and larger eye movements.

3.2 External Environment

The external environment is a semi-frame-by-frame environment adopted from the ALE [1] framework. Accordingly, it adds the possibility to pause the game at each step and thereby mirrors the setup as used by Zhang et al. [34], in particular a frame rate of 20 Hz. The agent takes two actions every timestep, one sensory and one motor action. To determine the time passed between two steps,

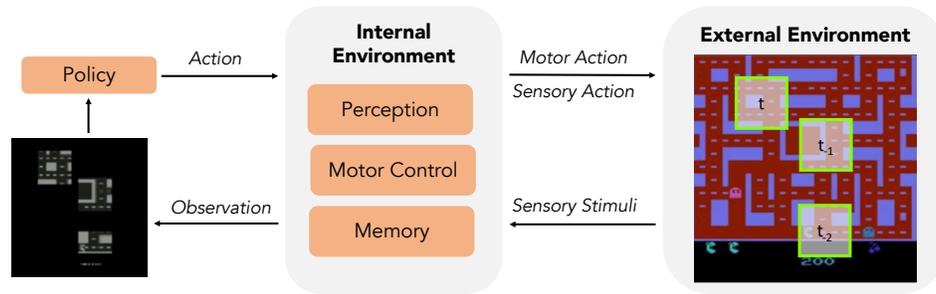


Figure 1: An overview over the CR-Eyes Architecture. Through the internal environment the agent performs a motor action (game action) and a sensory action (position to look at) at every timestep t in the external environment. The resulting observed patch is returned to the internal environment and integrated into the memory, consisting of the last n stacked observations. The memory is given to the RL agent as observation on which basis the agent will output the next motor and sensory action.

we use the EMMA model which allows us to calculate the time it took for the agent to fixate two subsequent patches of the screen.

3.3 Internal Environment

Vision is the main interface between the agent and the game. Human vision is limited by several factors such as muscular bounds that result in longer saccade times for larger eye movements, or fixation and encoding times of targets. To model realistic saccade times, we employ the EMMA model by Salvucci [24] with default parameters and replicating the setup by Zhang et al. [34]. The internal environment models foveated vision, so that the agent only perceives a fraction of the screen displaying the Atari game, similar to Shang and Ryoo [25], where the screen is divided into a 5×5 grid. These 25 discrete positions define the sensory actions available to the agent. Furthermore, we stack the last $n = 4$ observed patches to model working memory that serves as the agent’s observation.

3.4 Neural Network Architecture and Training

Based on the work of Shang and Ryoo [25] we utilize a similar architecture, namely a DQN with two separate action heads. Following best practices in Atari RL we re-scale the gameplay frames to 84×84 pixels, use sticky actions with $\xi = 0.25$, and frame skipping $k = 4$ [19]. To prevent indefinite pausing, we employ *invalid action masking* [12] and introduce a specific pause penalty. This penalty, along with saccade costs, is optimized via hyperparameter grid search to match human data.

4 Evaluation and Results

To evaluate CR-Eyes, we compare gameplay performance and gaze behavior against human benchmarks and datasets. We evaluate CR-Eyes on three Atari games, namely *Asterix*, *H.E.R.O.*, and *Seaquest*. For these games, eye-tracking data [34] as well as non-expert human scores are available [1].

4.1 Gameplay Behavior

We report game scores in both the standard setting [1] and the pausing enabled Atari-HEAD setting [34] (15 minutes scores), testing whether CR-Eyes can leverage pausing to acquire additional information and improve action selection, similar to humans. While

the goal of CR-Eyes is to achieve human-like scores, we also compare our results against classical Reinforcement Learning (RL) approaches such as DQN [21] and the world model approach Dreamer [10]. Table 1 shows that CR-Eyes achieves game scores similar to non-expert humans in *Asterix* and *H.E.R.O.* for the regular game setting. However, for *Seaquest* CR-Eyes’ scores are lower. While CR-Eyes achieves better game scores in all three games when it’s given the possibility to pause the gameplay, the improvement is significantly lower than that observed in human players. The scores reported in the Atari-HEAD dataset [34] are several orders of magnitude higher than the best CR-Eyes scores. The DQN with full state knowledge also achieves scores levels of magnitude lower [21] than those reported in Atari-HEAD [34]. In contrast Hafner et al. [10] report that their approach *Dreamer* achieves scores much higher than those reported in Atari-HEAD.

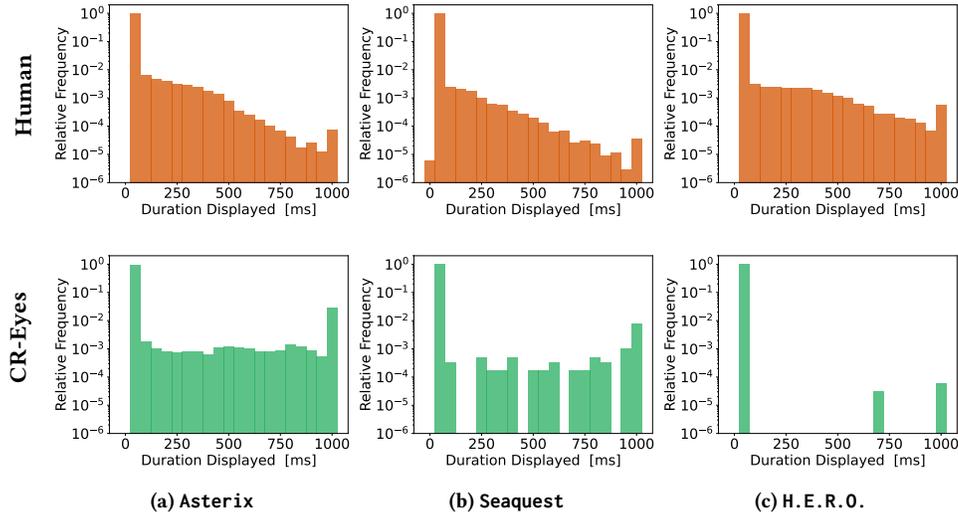
The *pause frequency* is inferred by analyzing the time a frame was displayed (Figure 2). The human data (Figure 2 (top)) was recorded by Zhang et al. [34] and the durations for CR-Eyes (Fig. 2 (bottom)) are calculated using the EMMA model [24]. Notably, the game remains unpaused for the majority of the time for both human and CR-Eyes (50ms display duration). Analyzing the deviations shows that humans prefer shorter pauses. In contrast, the agent has a roughly equal distribution of pause lengths. In general, the agent replicates the human trend to only pause the game seldom. However, for H.E.R.O. the agent learned to not pause at all. Qualitative evaluations of gameplay videos show that, while humans mainly pause in complex situations, the agent primarily pauses shortly before losing the game.

4.2 Gaze Behavior

We compare CR-Eyes saliency predictions to those of human players in the Atari-HEAD dataset [34]. Due to the stochastic environment and the environment changing actions, the game roll-outs of CR-Eyes do not correspond to the ones in the dataset. Accordingly, we cannot directly compare the saliency maps generated by CR-Eyes to those available in the dataset. To solve this issue, we train a supervised saliency predictor using the Atari-HEAD data to generate synthetic saliency maps for the roll outs generated by CR-Eyes. Saliency prediction is a common technique used in RL and imitation learning to evaluate an agent’s visual attention allocation against

Table 1: Game scores for different human (orange) and RL (blue) baselines and our agent (green).

Game	non-expert Human [1]	Atari-HEAD [34]	DQN [21]	Dreamer [10]	CR-Eyes (no pause)	CR-Eyes (pausing)
Asterix	620	135,000	6012	441,763	495	619
Seaquest	156	64,710	5286	356,584	20	68
H.E.R.O.	6087	77,185	19,950	40,677	5112	8649

**Figure 2: Durations how long frames were displayed in ms for humans (orange) [34] and our agent (green). The agent was trained for three million training steps. The long tail of the distribution is summed up in the last bin and the y-axis is logarithmically scaled.**

that of humans [15, 21, 34]. Saliency maps over the whole episode reveal that unlike humans who mostly focus at the center (Fig. 3a), the agent fixates on the upper-left spawn point (Fig. 3b), missing the next key Area of Interests (AOIs) relevant for game progress.

However, truncating human episodes to the mean length of agent episodes (Fig. 3c) reveals systematic differences in saliency. While the agent identifies the avatar as the relevant object on the screen, it fails to translate this knowledge into effective gameplay. This behavior can also be observed in the scanpaths as shown in Figure 4. In this scene the avatar is in the top left corner, but the agent focus remains in the middle right and empty part of the screen.

5 Discussion and Future Work

Our results show that CR-Eyes is a first step toward a computational model grounded in cognitive science that can simulate human gaze behavior by closing the perception-action loop in dynamic environments solely based on pixel streams. To evaluate if CR-Eyes simulates human behavior, we compare it to human data along four dimensions: (1) game scores, (2) saliency predictions, (3) scanpaths, and (4) pausing behavior. CR-Eyes’s game performance is comparable to non-expert humans in a regular setting without pausing. Like humans, CR-Eyes leverages the possibility to pause the gameplay to reach higher game scores. However, this improvement is magnitudes smaller compared to that of human players. Furthermore, our approach does not reflect human preferences for short pauses. We

further compared saliency predictions against human data from Atari-HEAD [34]. While CR-Eyes aligns closely to human saliency in early game stages it slightly diverts in later stages. Simulated scanpaths however do not match human data. The agent does not focus on the avatar and regularly misses information considered relevant by humans such as approaching enemies.

From a methodological perspective, these results suggest that saliency maps alone are an insufficient metric for evaluating visual sampling behavior as saliency maps can mask fundamentally different sampling behavior. From a behavioral perspective, the agent’s scanpaths are primarily optimized for reward maximization, whereas human visual behavior reflects an understanding of task dynamics, leading to anticipatory glances based on expectations rather than immediate information gain. This indicates that the agent does not acquire game dynamics in a human-like manner. Differences in pausing behavior further suggest that the agent fails to integrate multiple observations to better infer and exploit environment dynamics. This limitation likely stems from the model-free nature of DQN. Even in standard fully observable Atari environments, DQN scores magnitudes lower than humans in Atari-HEAD. To bridge this gap, model-based algorithms like Dreamer [10] appear promising, as they learn world dynamics to facilitate strategic planning. Another reason for the observed differences in visual sampling behavior may be that the agent relies on foveated vision only, whereas humans use peripheral vision as

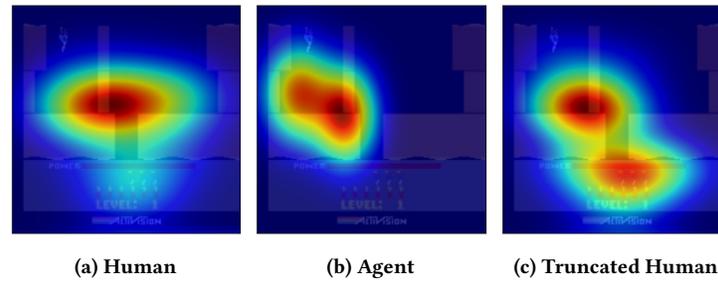


Figure 3: Agent and human saliency throughout an entire *Seaquest* episode. Human data is taken from Atari-HEAD. Subfigure c) shows saliency for the human episode truncated to the length of the agent’s episode.

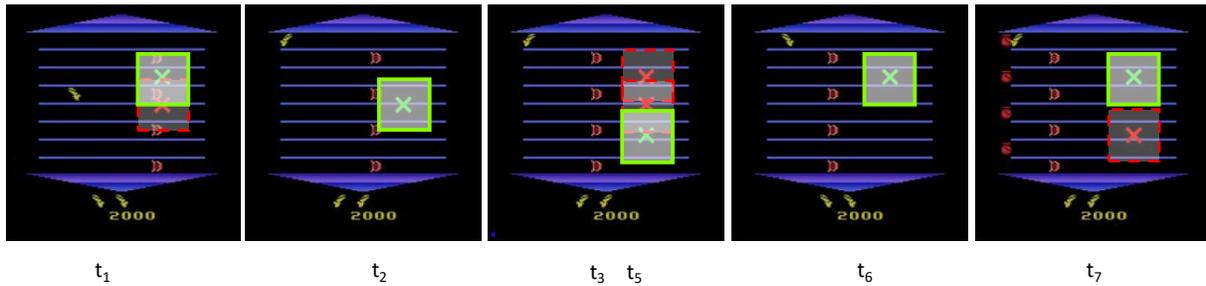


Figure 4: Resulting scanpath for *Asterix*. The green squares are the fovea position at step t_i . The red squares represent the agent’s sensory action, that will be executed in the next step. At t_3 the game is paused until t_5 .

well. Research shows that this affects performance and may result in different gaze patterns and scanpaths [6].

We propose the following steps for future work: (1) use a model-based algorithm like Dreamer [10] to allow the agent to generate an internal representation of game dynamics, (2) perform parameter inference with human behavior as cost function, and (3) employ models of peripheral vision [23], bringing agent perception closer to human perception. While gaming behavior is achievement-oriented, it shares the core cognitive challenge of information acquisition in dynamic environments with high-stakes domains. We argue that the gap between entertainment and safety critical tasks is defined by the cost of failure. As CR models explicitly optimize against a utility function, this can be bridged by adjusting the penalty for errors to reflect high-stakes consequences. Thus, our model is a first prototype and could be further applied to other domains such as automated testing of interface designs in driving [18], aviation, or video game design, complementing user studies. Additionally, in high-stake monitoring tasks like air traffic control [8], these models could serve as a safety mechanism by detecting irregular scanning behavior and preventing *Out-of-the-Loop* [20] effects.

Acknowledgments

The authors acknowledge the financial support by the Federal Ministry of Research, Technology and Space of Germany and by Sächsische Staatsministerium für Wissenschaft, Kultur und Tourismus in the programme Center of Excellence for AI-research „Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig“, project identification number: ScaDS.AI.

References

- [1] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. 2013. The Arcade Learning Environment: An Evaluation Platform for General Agents. *Journal of Artificial Intelligence Research* 47 (June 2013), 253–279. <https://doi.org/10.1613/jair.3912>
- [2] Michael D. Byrne. 2001. ACT-R/PM and Menu Selection: Applying a Cognitive Architecture to HCI. *International Journal of Human-Computer Studies* 55, 1 (July 2001), 41–84. <https://doi.org/10.1006/ijhc.2001.0469>
- [3] Antoine Coutrot, Janet H. Hsiao, and Antoni B. Chan. 2018. Scanpath modeling and classification with hidden Markov models. *Behavior Research Methods* 50, 1 (Feb. 2018), 362–379. <https://doi.org/10.3758/s13428-017-0876-8>
- [4] Arturo Deza, Jeffrey R. Peters, Grant S. Taylor, Amit Surana, and Miguel P. Eckstein. 2017. Attention Allocation Aid for Visual Search. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, Denver Colorado USA, 220–231. <https://doi.org/10.1145/3025453.3025834>
- [5] Gautier Drusch, J.M. Christian Bastien, and Stéphane Paris. 2020. Analysing Eye-Tracking Data: From Scanpaths and Heatmaps to the Dynamic Visualisation of Areas of Interest. In *International Conference on Applied Human Factors and Ergonomics (AHFE)*. <https://doi.org/10.54941/ahfe100394>
- [6] Yke Bauke Eisma, Ahmed Bakay, and Joost De Winter. 2024. Expectancy or Saliency?—Replicating Senders’ Dial-Monitoring Experiments With a Gaze-Contingent Window. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 66, 6 (June 2024), 1770–1785. <https://doi.org/10.1177/00187208231176148>
- [7] Jerry A Fodor. 1983. *The modularity of mind*. MIT press.
- [8] Anton Gasse, Alexander Lingler, Martin Lorenz, Antti Oulasvirta, Philipp Wintersberger, and Patrick Ebel. 2025. Evaluating Attention Management Systems for Dynamic Monitoring Tasks. In *Adjunct Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work*. ACM, Amsterdam Netherlands, 1–6. <https://doi.org/10.1145/3707640.3731920>
- [9] Samuel J. Gershman, Eric J. Horvitz, and Joshua B. Tenenbaum. 2015. Computational Rationality: A Converging Paradigm for Intelligence in Brains, Minds, and Machines. *Science* 349, 6245 (July 2015), 273–278. <https://doi.org/10.1126/science.aac6076>
- [10] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. 2023. Mastering Diverse Domains through World Models. <https://doi.org/10.48550/ARXIV.2301.04104>
- [11] Tim Halverson and Anthony J. Hornof. 2011. A Computational Model of “Active Vision” for Visual Search in Human–Computer Interaction. *Human–Computer*

- Interaction* 26, 4 (Dec. 2011), 285–314. <https://doi.org/10.1080/07370024.2011.625237>
- [12] Shengyi Huang and Santiago Ontañón. 2022. A Closer Look at Invalid Action Masking in Policy Gradient Algorithms. *The International FLAIRS Conference Proceedings* 35 (May 2022). <https://doi.org/10.32473/flairs.v35i.130584> arXiv:2006.14171 [cs]
- [13] David E. Kieras. 2019. Visual Search Without Selective Attention: A Cognitive Architecture Account. *Topics in Cognitive Science* 11, 1 (Jan. 2019), 222–239. <https://doi.org/10.1111/tops.12406>
- [14] Davis E. Kieras and Davis E. Meyer. 1997. An Overview of the EPIC Architecture for Cognition and Performance With Application to Human-Computer Interaction. *Human-Computer Interaction* 12, 4 (Dec. 1997), 391–438. https://doi.org/10.1207/s15327051hci1204_4
- [15] Olivier Le Meur and Thierry Baccino. 2013. Methods for comparing scanpaths and saliency maps: strengths and weaknesses. *Behavior Research Methods* 45, 1 (March 2013), 251–266. <https://doi.org/10.3758/s13428-012-0226-9>
- [16] Alexander Lingler, Dinara Talypova, Jussi P. P. Jokinen, Antti Oulasvirta, and Philipp Wintersberger. 2024. Supporting Task Switching with Reinforcement Learning. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 82, 18 pages. <https://doi.org/10.1145/3613904.3642063>
- [17] Angela Lopez-Cardona, Parvin Emami, Sebastian Idesis, Saravanakumar Duraisamy, Luis A. Leiva, and Ioannis Arapakis. 2025. A Comparative Study of Scanpath Models in Graph-Based Visualization. In *Proceedings of the 2025 Symposium on Eye Tracking Research and Applications*. ACM, Tokyo Japan, 1–11. <https://doi.org/10.1145/3715669.3725882>
- [18] Martin Lorenz, Tiago Amorim, Debargha Dey, Mersedeh Sadeghi, and Patrick Ebel. 2024. Computational Models for In-Vehicle User Interface Design: A Systematic Literature Review. In *Proceedings of the 16th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. 204–215.
- [19] Marlos C. Machado, Marc G. Bellemare, Erik Talvitie, Joel Veness, Matthew Hausknecht, and Michael Bowling. 2018. Revisiting the Arcade Learning Environment: Evaluation Protocols and Open Problems for General Agents. *Journal of Artificial Intelligence Research* 61 (March 2018), 523–562. <https://doi.org/10.1613/jair.5699>
- [20] Natasha Merat, Bobbie Seppelt, Tyron Louw, Johan Engström, John D. Lee, Emma Johansson, Charles A. Green, Satoshi Katzaki, Chris Monk, Makoto Itoh, Daniel McGehee, Takashi Sunda, Kiyozumi Unoura, Trent Victor, Anna Schieben, and Andreas Keinath. 2019. The “Out-of-the-Loop” concept in automated driving: proposed definition, measures and implications. *Cognition, Technology & Work* 21, 1 (Feb. 2019), 87–98. <https://doi.org/10.1007/s10111-018-0525-8>
- [21] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmashan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. 2015. Human-Level Control through Deep Reinforcement Learning. *Nature* 518, 7540 (Feb. 2015), 529–533. <https://doi.org/10.1038/nature14236>
- [22] Antti Oulasvirta, Jussi P. P. Jokinen, and Andrew Howes. 2022. Computational Rationality as a Theory of Interaction, In CHI Conference on Human Factors in Computing Systems. *CHI Conference on Human Factors in Computing Systems*, 1–14. <https://doi.org/10.1145/3491102.3517739>
- [23] Ruth Rosenholtz. 2016. Capabilities and Limitations of Peripheral Vision. *Annual Review of Vision Science* 2, 1 (Oct. 2016), 437–457. <https://doi.org/10.1146/annurev-vision-082114-035733>
- [24] Dario D Salucci. 2001. An Integrated Model of Eye Movements and Visual Encoding. *Cognitive Systems Research* 1, 4 (Feb. 2001), 201–220. [https://doi.org/10.1016/S1389-0417\(00\)00015-2](https://doi.org/10.1016/S1389-0417(00)00015-2)
- [25] Jinghuan Shang and Michael S. Ryoo. 2023. Active Vision Reinforcement Learning under Limited Visual Observability. <http://arxiv.org/abs/2306.00975> arXiv:2306.00975 [cs].
- [26] Danqing Shi, Yao Wang, Yunpeng Bai, Andreas Bulling, and Antti Oulasvirta. 2025. Chartist: Task-driven Eye Movement Control for Chart Reading. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–14. <https://doi.org/10.1145/3706598.3713128>
- [27] Danqing Shi, Yujun Zhu, Francisco Erivaldo Fernandes Junior, Shumin Zhai, and Antti Oulasvirta. 2025. Simulating Errors in Touchscreen Typing. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1086, 13 pages. <https://doi.org/10.1145/3706598.3713153>
- [28] Danqing Shi, Yujun Zhu, Jussi P. P. Jokinen, Aditya Acharya, Aini Putkunen, Shumin Zhai, and Antti Oulasvirta. 2024. CRTypist: Simulating Touchscreen Typing Behavior via Computational Rationality. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–17. <https://doi.org/10.1145/3613904.3642918>
- [29] Steven D. Whitehead and Dana H. Ballard. 1990. Active Perception and Reinforcement Learning. In *Machine Learning Proceedings 1990*. Elsevier, 179–188. <https://doi.org/10.1016/B978-1-55860-141-3.50025-0>
- [30] Daniel M Wolpert, Jörn Diedrichsen, and J Randall Flanagan. 2011. Principles of sensorimotor learning. *Nature reviews neuroscience* 12, 12 (2011), 739–751.
- [31] Pingmei Xu, Yusuke Sugano, and Andreas Bulling. 2016. Spatio-Temporal Modeling and Prediction of Visual Attention in Graphical User Interfaces. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 3299–3310. <https://doi.org/10.1145/2858036.2858479>
- [32] Zhibo Yang, Lihan Huang, Yupei Chen, Zijun Wei, Seoyoung Ahn, Gregory Zelinsky, Dimitris Samaras, and Minh Hoai. 2020. Predicting Goal-Directed Human Attention Using Inverse Reinforcement Learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Seattle, WA, USA, 190–199. <https://doi.org/10.1109/CVPR42600.2020.00027>
- [33] June-Seop Yoon, Hee-Seung Moon, Ben Boudaoud, Josef Spjut, Iuri Frosio, Byungjoo Lee, and Joohwan Kim. 2025. Modeling visually-guided aim-and-shoot behavior in first-person shooters. *International Journal of Human-Computer Studies* 199 (May 2025), 103503. <https://doi.org/10.1016/j.ijhcs.2025.103503>
- [34] Ruohan Zhang, Calen Walshe, Zhuode Liu, Lin Guan, Karl S. Muller, Jake A. Whritner, Luxin Zhang, Mary M. Hayhoe, and Dana H. Ballard. 2019. ATARI-HEAD: Atari Human Eye-Tracking and Demonstration Dataset. <https://doi.org/10.48550/arXiv.1903.06754> arXiv:1903.06754 [cs, stat]