



Deep Learning for Sensor-Based Human Activity Recognition

Rashid Ayub¹, Prasadu Peddi²

¹Research Scholar, Sunrise University, Alwar

²Professor, Sunrise University, Alwar

Article Info: Received: 2-07-2023 / Revised: 04-08-2023 / Accepted: 20-08-2023

Address for Correspondence: Rashid Ayub

Conflict of interest statement: No conflict of interest

Abstract

The apparatus we used to test and refine our Bounded Rationality HMM answer to the BRAI. In particular, we looked at two simulation settings, each of which has its own quirks. To begin, we used Deep learning, a fully visible robotic mining simulation in which an agent must use a microscope to detect the mineral composition of a mine, and then gather those minerals in order to perform tasks. The microscope's accuracy degrades during operation (depending on the sort of test being conducted), but recovers over time when it is not in use. Second, we used a simulation for partly observable user preference elicitation dubbed User Rec, which is based on the work of Doshi and Roy (2008) and describes a scenario in which an intelligent user interface agent must ascertain a user's choice through interruptions that solicit data from the user.

Keywords: Deep Learning, Recognition, Sensor-Based, Human Activity

Introduction

Human activity recognition has recently advanced, opening up a wide variety of potential uses, including in smart homes, healthcare, and improved manufacturing. Activity recognition is crucial to society because it provides a record of people's actions that can be used by computers to keep tabs on them, evaluate them, and help them in their everyday lives. Video-based systems and sensor-based systems are the two primary types of human activity identification technology. Cameras in video-based systems are used for behavior recognition. In order to dead reckon a person's mobility data or register their activity tracks, sensor-based systems use sensors on the body or in the environment. Due to concerns about invasion of privacy, sensor-based systems have largely replaced the use of cameras in tracking our everyday activities. In addition, sensors make use of this pervasiveness. The widespread availability of smart devices and the Internet of Things has made it possible to put sensors in previously inaccessible places, such as in automobiles, walls, and even furniture.

Everywhere we go, sensors record data about our movements without disturbing us.

Human activity recognition is an area where machine learning techniques have been widely used in prior publications. Time-frequency transformation, statistical methods, and symbolic representation are only a few examples of the feature extraction techniques used extensively. However, the recovered characteristics are heuristic and the extraction process is carefully developed. There was a lack of systematic or universal feature extraction methods that could reliably capture unique aspects of human actions. In several fields, including computer vision, natural language processing, and voice processing, deep learning has recently seen striking success in modeling high-level abstractions from complex data. There has been an explosion of research into the application of deep learning to the problem of human activity identification since the publication of seminal papers such as. New efforts are being made to tackle these unique obstacles as deep learning continues to advance

in the field of human activity detection. Researchers are hesitant to embrace deep learning because of its rapid rise to prominence, rapid pace of invention, and lack of theoretical foundation.

Literature Review

Marwala, Tshilidzi. (2013). We revisit an extension of the bounded rationality theory called flexibly-bounded rationality in this study. Making judgments using rational methods necessitates employing defective and partial knowledge in conjunction with an intelligent machine that, if it is a person, is inconsistent. In limited rationality, the choice is taken within the confines of these restrictions, despite the fact that the information to be employed is partial and imperfect and that the human brain is inconsistent. The notion of flexibly-bounded rationality makes use of AI to make better judgements and sophisticated information analysis tools like the correlation machine to fill in any gaps in knowledge. Therefore, the scope of rational thought is broadened by the use of flexibly-bounded rationality. This study introduces the idea of marginalization of irrationality in decision making to address the challenge of satisficing when irrationality is present, given that human decision making is inherently illogical.

Gama, João. (2013). New information and communication technologies have resulted in a sea shift in how data is gathered and processed. The age of constrained rationality has arrived in data mining. In this paper, we explore how the data stream computational model's resource restrictions affect the development of learning algorithms. We provide an in-depth analysis of how stream mining algorithms function and outline potential avenues for further study, such as ubiquitous stream mining and self-adaptive models.

Bettis, Richard et.al. (2018). The Turing Award in Computing was shared by Alan Newell and Herbert A. Simon for their seminal contributions to the field of Artificial Intelligence. In addition to his Peace Prize, Simon took home the Nobel Prize in Economics for his work on "bounded rationality." The same core heuristic, "search till a satisfactory solution is found," was used in both instances. We

suggest that the field of behavioral strategy has a lot to gain from the study of computational complexity and AI. These areas of study may strengthen the theoretical underpinnings of constrained rationality and the need for and use of heuristics. Last but not least, it may be helpful to apply a notion of "organizational intractability" inspired by the metaphor of the Theory of Computational Complexity to figure out which analytical decision-making tools are really unworkable in real-world settings due to time and management attention limits.

Russell, Stuart. (2016). AI's ultimate objective is to enable the development of, and insight into, human intellect. For this, we need a clear definition of intelligence that will enable the construction of reliable systems and broad conclusions over time. The idea of rational agency has been a frontrunner for this position for some time. This study, an updated version of one first published in 1997 (Russell, *Artif Intell* 94:57-77), examines the line of thinking that eventually led to an alternative contender, limited optimality, that is more in line with our everyday understanding of intelligence. Some encouraging progress made in recent times is also discussed.

Nobre, Farley et.al. (2019) This article discusses a novel and cohesive viewpoint that helps participants (particularly the new economic men) in the organization push the limits of their reasoning. This view has its roots in the literatures of fuzzy logic, limited rationality, and cognitive psychology, all of which emerged about the same time in the second half of the twentieth century. This idea is bolstered by an interview with Professor Lotfi A. Zadeh (the founder of fuzzy logic) conducted in 2012 at the University of California, Berkeley. The study's findings suggest that (a) fuzziness, rather than probability, is the type of uncertainty that most pervades decisions in organizations; (b) fuzzy logic goes beyond bounded rationality by supporting the latter with new mathematics to solve non-programmed and ill-structured problems of unknown probability distributions; and (c) bounded rationality and fuzzy logic are complementary to one another.

Bounded Rationality HMM Instantiation

Each microscope in a Deep Learning environment is a stateful resource used for sensing, meaning its behavior is contingent on its current state (its energy level). This is because the usage of a microscope depletes its battery, resulting in less precise observations. Because of the inaccurate observations introduced by Bounded Rationality, an agent must weigh the importance of refining its knowledge to complete its present duties against the risk of compromising its knowledge. As indicated in Table 1, we model the decision-

making process for implementing sensing actions using a microscope as a Bounded Rationality HMM to account for this trade-off. To emphasize the importance of the constant prior $u=1$, we model the current state of knowledge as the number of preceding observations with a belief equal to $b + d$ in equation (16). Finally, we use equation (16) to describe the value of knowledge refinement as the difference between the agent's initial and updated maximum expectations for the mineral type in each supply.

Table 1: Deep learning Bounded Rationality HMM

| Bounded Rationality HMM | Deep learning Characteristics |
|--|---|
| Stateful Resource | Electronic Microscope |
| Sensing States $S = \{ \langle R_s, K_s \rangle \}$ | R_s = Microscope Energy, K_s = # of Previous Observations |
| Activity Choices A | Advanced/Basic Mine Test, Wait |
| Transition Probabilities $T(s, a)$ | Change in energy from a test or self-recharge during wait, increased number of observations |
| Knowledge Refinement Reward $R(s, a)$ | Increase or decrease in they believed possibility of the true mineral in the supply |

We take into account three RL methods to learn how to function in this Observation Selection HMM: The first is Q-Learning (Watkins, 1989), second is R-Max (Brafman and Tenen Holtz, 2002), and third is REINFORCE (Williams, 1992). These algorithms were selected not because we wanted to conduct a comprehensive analysis of cutting-edge RL methods, but because they reflect a variety of RL approaches and are either widely used or easy to understand or have unique qualities.

Since Q-Learning and R Max have already been detailed in Sections 3.3.1 and 3.3.2, respectively, we will not repeat that information here. Conversely, REINFORCE (Williams, 1992) is a family of model-free RL algorithms that employ neural networks to train both an action selection controller and the reward function. The agent learns two neural networks: 1) a stochastic neural network that estimates the reward function based on the current state and the action taken, and 2) a neural network that approximates the reward function based on the

current state and the action taken. The latter is taught using conventional supervised learning techniques, whereas the former is taught by adding reinforcement to the network's weights in proportion to the value of the reward for performing the desired action. To train the network using back propagation, we employ the following update function in our experiments

$$\alpha(r(s, a) - R(s, a))e$$

Where a , $r(s, a)$, and $R(s, a)$ are as defined previously (c.f., Section 3.3.1) and is the computed eligibility of weight updates (Williams, 1992). Of note, we reinforce using $r(s, a) - R(s, a)$ rather than just the recent $r(s, a)$ as our learned $R(s, a)$ gives the algorithm a starting point for rewards, reducing learning-time variation.

Q-Learning and R Max, two of the most popular model-free and model-based RL algorithms, are quite straightforward. Both of these methods, however, rely on discrete states,

which leads to the loss of some information when applied to the genuine sensing state. Since the energy state of the microscope is continuous, we discretize the data into bins spanning the range [0, 1] for use in deep learning. However, REINFORCE's neural network foundation means that we may preserve continuous values even while we learn.

Human Activity Recognition: A User Preference Elicitation Simulation

Environment Description

Human activity recognition is the second simulated environment we utilize in our tests; it represents a user choice elicitation issue. Doshi and Roy (2008) suggested this simulation setting to test a novel PBVI-based approach for resolving the preference elicitation POHMM (Boutilier, 2002; Doshi and Roy, 2008). This part starts off by and then we add more to it to make it more like a model of the actual world by integrating things like user anger when things don't work out the way they're supposed to.

Table 2 Deep learning Experiment Parameters

| Parameter | Value |
|-----------------------------|----------------------------|
| Grid Size | 20 x 20 |
| # Supplies | 20/type |
| # Agents | 30 |
| # Tasks | 50 |
| # Agents/Task | 5 |
| Microscope Recharge Rate | 10%/time unit |
| Max Advanced Test Energy | 50% |
| Max Basic Test Energy | 25% |
| Advanced Test Accuracy | 0.8 |
| Basic Test Accuracy | 0.4 |
| Belief Confidence Threshold | 0.65 |
| Learning Rate α | 0.3 |
| # Discrete Sensing States | 400 |
| <i>NF</i> | 0, 0.1, 0.2, 0.3, 0.4, 0.5 |

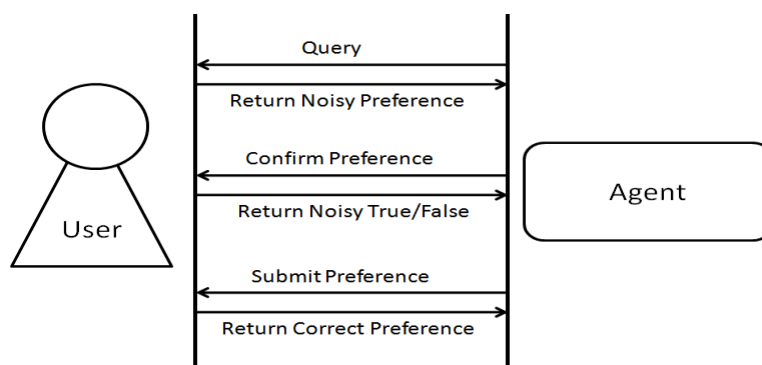


Figure 1: Human-Agent Interaction in Human activity recognition

Doshi and Roy’s World. A human user is supported by an intelligent agent in the initial setting presented by Doshi and Roy (2008). For an illustration of how an agent communicates with its human user, see Figure 1 below. An intelligent agent interacting with a human user must elicit the user's preference over a set of items (e.g., goods, scenarios, goals) in order to provide support to the user. The agent may use two sensory operations to achieve this goal:

1. *query*, which asks the user to state their current preference, and

2. *confirm*, whereby the agent asks the user if its belief about the user’s top preference is correct.

When the agent asks for confirmation of actions based on the user's real preferences, the user gives feedback in the form of observations that are utilized to update the agent's beliefs (i.e., knowledge) about the current user preference. Once the agent is certain it has grasped the user's choice, it may submit a third action to carry out the intelligent help the user has requested. A user's choice is elicited via a sequence of interactions, and each episode concludes with the submit button being clicked.

After a user clicks "submit," their preferences are cleared and a fresh episode is played.

To more closely reflect the actual situation, the user's input is noisy in the sense that it contains a small but non-zero likelihood of being erroneous (according to a fixed probability for each of the two sensory activities indicated above, as given in Table 3). This noise causes the agent to be less confident of its beliefs, so it may execute a series of inquiry activities before

committing to a confirm action. Also, throughout the course of a contact with the agent, there is a small chance (also shown in Table 3) that the user will change its choice. For this reason, it might be challenging for an agent seeking to model its user's desire to reconcile all replies in its beliefs about user preference, even if it first obtains a correct response from the user.

Table 3: Example Environment Parameters (Doshi and Roy, 2008)

| Environment Parameter | Value |
|-------------------------------------|-------|
| Correct Query Response Likelihood | 50% |
| Correct Confirm Response Likelihood | 80% |
| User Preference Change Likelihood | 1% |
| Number of Possible User Preferences | 10 |

Table 4: Example Task-Level Reward Structure for Agent Actions (Doshi and Roy, 2008)

| Action | Reward |
|------------------------------|--------|
| Query Preference | -2 |
| Confirm Correct Preference | -1 |
| Confirm Incorrect Preference | -5 |
| Submit Correct Preference | 100 |
| Submit Incorrect Preference | -200 |

In addition, there is a constant penalty (i.e., punishment) for any sensory actions done by an agent. Since the agent wants to submit the user's choice as accurately as possible, it will try to avoid executing unnecessary sensing activities if it already has enough information about the user's desire. Table 4 reproduces an example reward structure provided in to aid an agent in deciding whether to undertake an action and which actions to take. Because it requires more work from a user to respond over a set of preferences than to agree or disagree with the agent's top belief, the given reward structure penalizes an agent more when its confirm and submit actions are incorrect (i.e., the wrong preference is identified).

Doshi and Roy (2008) take into account a POHMM model of the environment dubbed the Preference Elicitation PODMP (Boutilier, 2002), which is detailed in Table 5, while designing agent behavior to address the preference elicitation issue in this environment. In this context, an agent is provided with a fully parameterized PODMP model of the environment (developed by an assumed domain expert) and then uses versions of the PBVI algorithm to choose which actions to take in light of its present belief state. Here, the agent's belief state is representative of its understanding of user preferences, as it quantifies the probability that the agent has correctly identified the user's true choice among all potential options.

Table 5: Preference Elicitation POHMM Model

| POHMM Model | Preference Elicitation Problem |
|---------------------------------|--|
| States | User preferences |
| Actions | Query, Confirm, Submit |
| Observations | User preference or true/false |
| State transitions probabilities | Small chance of preference change during an episode |
| Observation probabilities | Likelihood of correct/incorrect responses to sensing actions |
| Reward | Reward structure |

Our Extension: User Frustration. For a more realistic simulation of the environment's dynamics, we included user irritation that builds up across several elicitation episodes and is exacerbated by the agent's disruptions. As faulty behaviors reduce user confidence in the system and motivation to utilize the system, they also raise frustration when the agent wrongly acts on its ideas about user choice. As a result, she becomes more irritable, which in turn causes a disruption in her cognitive state, and she responds more rapidly because she feels rushed to make up for the time she missed because of the interruption. In contrast, when the agent delivers appropriate intelligent assistance based on true ideas about user desire, frustration and its side effects are reduced.

There are no measurable mathematical models for computer user irritation that we could

include into our implementation after searching the human-user interaction (HCI) literature and the research done by the intelligent user interface (IUI) community. To kick off our investigation, we recommend the following as a rough estimate of customer dissatisfaction.

We represent user dissatisfaction as a cumulative effect, where agent actions either raise or reduce the user's frustration, to keep things simple while yet include user frustration inside our simulated program. We express a user's degree of annoyance as a number between zero and one hundred, with one hundred being the most extreme level of dissatisfaction. In particular, the cumulative effects of the negative outcomes of each agent's decisions contribute to the user's sense of exasperation. In Table 6, we illustrate one such value.

Table 6: Example Frustration Structure for Agent Actions

| Action | Frustration Increase |
|------------------------------|----------------------|
| Query Preference | 1 |
| Confirm Correct Preference | 0.5 |
| Confirm Incorrect Preference | 2.5 |
| Submit Correct Preference | -10 |
| Submit Incorrect Preference | 10 |

What makes our add-on so intriguing is that the user's reaction time (she reacts quicker to get back to her task) and accuracy (she responds less correctly due to disturbed cognitive state) both fluctuate depending on her current annoyance level. In this work, we describe the evolution of the reaction time and the accuracy drop as being linearly related to the level of user annoyance at the moment. In particular, the

delay timer has a maximum setting by default. (MAX_DELAY) when the user is not experiencing any frustration, and is determined by subtracting the frustration multiplied by the frustration delay factor (FDP) from the previous maximum:

$$delay = MAX_DELAY - FDP * Frust$$

(Capped to provide an integer number of simulation ticks representing time), whereas the noise introduced by an increase in user irritation is determined by multiplying frustration by a frustration noise factor (FNF):

$$\text{noise} = \text{FNF} * \text{Frustr}$$

Finally, we halt our simulations early (i.e., before all episodes are done) if the user reaches maximal frustration at the conclusion of a particular number of episodes in a row, since a human user cannot be expected to continue working when utterly irritated. This is the user's boiling point count, and it indicates how long they can remain frustrated at a level higher than their personal threshold (100% frustration).

Conclusion

The user's patience becomes thin, and her reaction speed and precision suffer as a result. As a result, the agent in both settings must solve the BRAI in order to make informed decisions about what sensing activities to engage in because the state of the resource (energy for the microscope and frustration for the user) influences the accuracy of the sensing outcomes. Our modeling of the Bounded Rationality HMM and the reinforcement learning techniques used to regulate agent sensing have been detailed for both simulations.

Reference

1. Marwala, Tshilidzi. (2013). Flexibly-bounded Rationality and Marginalization of Irrationality Theories for Decision Making.
2. Eck, Adam. (2010). Agent Sensing with Stateful Resources. 1283-1284.
3. Gama, João. (2013). Data stream mining: The bounded rationality. *Informatica (Slovenia)*. 37. 21-25.
4. Bettis, Richard & Hu, Songcui. (2018). Bounded Rationality, Heuristics, Computational Complexity, and Artificial Intelligence. 10.1108/S0742-332220180000039010.
5. Viale, Riccardo & Gallagher, Shaun & Gallese, Vittorio. (2023). Bounded rationality, enactive problem solving, and the neuroscience of social interaction OPEN ACCESS EDITED BY. *Frontiers in Psychology*. 14. 1152866. 10.3389/fpsyg.2023.1152866.
6. Russell, Stuart. (2016). Rationality and Intelligence: A Brief Update. 10.1007/978-3-319-26485-1_2.
7. Nobre, Farley & Walker, David. (2019). Advancing Bounded Rationality with Fuzzy Logic and Cognitive Machines in the Organization.
8. Lewis, Richard & Howes, Andrew & Singh, Satinder. (2014). Computational Rationality: Linking Mechanism and Behavior Through Bounded Utility Maximization. *Topics in Cognitive Science*. 6. 10.1111/tops.12086.
9. Marwala, T.. (2015). Causality, correlation and artificial intelligence for rational decision making. 10.1142/9356.
10. Jarrahi, Mohammad Hossein. (2018). Artificial Intelligence and the Future of Work: Human-AI Symbiosis in Organizational Decision Making. *Business Horizons*. 61. 10.1016/j.bushor.2018.03.007.