

A VR-Based System and Architecture for Computational Modeling of Minds

Saty Raghavachary¹ and Lurong Lei²

¹ University of Southern California, Los Angeles CA 90089, USA

² Chongqing University of Arts and Science, Chongqing 402160 China

saty@usc.edu¹, 20030051@cqu.edu.cn²

Abstract. Computational modeling of natural cognition is a crucial step towards achieving the grand goal of human-level computational intelligence. Successful ideas from existing models, and possibly newer ones, could be assembled to create a unified computational framework (eg. the Standard Model of the Mind, which attempts to unify three leading cognitive architectures) - this would be of great use in AI, robotics, neuroscience and cognitive science. This short position paper proposes the following: a VR-based system provides the most expedient, scalable and visually verifiable way to implement, test and refine a cognitive mind model (which would always be embodied in a character in a virtual world). Such a setup is discussed in the paper, including advantages and drawbacks over alternative implementations.

Keywords: Virtual Reality, Embodiment, Computational Mind Modeling, Cognitive Architectures, Artificial General Intelligence, Artificial Intelligence.

1 Introduction

In this paper, we describe a system for creating an artificial general intelligence (AGI) agent - one that would be capable of carrying out cognitive tasks such as perception and concept formation. There has been continued interest in developing such an agent capable of human-level intelligence, but there has been limited success in achieving it. We hope that the system we are proposing, can help the effort along.

Our key idea is this: an agent capable of achieving human-level intelligence, necessarily needs to inhabit a body, be it physical or virtual - it cannot be a disembodied presence (a "brain in a jar", as it were). Further, such an agent needs to be educated and trained from 'infancy', similar to how a human baby would be. The agent's concept acquisition, the most essential aspect of intelligence, needs to be grounded via immediate experience of reality, and tied to relevant contexts; again, embodiment is the most direct way to ensure this. The agent would need to learn continuously and incrementally, by interacting with its environment, which includes acquisition of basic skills and capabilities, leading up to language acquisition, and higher levels skills related to reasoning and problem-solving; once again, it is difficult to achieve this, without embodiment.

Given the choice between physical and virtual embodiment, we are advocating for a virtual setup, for reasons we provide in an upcoming section of the paper.

In the following sections and subsections, we describe the system architecture, agent design, the agent's (virtual) environment, implementation, and what we might gain from creating such a system.

2 Architecture

We are about to describe a 'virtual reality operating system' (VR-OS) that would create both a virtual environment/world (VR-EN), as well as a virtually real agent (VR-AG) that would inhabit the environment. The VR-AG would learn and grow by interacting with its environment, and with a human 'teacher' who would appear to the VR-AG as another inhabitant of the VR-EN.

2.1 VR-OS

Much like a traditional OS, VR-OS is a platform for creating and executing multiple concurrent processes, scheduling tasks/events, managing resources (such as memory), providing I/O capabilities, etc. It would provide the following:

- a facility for creating and populating a 3D virtual world (VR-EN) - this includes being able to populate the world with objects that can exhibit behaviors, control the passage of time, schedule events (such as sunrise and sunset, growth of plants, weather events, etc), include 'physics' behavior such as gravity and collision detection
- facility for creating a 3D virtual embodiment ('VR-AG') that would live in the world, and interact with it; the VR-AG would contain a 'mind' implementation which will let it perceive, learn, grow, reason, etc.
- a human 'teacher' agent who interacts with the VR-AG, by being embodied as an avatar in the VR-EN
- 'God-like' ability to modify the VR-EN from outside the environment, and modify the VR-AG's internal state
- time-stamping and recording facility for capturing activities, events, interactions, VR-AG behaviors, VR-AG's mental processes, etc.
- editing, and playback, of recorded sessions
- real-time observation of happenings in the VR-EN and inside the VR-AG, by a human observer outside the virtual world

2.2 VR-AG

As mentioned earlier, the VR-AG is an agent with embodied cognition. Specifically, it would live in the VR-EN, experiencing it via its senses, embedded in a small human-like body with articulated limbs. The VR-AG would include a binocular first-person view of the world as seen through its eyes, binaural audio perception via microphones inside its ears, and be able to sense contact and touch via its skin. The agent would be able to observe itself, interact with the world (eg. navigate, pick up

small objects) and interact with its teacher as well (see, hear, get near, hold hands etc). Most importantly, the VR-AG would start out ('be born') with an almost "tabula rasa" brain/mind (pre-programmed with some basic instincts and behaviors, explained in an upcoming section), and be able to store in its memory, experiences, concepts, etc. as it learns and grows. The VR-AG would need to nourish itself via an energy 'food' source when it runs low on energy. The VR-AG's mind would be totally transparent, visible to an observer outside the VR-OS. Every change to the mind's memory, every cognitive process (such as perception, thought and emotion) and every interaction with the VR-EN would be recorded with a timestamp, for purposes of analysis, playback, etc.

2.3 VR-EN

The virtual world the VR-AG lives in, would contain geometry (large objects such as buildings, medium objects such as chairs, and small objects such as balls and books), materials that are associated with the geometry, cameras (situated in the VR-AG, and elsewhere desired), and photoreal rendering. The world would also contain 'game physics' phenomena such as non-interpenetration, gravity, sunlight, fracturing, bouncing, etc - such phenomena would be part of the VR-AG's 'qualia' and be recorded as experiences in its mind.

2.4 Implementation

We surveyed existing implementations of AGI-oriented as well as regular VR platforms, hoping to leverage the one closest to our proposed one. But to the best of our knowledge, such a system does not exist. What we found was too specific, or lacked functionality:

- Deepmind's Gym [1] is an environment for comparing reinforcement learning algorithms
- DeepMind Lab [2] is a first-person 3D game platform, for experimenting with deep reinforcement learning
- CogPrime [3] is a platform for embodied AGIs, much like what we are proposing - but it is not a VR system, and the underlying agent architecture is not easily replaceable with alternatives
- BabyX [4] is photoreal simulation of a baby's face, backed by biologically-accurate neural network to perform AGI-like functions such as perception and language acquisition - not open-source, and is not a VR environment
- there are several smaller implementations of specialized worlds (eg. a visually-guided 2D agent [5] but they are not extensible, by design

As a result, we are proposing building a system from the ground up. Here are some specifics regarding the implementation:

- the underlying VR-OS would be built using Node.js, a scalable, cross-platform, JavaScript-based runtime engine that can run code on multiple platforms (laptops, browsers, servers); it is a non-blocking system that runs

asynchronous code on multiple threads (eg. Pomelo[6] is a 3D game engine framework that runs on Node.js)

- WebGL, Three.js, BabylonJS and other graphics and game libraries (in JavaScript) could be used to create browser-based clients that can interface with VR-OS, making it possible to interact with the VR-AG via a phone browser attached to Cardboard VR, for example
- the VR-EN and VR-AG architectures would borrow heavily, features and design aspects from game engines such as Unreal and Unity
- the OS would provide dashboards with controls for monitoring and altering the VR-AG's internal states and configuration, for observing the VR-AG/VR-EN interaction as an invisible observer (so that the VR-AG is not aware of the extra presence), and for altering the VR-EN (eg. causing an event, re-positioning objects, etc)
- the OS would also provide an interface for exporting and importing JSON-based files that contain VR-AG/VR-EN interaction logs, the VR-AG's internal state (including its memory), the VR-EN's configuration, etc - in short, everything for modifying or repeating world events, interactions and agent behavior

2.5 Comparisons with unembodied AGI and physically-embodied AGI

We strongly agree with the following characterization by Pfeifer and Scheier [7], that 'Intelligence cannot merely exist in the form of an abstract algorithm but requires a physical instantiation, a body.' Without being able to experience physical sensation, navigation, self-directed gaze and other body-based phenomena, an AGI is simply not going to be as relatable to human, as an embodied equivalent would be - in the real world, we do not have human intelligence that functions without a body.

Physical embodiment of an agent in the form of a robot, is certainly an alternative to a VR-AG. However, robots are expensive to create, are prone to mechanical failures, have safety issues when interacting with humans, and are not trivially upgradeable as they 'grow up'. Further, robots operate in the real world which can be unpredictable and not controllable (events cannot be isolated, prevented, or replayed, for example); extra actors and objects in the scene might interfere with an AGI's learning, memory formation and problem-solving.

2.6 Shortcomings of the virtual AGI approach

In the above sections, we laid out all the advantages of being able to interact with a virtual agent in a virtual world - these chiefly come down to lower cost, total controllability, configurability, and repeatability.

On the flip side, it might not be easy to replace an object with another, for example (we might not have the replacer object in our assets library, or if we had one, it might lack the functionality we desire, because that was not coded in during creation). Further, the environment's capabilities would be only as good as the simulation (eg. we might not be able to tip over a paint can to have our VR-AG watch the paint ooze

out, because fluid simulation is not part of our physics yet). In our opinion, such drawbacks are not insurmountable (and need to be addressed just once), so we do strongly prefer the virtual approach.

3 Usage

In this section, we itemize, under subsections, a wide variety of experiments related to teaching the VR-AG, studying the VR-AG's mental states and behavior, simulating animal physiology and animal behavior (including body-related ones), evolution simulation, comparative studies between VR-AGs, etc., all of which can be carried out in the proposed platform. While many of these can be done in a non-VR environment as well, the VR platform provides richer possibilities, namely first-person viewpoint and 'natural' interactions for the agents (to the extent that their anatomy and brain modeling permits), and for the participating humans, observability (of external behavior, and internal states) from outside the VR-EN, and interactability from within.

The possible experiments are listed in the form of a wish-list, with no details on how to set them up (since specifics of the physical and mental construction of the agents would determine the setup).

3.1 Instincts and built-in capabilities

In a manner similar to how human and animal/bird/fish/insect babies are born with pre-wired instincts, a VR-AG could be outfitted with them from 'birth' so that the capabilities and behaviors tied to those instincts could be carried intuitively and automatically, without needing thought or memory. Each instinct would be coded as a pair: what conditions would activate it, and what the resulting behavior should be. In addition to instincts, the VR-AG would possibly need some innate capabilities that are related to its survival. Possible instincts and built-ins include:

- self-preservation (no damage to self, mostly related to bodily harm)
- dealing with hunger
- face recognition
- avoiding discomfort (eg. loud noise, bright light, shaking, too hot or cold exterior)
- fear of unfamiliar faces or situations
- motivation, curiosity to explore
- basic capacity for pleasure and pain
- tendency to seek reward (this could form the basis for reinforcement learning tasks)
- urge to socialize, and communicate

3.2 Things that could be taught to a 'toddler' VR-AG

A human toddler who attends pre-school, gets immersed in an environment that is conducive to learn a variety of foundational knowledge and gain basic experiences. A

VR-AG would likewise benefit from learning and experiencing similar items, which include:

- spatial navigation
- object invariance with respect to head rotation
- object permanence (for older agents)
- a sense of time
- language
- things: labels (names), parts/structure
- simple physical phenomena
- concepts (such as big, outside)
- feelings
- motor responses
- opposites
- number sense, counting
- basic weather
- handedness
- simple mechanisms (eg. via manipulatives such as toys)
- simple heuristics

3.3 Ways to train/teach a VR-AG

Some of the basic ways by which to train or teach a VR-AG (identical to techniques employed on human children) include the following [note that they all presuppose language understanding and basic speech generation, on part of the VR-AG]:

- showing and naming (objects, pictures on board books) - repetition would cause reinforcement
- using manipulatives - learning occurs by grasping, turning, etc
- performing simple experiments - for the VR-AG to watch, or repeat
- catechism - question and answer is a very effective technique that works on humans
- verbal exposition, including reading stories and singing nursery rhymes
- using reward and consequences, to reinforce and suppress behavior, respectively

3.4 Physical and mental states to simulate and observe

If our VR-AG is created with a need to consume energy ('eat'), and spend it by performing physical and mental tasks, we could invoke feelings of hunger, tiredness, sleepiness, and inattention by having the VR-AG deplete its stored energy.

Since we can directly observe the mental processes in action, we would be able to study the following aspects [if they are built into the VR-AG]:

- thinking - this would be the process of utilizing perceived inputs, memory-retrieved facts/skills/events, guided by expectation, influenced by feelings, to generate outputs ("thoughts") and actions
- feeling - eg. the feeling of fear, or peace

- memory formation, including semantic, procedural and episodic memories
- concept formation - including generalization by induction, and deductively situating new concepts under existing ones
- memory recall, including partial recall
- mis-remembering
- memory modification
- forgetting

3.5 Simulation of animal/bird behavior

Our VR-AGs do not have to be humans, they could be other animals, birds, fish, etc, whose nervous systems have been sufficiently studied so their anatomy, physiology and behavior could be accurately modeled and simulated in the VR-EN. Following are some interesting possibilities:

- panoramic vision in birds (coupled with limited binocular vision)
- echolocation in bats
- gaze holding, and quick gaze shifting (saccadic motions)
- footedness in birds, eg. parrots
- head bobbing motion in birds, eg. pigeons
- pecking, in birds
- 'plunge diving' by waterfowl, to catch fish
- prey centering (foveated imaging), eg. by eagles
- communication between bees, using 'waggle' dance [8]
- firefly flash synchronization, using a 'spiking' neuronal model for example [9]
- recreation of complete neural circuitry, for creatures whose entire neural atlas is known, eg. *C.elegans* [10] - for the purposes for studying stimulus response behavior, for example

3.6 More fanciful experiments

As the VR-AG becomes increasingly intelligent, with attendant brain mechanisms and structures modeled in, it would be interesting to indulge in more wishful pursuits such as these:

- teach the VR-AG to produce creative output: compose music, make art, generate poetry
- study the aesthetic response by having the VR-AG look at beautiful art or listen to a lovely tune
- have multiple VR-AGs interact, study socialization
- introduce multiple real-life avatars into the VR-EN, for the VR-AG(s) to interact with
- create a collective brain, by linking multiple VR-AG brains, study its behavior and capabilities
- graft parts of a VR-AG's brain on to another, and study the change in behavior of the target agent

- transplant a VR-AG brain on to a physical humanoid robot
- compare various cognitive architecture implementations [Alexei S], under identical VR-EN configurations
- using brain-computer interfaces (BCI), connect a human brain to a VR-AG's, study the effects on the human as well as the VR-AG
- provide additional senses to the VR-AG, eg. extra pair of eyes on the back of the head, extra sensing capability (eg. radio waves)
- study metacognition, and epistemology
- study the effects of brain damage and disease, on the VR-AG's brain

While the above might sound outlandish at this point, it is worth remembering that the goal of BICA researchers is nothing less than full-scale AGI, achieving which would make these experiments become feasible.

4 Summary

In this paper, we have made the case for pursuing AGI development using a virtual platform (instead of a disembodied, or robot-based one) to set up agent/environment interactions.

We discussed the various components of the proposed system (VR-OS, VR-EN, VR-AG), implementation specifics, and mentioned the pros and cons of our approach; then we briefly touched on a number of items that could be testable on the platform (VR-AG's instincts, mental states, memory formation and memory access; non-human cognition; more fanciful sensor and perception experiments; body-based cognitive behaviors, animal behaviors, etc).

5 Future Work

We note that the design we have discussed above, is at proposal stage; it has yet to be implemented, so that would constitute future work. That said, a 'version 2' rollout could expand the possibilities, via the enhancements listed below.

Future capabilities of VR-OS would include:

- access via the Internet, for remote 'teacher' (or teachers) participation
- capability to set up multiple, simultaneous human avatars that can populate the VR-EN and engage with the VR-AG
- a 'cloud VR' setup, where the VR-EN would be hosted as well as rendered on a cloud server, which would stream rendered video to its remote clients - clients would not need any setup or installation, they would simply participate via web browsers

Future capabilities of VR-AG would include the following.

Body:

- more senses, such as being able to feel hot and cold, feel the wind, feel surface texture
- enhanced biomechanics, including a muscle and skin layer, bipedal gait synthesis

Brain:

- proprioception
- depth perception

And finally, here are new VR-EN-related features to add:

- more VR-AGs, for social interaction and communication
- more geometry, for a more interesting world
- more physics, eg. weather phenomena
- more places, for the VR-AG(s) to visit

References

1. Brockman, G., et. al.: OpenAI Gym, <https://arxiv.org/abs/1606.01540> (2016)
2. Beattie. C., et. al: Deepmind Lab, <https://arxiv.org/abs/1612.03801> (2016)
3. Goertzel.B: CogPrime: An Integrative Architecture for Embodied Artificial General Intelligence, https://wiki.opencog.org/w/CogPrime_Overview (2012)
4. Lawler-Dormer. D: BABY X: Digital artificial intelligence, computational neuroscience and empathetic interaction, ISEA 2013 Conference proceedings (2013)
5. Beer, R. D.: Toward the evolution of dynamical neural networks for minimally cognitive behavior. In: From animals to animats 4, 421-429 (1996)
6. Pomelo, <https://github.com/NetEase/pomelo>, last accessed 2019/5/15
7. Pfeifer. R., Scheier., C: Understanding Intelligence, MIT Press, Cambridge, MA (1999)
8. Fernando. S., Kumarasinghe, N.: Modeling a Honeybee using Spiking Neural Network to Simulate Nectar Reporting Behavior. In: International Journal of Computer Applications (0975 –8887), Volume 130, No.8 (2015)
9. Kim. D.: A spiking neuron model for synchronous flashing of fireflies. In Bio Systems 2004, DOI:10.1016/j.biosystems.2004.05.035 (2004)
10. WORMATLAS, <https://www.wormatlas.org/neuronalwiring.html>, last accessed 2019/5/15