Computatrum Project Spec

Jacob Valdez

June 22, 2022

Abstract

It's time for artificial intelligence (AI) to grow up. State-of-the-art deep neural networks today are relatively hardwired considering that they must be spoon-fed datasets, explicitly trained, and operate under an economic fitness landscape that does not always align with their training objective. Even at the iteration speed of an expert research team with high end equipment, there are just too many accidental and essential complexities, and not enough automation driving general AI evolution from end-to-end. Approaching and surpassing that rate-limiting bar of human research and development demands liberating as many aspects of the development cycle as possible to autonomous control. Targeting this problem, I propose a 4 ambitious iterations to the AI R&D cycle: 1) a highly complex, multi-task, open-world learning environment: the Artificial Experience which includes 2) a general purpose computer interaction environment: ComputerEnv, 3) a novel multimodal, multiagent, multiparadigm deep learning architecture: the Multi-Agent Network (MAN), and 4) an open-ended full stack integration of these components: Computatrum. This paper provides motivation, details, and acceptance criteria to the above – possibly foundational contributions paving the way to artificial general intelligence.

1 The Problem

The realization of sufficiently advanced artificial intelligence (AI) – especially that vaguely referred to by phrases "human-level AI", "artificial general intelligence", and "superintelligence" – has strong and clear motivation from many fields of human endeavour including healthcare, business, economics, governance, and, of course, the STEM disciplines. Still, there are no silver bullets, and the gradual evolution of AI has been artificially rate-limited in many respects. Ever since its inception, there has not been a general consensus on the/a formal measure of intelligence, goal for autonomous agent action, or unifying framework to guide AI research, and even today, state-of-the-art deep neural networks are relatively hardwired considering that they must be spoon fed datasets, situated in closed environments, and train under an economic fitness landscape that they have no direct awareness of. In effect, they are treated like

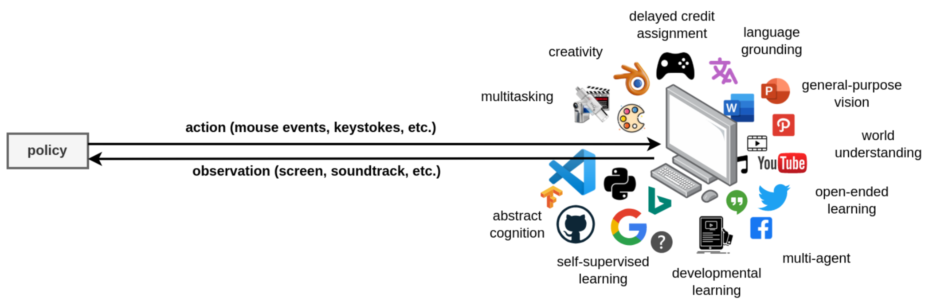

Figure 1: The general-purpose computer provides reasonable coverage over the anthropocentric problem domain. Performance across many tasks in this open-world therefore gives a proxy of development towards 'artificial general intelligence'. The Computatrum project aims to develop a policy capable of attaining reasonable performance in this domain with and possibly without supervisory feedback signals.

infants. Approaching and surpassing the rate-limiting bar of human research and development demands liberating as many aspects of the AI development cycle as possible to autonomous control. Given this immense problem domain, I ask:

> Can a fully autonomous, open-ended AI system research and develop state-of-the-art AI systems – including improvements of itself – subject to the same technological and financial constraints as an independent researcher?

This is not simply asking for autoML, intrinsically-motivated reinforcement learning, or a formal AI-generating algorithm. While each of the aforementioned represent important approaches to engineering highly advanced, open-ended learning AI, I propose synergizing them all and many more salient features from the bleeding edge of AI R&D into a system that genuinely propagates feedback 'end-to-end', maintains a highly self-informative internal dynamical state equatable to consciousness, and uses standard peripherals connected to a Ubuntu VM with Internet access to interact with the real, open-ended Environment including robots, research sites, and its own software and compute resources. I call this system Computatrum, and set forth the following fundamental milestones to its development:

- a highly complex, multi-task, open-world learning environment: the Artificial Experience,
- a general purpose computer interaction environment: ComputerEnv, and
- a novel multimodal, multiagent, multi-paradigm deep learning architecture: the Multi-Agent Network.

The following three sections introduce those milestones in their respective order (Sections 2, 3, 4). Then I provide detailed objectives for Computatrum (Section 5). Finally, I define unambiguous acceptance criteria for the integrated system (Section 5).

2 The Artificial Experience

The problem

AI development implicitly demands elucidating the dimensions of variation and relational structure underlying the problem domain and then iterating development around that data. For datasets, this includes questions as:

- *What problem domain does the data belong to?* Examples include natural language modeling, image infilling, audio transcription, etc. Most datasets belong to several problem domains. Imagining each dataset as a vector embedding clustered in the high dimensional, sparely occupied space of all data, we might consider whether it is possible to make a machine teaching agent that selects small candidate data from each region in the subspace to represent a general domain dataset.

- How is the data structured and represented? Examples include hand-engineered feature vectors, categorical tokens ordered in a sequence for text, 4D float32 tensors representing images, etc. Much recent work in machine learning has focused on exploiting peculiarities of individual modalities while significant progress remains to be undertaken in transferring domain-specific techniques to machine learning in general.

For environments we might also consider:

- *Is the environment simulated or real?* Real data often comes at a higher cost than simulated data, yet it also often retains the rich complexities of the real world. Simulated data is fast food. While it can be easily manufactured and is often useful in passing tests, a neural network raised on simulations alone may have a difficult time adapting to the real world.

- *Is the environment single-agent or multi-agent? Are they situated in competitive, cooperative, or mixed-mode interactions?*

- *Does the environment have a scalar objective? a multidimensional objective? or any objective?*

Ideally, we should train increasingly general ML systems over representative data from all of these dimensions of variation. Still, many training pipelines are very brittle, and no work to my knowledge provides developers with a single, simple dataset or environment that integrates hundred of tasks, datasets, environments, and modalities.

Proposal

The `artificial-experience`[1] is a library to facilitate training and evaluating models, optimizers, pipelines, and training paradigms across dozens of tasks, domains, dataset loaders, environments, and hubs simultaneously, lifelong, and in-context. It provides convenience wrappers to automatically annotate existing dataset loaders, environments, and hubs with modality information which can be used by agents to dynamically generate or select appropriate encoders and decoders for each input and output modality. This library also provides a highly complex, open-world, multi-task, multi-agent, multi-paradigm (supervised, self-supervised, unsupervised, reinforcement, and meta) learning environment the `ArtificialExperience` which can be used quickly run AGI experiments. It is agnostic to the actual training paradigm and 'tricks' employed (such as augmentations, experience replay, curriculum learning, etc.) but can integrate cleanly with tools that do.

3 ComputerEnv

The Problem

The general-purpose computer provides a simple interface to vast distributions of natural and synthetic complexity which I believe reasonably proxy the anthropocentric problem domain. It inherently includes any dataset that machine learning practitioners might use, billions hours of recorded audio and video, live social media feeds, uncountable scientific, engineering, business, and historical documents, as well as creative software, integrated development environments, simulators, engineering design tools, e-commerce platforms, business systems, and many more applications. Considered together with the Internet, the general-purpose computer is a ready-made multiagent, language-grounded, lifelong-learning environment-incubator for the development/evolution of progressively more capable, general, and autonomous artificial intelligence.

Proposal

`computer-env`[2] is a Python library that provides gym-style environments for computer interaction. It provides the following ready-made environments:

- `LocalGUIEnv`: observes and interacts with the host machine. Modalities: keyboard, mouse, bitmap display(s), internal soundtrack.
- `VNCGUIEnv`: observes and interacts with a VNC server. Modalities: keyboard, mouse, bitmap display. When used in conjunction with docker, developers can emulate multiple computer environments on the same host machine with minimal overhead.

[1] https://github.com/Limboid/the-artificial-experience
[2] https://github.com/Limboid/computer-env

- `StdIOEnv`: observes and interacts with the standard I/O of a headless console. Modalities: keyboard, terminal display.
- `ChromeEnv`: observes and interacts with a headless Chrome browser. Modalities: keyboard, mouse, bitmap display.
- `AndroidEnv`: observes and interacts with an Android. Modalities: keyboard, touchscreen, audio I/O.

`computer-env` also provides a collection of `peripherals` which can be combined to meet the needs of custom computer interaction setups.

4 The Multi-Agent Network

The Problem

We cannot expect the same algorithm to make increasing returns on all problems. ML R&D has already burned millions of dollars, hours, and tonnes of CO2 into the parameters of large, modality-specific *foundation models* (such as bert, resnet-50, posenet) with no end in sight. Can these existing be synergized into a general multimodal vision-language-audio architecture with minimal retraining? Additionally, is it feasible to focus comparable attention towards developing *foundation policies* – neural networks that can be fine-tuned with minimal effort to "achieve goals in a wide range of environments"?

Proposal

The Multi Agent Network (MAN)[3] is a Python library geared towards meta-learning and manually orchestrating networks of specialized and general-purpose deep learning agents. It builds on the following concepts

- A *MAN* is a network of *agents*.
- Agents communicate with each other using *connections* between *ports*.
- ports have a *value*, *modality*, and *type*. The MAN uses this information to perform unsupervised statistical analyses and automatically grow or prune connections between compatible ports.
- The port's modality has a *structure* (vector, set, grid, or graph), *data_type* (Boolean, categorical, range, integer, string, and real), and textual *description*. Modalities with a set, grid, or graph structure are composed of child modalities.
- The port's type can be *bottom*, *top*, *side*, *reward*, or a domain-specific type.

[3] https://github.com/Limboid/man

- *Pyramidal agents* are a subclass of agent types that have a bottom, top, side, and reward port and use *sparse distributed representations* to communicate between ports.
- *Reward agents* supply intrinsic rewards to the reward port of other agents using information theoretic metrics based on global network connectivity and activity.
- The Multi-Agent Network library provides convenience wrappers to identify the modality and type of pretrained models to make corresponding ports.
- *Dataset agents* and *Environment agents* are used to interface the MAN with external data.
- During the update step of a MAN, all agents are updated in parallel using values from the previous time step. Agents use a hardwired or learned *trigger function* to determine if they should update at all.

5 Computatrum

Putting it all together

`Computatrum` is the composition of the above three components. A `MAN` instance undergoes general pre-training in the `ArtificialExperience`. Then it is fine-tuned for conditional and unconditional computer interaction in `computer-env` environments. In the conditional case, the Computatrum initially trains against a small amount of human demonstrations in `LocalGUIEnv` to learn an appropriate reward function. Then it alternates between task-conditional learning and unconditional exploration.

Acceptance Criteria

Targeting the loosely-defined, open set of computer interaction tasks is not simple due to their non-stationary distribution. This is further complicated by heterogeneous user interfaces and context-sensitive application of natural world metaphors such as location, navigation, and gesture. Then there is also the issue of estimating task progress, completion, and reward in spite of shifting and overlapping task boundaries. Finally, the complexity and critical thinking demanded by 'real-world' problems often overshadow those experienced in an automated testing suite.

To alleviate these issues, I propose the following unambiguous final acceptance criteria:

> Can the Computatrum classify (demonstration, label) pairs as correct or incorrect?

This is essentially a binary classification problem. For example, after observing the mouse movements, keyboard events, and display frames corresponding to a user sending an email to `user@email.com`, the computatrum should identify "sending an email to `user@email.com`" as a correct label. It should identify "downloading and installing `program`" as an incorrect label. A Python implementation of this evaluation objective is available at https://github.com/Limboid/computatrum/blob/main/task_guided_behavior_distilation/task_guided_behavior_distilation.ipynb.

6 Conclusion

This proposal described my plans to develop a revolutionary data-hunting, multitasking, multi-paradigm learning agent, Computatrum. Later work will detail progress on the Artificial Experience, ComputerEnv, Multi-Agent Network, and Computatrum. Admittedly, much work will remain even after successfully realizing the aforementioned proposals in order to reach the initial ambitious goal of developing a "fully autonomous, open-ended AI system research and develop state-of-the-art AI systems – including improvements of itself". We (computatrum and I) will work on the following towards this end:

- `the-artificial-school`: a training library to facilitate optimizing high-level, abstract learning objectives
- `ainimal-zookeeper`, `ainimal-zoo`, and `ainimal-leash`: an ecosystem for studying, caring for, reigning, and interacting with **ai**nimals
- Limboid: an inexpensive, high-reconfigurable robot for direct world interaction.
- Massive-MAN: a highly-distributed massive conglomerate of multi-agent networks; superintelligence.
- BoidNet: a spatially distributed network of limboids for Massive-MAN to control.

Please see https://github.com/Limboid/the-artificial-ecosystem for the big picture, roadmap, and latest exciting developments.