CogEnv: A Reinforcement Learning Environment for Cognitive Tests

Morteza Ansarinia (morteza.ansarinia@uni.lu)

Brice Clocher (brice.clocher@uni.lu)

Aurélien Defossez (aurelien.defossez@uni.lu)

Emmanuel Schmück (emmanuel.schmuck@uni.lu)

Pedro Cardoso-Leite (pedro.cardosoleite@uni.lu)

Department of Behavioural and Cognitive Sciences, University of Luxembourg
4366 Belval, Luxembourg

Abstract

Understanding human cognition involves developing computational models that mimic and possibly explain behavior; these are models that "act" like humans and produce similar outputs when facing the same inputs. To facilitate the development of such models and ultimately further our understanding of the human mind we created CogEnv: a reinforcement learning environment where artificial agents interact with and learn to perform cognitive tests and can then be directly compared to humans.

By leveraging CogEnv, cognitive and AI scientists can join efforts to better understand human cognition: the relative performance profiles of human and artificial agents may provide new insights on the computational basis of human cognition and on what human-like abilities artificial agents may lack.

Keywords: Cognitive Tests; Reinforcement Learning; Human versus Artificial Agents

Understanding the computations underlying human cognition is vital for scientific progress. Most efforts in cognitive sciences to understand how people perform cognitive tests focus on models that describe the data (e.g., factor analysis). There are only a few models describing the mechanisms underlying the performance of a task (i.e., models that "act" like humans and produce responses) and even fewer models that can account for performance across many tasks.

One productive strategy has been to develop cognitive architectures (e.g., ACT-R; Anderson et al., 2004). Alternatively, recent developments in AI allow the application of flexible, generic architectures to solve a wide variety of problems. Their ability to do (or not do) so may reveal computational constraints underlying specific tasks. Reinforcement Learning (RL) seems particularly well suited to model performance in cognitive tests as they typically involve the presentation of a stream of stimuli and the execution of a discrete set of actions followed by a reward signal that may drive learning.

Despite the relevance and potential of RL to model cognition, there is currently no easy way to train RL models on the same cognitive tests that are used to assess humans.
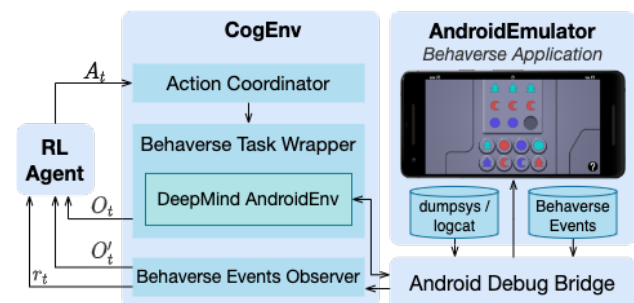

Figure 1: Overall architecture of the CogEnv. CogEnv communicates with AndroidEnv via Protocol Buffer messages and manages access to the Behaverse events. O_t is the screenshot of the task at time t, O'_t is the extra observations extracted from the Behaverse events including information about the task and stimuli, r_t is the reward, and A_t is the agent's action.

Here we present CogEnv, a configurable multi-task environment for RL agents to emulate cognitive tests. Under the hood, CogEnv uses DeepMind's AndroidEnv (Toyama et al., 2021) to run the Behaverse cognitive assessment battery (see `behaverse.org`). Behaverse tasks are customizable at many levels, allowing the construction of a large number of randomized trials for training RL agents. In the following sections, we present the technical details of the CogEnv and its capability to run RL agents on cognitive tests.

Technical Specifications

We simulate a real-time RL environment, where the environment, upon receiving an observation, invokes a callback method in the agent. We use AndroidEnv to run and manage the Behaverse cognitive assessment battery in a virtual Android device. A set of task-specific parsers then decodes screenshots, event streams, and system logs to extract numerical rewards and symbolic observations (see Figure 1). The reward and the observed state are then sent to the agent via the callback. CogEnv then waits for the agent to respond with an action, and issues a timeout if no response occurred within a duration specified by the cognitive test.
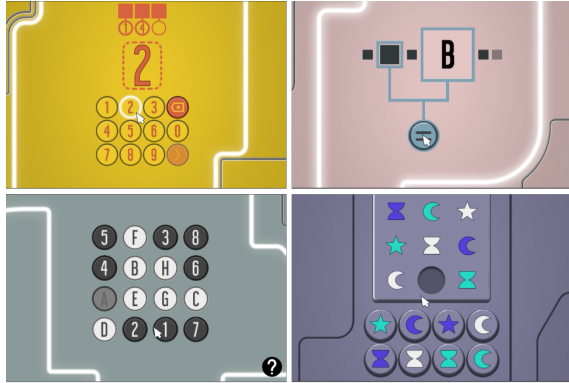
Figure 2: Screenshot (O_t) of four Behaverse tasks. A) Digit Span (working memory), B) N-Back (working memory), C) Trail Making Test (cognitive flexibility), and D) Belval Matrices (matrix reasoning). See `https://behaverse.org`.

Tasks CogEnv currently runs four Behaverse tasks (see Figure 2) selected to cover main components of cognitive control (Ansarinia, Schrater, & Cardoso-Leite, 2022). In the Belval Matrices test for example, agents are shown a matrix of symbols on a 3x3 grid, where one of the cells of the matrix has been removed, and they are tasked to identify the missing cell from a set of eight options (panel D of Figure 2). The Belval Matrices can randomly generate a large number of test items of varying difficulty and structure, which makes this test interesting for human and artificial learning studies.

Timing CogEnv supports both step-lock (where the environment pauses between two consecutive actions) and real-time mode (where the environment runs asynchronously from the agent). A real-time environment is necessary to study the timing of actions: human behavior is typically evaluated in terms of both accuracy and speed.

Action Space Each test defines a discrete action space that is in fact bounded tap gestures on the buttons of the graphical interface. The Action Coordinator component (see Figure 1) automatically constructs a sequence of AndroidEnv gestures (TAP, TOUCH, and LIFT) that together perform the requested action as a set of movements in the emulated device.

Observation Space CogEnv asynchronously invokes and waits for the agent to act. The invocation is accompanied by a screenshot of the Behaverse screen, as well as the reward value and symbolic representations of the task state extracted from the logs and event streams.

Comparing humans and artificial agents

CogEnv allows us to compare human and artificial agents on the exact same cognitive tests, generating for both the same type of data that can be analyzed using a common data analysis pipeline. Figure 3 illustrates how such comparisons may yield new insights. We collected data (accuracy and response time) from 200 human participants completing 20 items of the

Belval Matrices (see Figure 3) and are currently training a selection of discrete control agents on the same test (i.e., DQN and R2D2 from the Acme Tensorflow library; see Hoffman et al., 2020; Toyama et al., 2021): agents are trained on 1000 randomly generated items and tested on a set of 20 unseen test items, the exact same used with human participants.

Contrasting human and artificial agents may yield one of the following main scenarii: **(A)** The artificial agent mimics the human performance profile well, suggesting it captures something fundamental about human cognition and that its study may help us better understand humans. **(B)** The artificial agent performs the task well but displays a different performance profile than humans. This could suggest that there are in fact several ways of solving the task and that the human performance profile has a characteristic computational signature. **(C)** The artificial agent performs like humans on some items but very differently on others. This may indicate that humans use a mixture of cognitive strategies or that the artificial agent needs to be augmented to perform human-like.

Whatever the case may be, it is clear that the comparison of human versus artificial agents, as well as the comparison among artificial agents provides a unique source of information that significantly augments our ability to make sense of human behavior in cognitive tests.
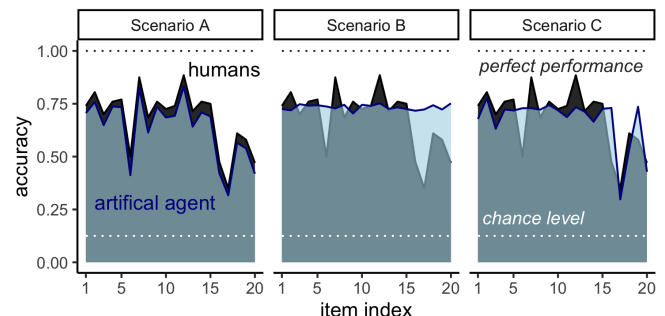

Figure 3: Hypothetical scenarios when comparing the performance of humans versus computational agents (see text).

Conclusion

Cognitive tests play a central role in the study of human cognition. We introduced CogEnv, a framework that runs cognitive tests within a virtual environment that enables training and evaluating artificial agents in a way that is directly comparable to human studies. CogEnv also provides a way to study cognitive tests and how learning to perform well in one cognitive test might transfer to others.

Environments like CogEnv have proven quite useful in other fields, e.g., AnimalAI 3 for animal cognition (Crosby et al., 2020) and RecSym for recommendation systems (Ie et al., 2019). We believe that CogEnv can complement other approaches (e.g., cognitive architectures) and hope it will yield new insights on human cognition and help coordinate efforts across disciplines to better understand the computational foundations of cognitive performance.

Acknowledgments

References

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological review*, *111*(4), 1036.

Ansarinia, M., Schrater, P., & Cardoso-Leite, P. (2022). Linking theories and methods in cognitive sciences via joint embedding of the scientific literature: The example of cognitive control. *arXiv preprint arXiv:2203.11016*.

Crosby, M., Beyret, B., Shanahan, M., Hernández-Orallo, J., Cheke, L., & Halina, M. (2020, 08–14 Dec). The animal-ai testbed and competition. In H. J. Escalante & R. Hadsell (Eds.), *Proceedings of the neurips 2019 competition and demonstration track* (Vol. 123, pp. 164–176). PMLR.

Hoffman, M., Shahriari, B., Aslanides, J., Barth-Maron, G., Behbahani, F., Norman, T., . . . others (2020). Acme: A research framework for distributed reinforcement learning. *arXiv preprint arXiv:2006.00979*.

Ie, E., Hsu, C.-w., Mladenov, M., Jain, V., Narvekar, S., Wang, J., . . . Boutilier, C. (2019). Recsim: A configurable simulation platform for recommender systems. *arXiv preprint arXiv:1909.04847*.

Toyama, D., Hamel, P., Gergely, A., Comanici, G., Glaese, A., Ahmed, Z., . . . Precup, D. (2021). Androidenv: A reinforcement learning platform for android. *arXiv preprint arXiv:2105.13231*.