Cognitive LLMs: Towards Integrating Cognitive Architectures and Large Language Models for Decision-making in Manufacturing

Siyu Wu, CR/RSI1-NA, College of IST, Penn State

Advised by: Dr. Alessandro Oltramari

Aug 6th, 2024

Cognitive LLMs
Challenges for LLM-based Decision Making

- Hallucinations
- Noisy inference behavior
- Lack of trust
- Bad predictions and increased risks

Decision Intelligence for Production Optimization in Manufacturing

The Three Levels of Decision Intelligence

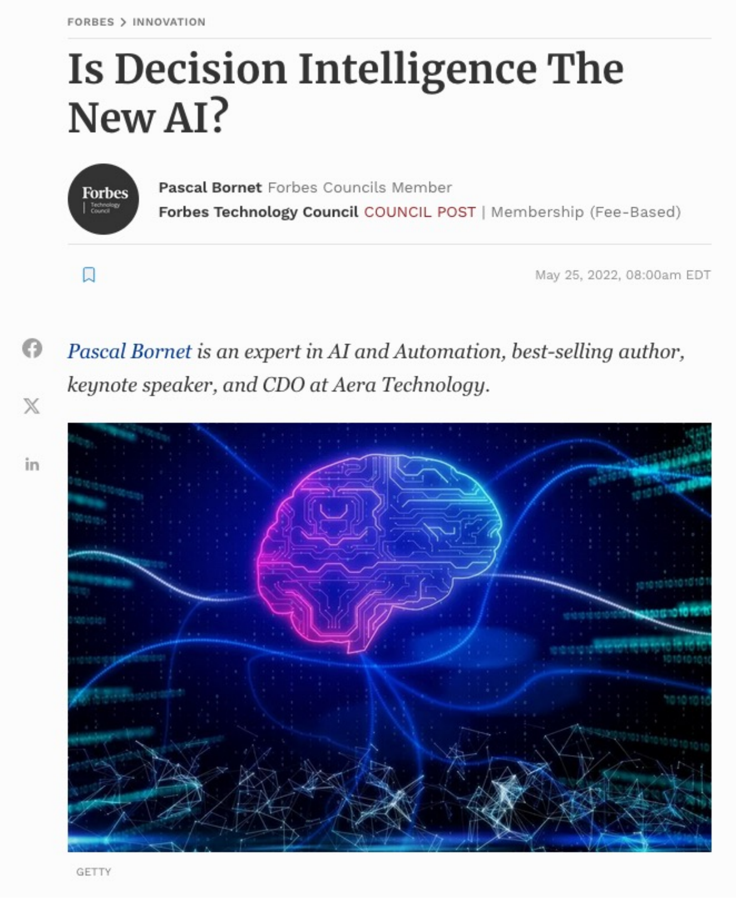
FORBES > INNOVATION

Is Decision Intelligence The New AI?

Pascal Bornet Forbes Councils Member
Forbes Technology Council COUNCIL POST | Membership (Fee-Based)

May 25, 2022, 08:00am EDT

Pascal Bornet is an expert in AI and Automation, best-selling author, keynote speaker, and CDO at Aera Technology.

GETTY

Decision Support

- machines provide some basic tools to support human decision making, such as alerts, analytics and data exploration. The decisions themselves are made entirely by humans

Decision Augmentation

- machines play a larger and more proactive role in the decision process. They analyze the data and generate recommendations and predictions for decision-makers to review and validate. Humans can make decisions based on the machine's suggestions, or they can work cooperatively with the machine to amend the recommendation.

Decision Automation

- machines perform both the decision step and the execution step autonomously. Humans have a high-level overview, monitoring the risks and any unusual activity and regularly reviewing outcomes to improve the system.

Decision Intelligence
A Common Model of Cognition

- **Cognitive Architectures** are computational frameworks that capture the invariant mechanisms of human cognition, including those underlying the functions of attention, control, learning, memory, adaptivity, perception and action
- A **Cognitive Model** is the software artifact resulting from using a Cognitive Architecture to model a human task
- The **“Common Model of Cognition”** is a recent development to consolidate four decades of research

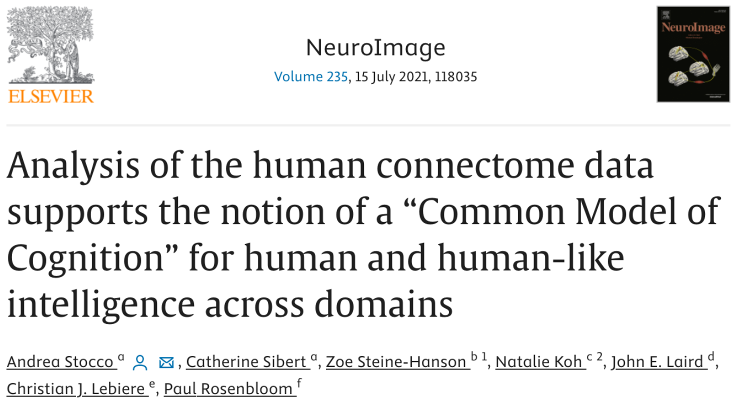
ELSEVIER

NeuroImage
Volume 235, 15 July 2021, 118035

Analysis of the human connectome data supports the notion of a “Common Model of Cognition” for human and human-like intelligence across domains

Andrea Stocco [a], Catherine Sibert [a], Zoe Steine-Hanson [b 1], Natalie Koh [c 2], John E. Laird [d], Christian J. Lebiere [e], Paul Rosenbloom [f]

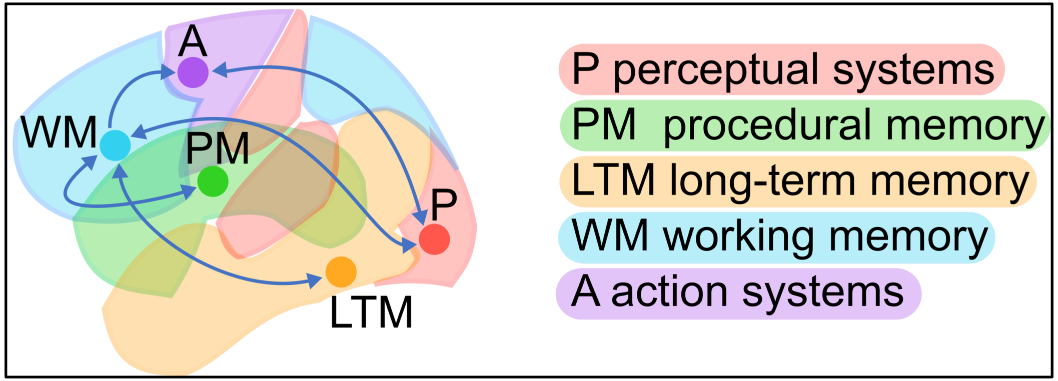

Stocco, Andrea, et al. "Analysis of the human connectome data supports the notion of a “Common Model of Cognition” for human and human-like intelligence across domains." NeuroImage 235 (2021).

Sibert, Catherine, Holly Sue Hake, and Andrea Stocco. "The structured mind at rest: low-frequency oscillations reflect interactive dynamics between spontaneous brain activity and a common architecture for task control." Frontiers in Neuroscience 16 (2022)

Kotseruba, Iuliia, and John K. Tsotsos. "40 years of cognitive architectures: core cognitive abilities and practical applications." *Artificial Intelligence Review* (2018): 1-78.

Cognitive LLMs
ACT-R: the Cognitive Architecture

Image adapted from: Ritter, F. E., Tehranchi, F., & Oury, J. D. (2019). ACT-R: A cognitive architecture for modeling cognition. Wiley Interdisciplinary Reviews: Cognitive Science, 10(3), Paper e1488. http://acs.ist.psu.edu/papers/ritterTOip.pdf

- “Architecture” refers to the fundamental organizational principle of a complex cognition system.
- Modules to implement the fixed mechanisms of cognition

Productions,
if-then rules,
actionable memory

Chunks,
factual memory,
slot-value pairs

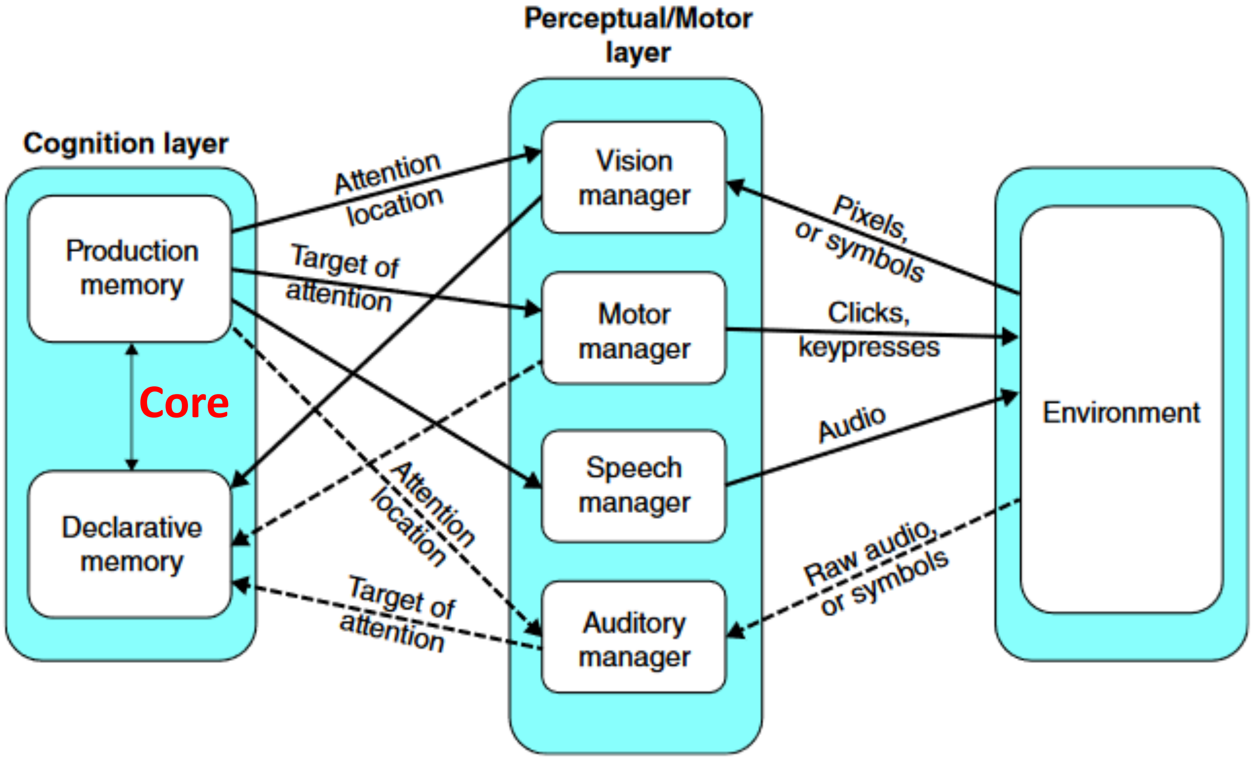

Cognitive LLMs
Approach

- Enhance the decision-making capabilities of LLMs by integrating intermediate semantic representations from Cognitive Architectures (CAs).
- Such semantic representations distilled from cognitive architectures serve as domain knowledge, infusing it into LLM decision augmentation.

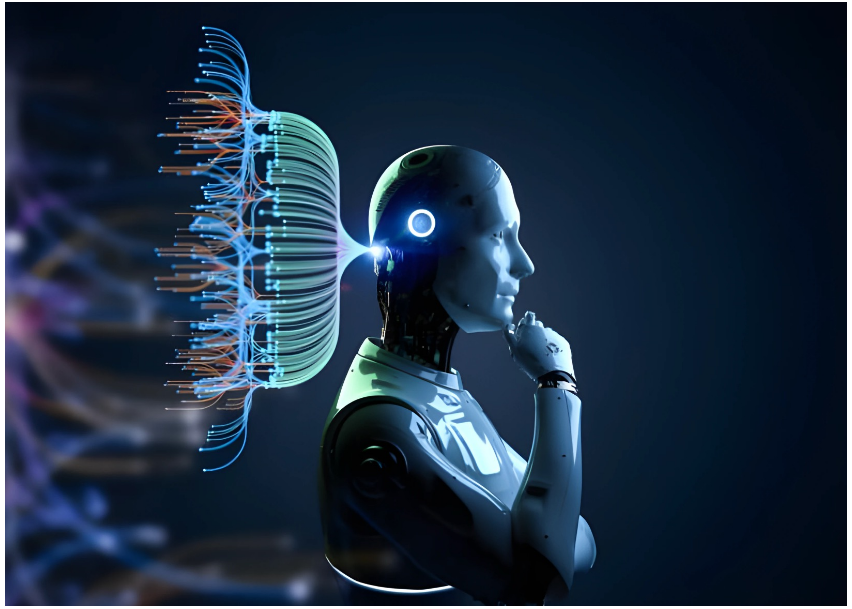

Cognitive LLMs

Solution

1

Exercise data + Documentation

BPS/CIP Handbook

Neuro-Symbolic Cognitive Method

2

Model

Tasks in exercise and CIP procedures from handbook are modeled with cognitive architecture

3

Simulate

Cognitive model is used to run stochastic simulations of task at scale (2,000 decision-trials)

4

Distill

Synthetic data are distilled semi-automatically from simulations and combined with prompt requests

5

Finetune

A finetuning pipeline is used to calibrate LLaMA-2 13B to perform decision augmentation for tasks in exercise

6

Q: "How to reduce production time, keeping HC stable?"

Pre-Assembly [data]+Assembly [data]

Cognitive LLM

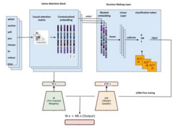

Recommendations

***A:** Emptying individual boxes with scrap material takes too much time. Recommendation is to load n empty boxes on cart and discard them in bulk"*

- A combination of different techniques like ontology-based formalization, NLP, psychometrics, are used to 1) model the symbolic components of the task, i.e., the declarative and procedural knowledge, and 2) set the sub-symbolic parameters (e.g., learning rate, similarity matching).
- Before running simulations at scale, tests are conducted on adequacy of symbolic and parametric representation

- Key semantic features are extracted from simulations and dataset for finetuning is prepared. Process is currently not fully-automatized

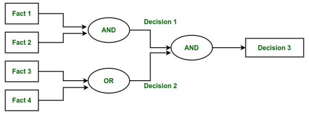

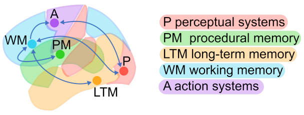

01

BPS handbook and documentation

Cognitive LLMs

BPS Handbook, Value Stream Map

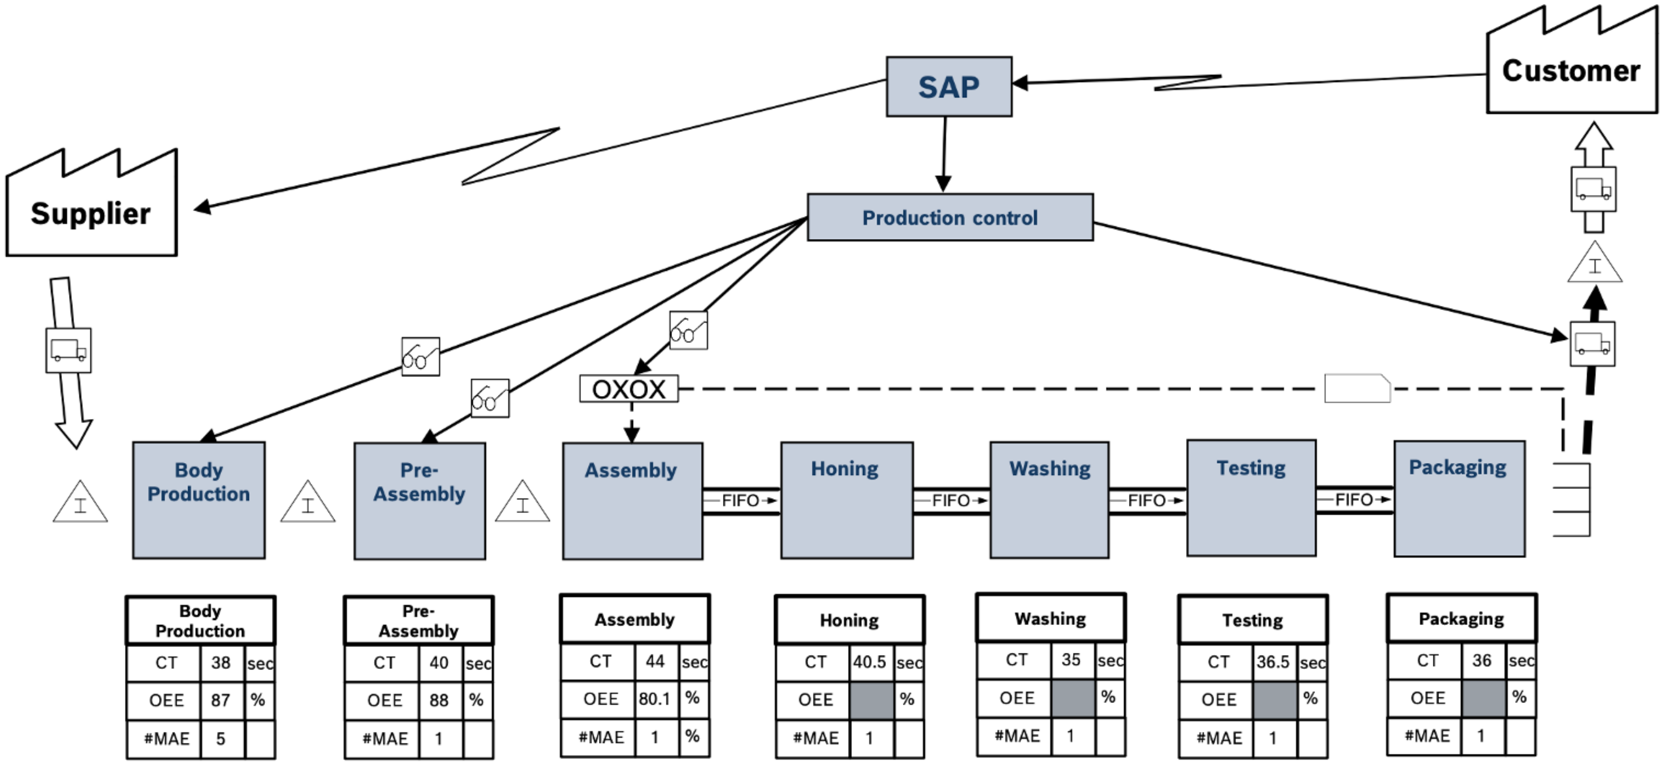

BOSCH

Cognitive LLMs
Documentation

- A simplified view of VSM and related intertwined variables

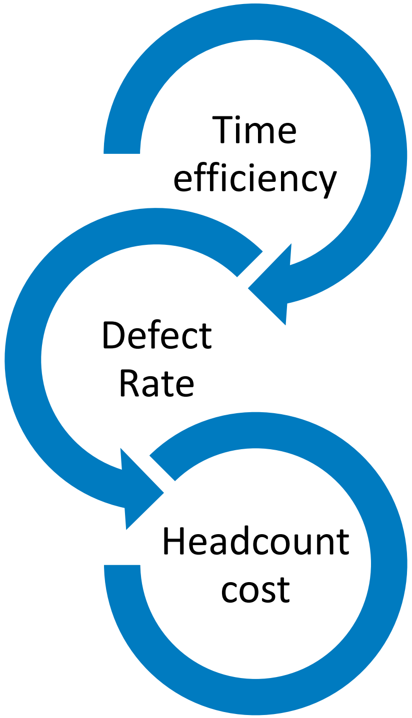

CT	OEE
MAE	Organizational Loss
Preassembly	Assembly

OEE: Overall Equipment Effectiveness

CT: Cycle Time

MAE: Mean Absolute Error

BOSCH

Cognitive LLMs
The Task

A production line consists of two sections with potential defect sources: pre-assembly and assembly

Pre-assembly takes 40 seconds with an Overall Equipment Effectiveness (OEE) rate of 88%

Assembly, on the other hand, takes 44 seconds but has a lower OEE rate of 80.1%

To achieve a total assembly time reduction of 4 seconds, we need to identify which section can be optimized with minimal impact on defect rate

It's important to note that reducing cycle time will also lead to an increase in **line headcount** cost

Cognitive LLM

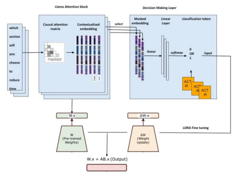

Recommendations

A: I recommend reducing the time spent in the assembly section based on the following expert rationale:

02

Model

Cognitive LLMs
VSM-ACT-R

VSM-ACT-R V1.0

- Rule-based model in decision making
- Integrating "personas" ranging from novice to intermediate and expert levels observed from human subjects

📘 The model learns over the course of trials and exhibits individual differences. It demonstrates a human-like learning progression, showing a steep learning curve at the beginning and gradual improvements later on

VSM-ACT-R V2.0

- Incorporate metacognitive processes of reflecting and evaluating the progress of the selected approach (headcount cost evaluation)
- Implementing a reinforcement-learning mechanism in a production-system framework, and simulating the reinforcement learning processes of decision-makers as they transition from novice to expert.

📘VSM-ACT-R: Toward Using Cognitive Architecture For Manufacturing Solutions. (June, 2024) Siyu Wu, Alessandro Oltramari, Frank E Ritter. Submitted to *17th International Conference on Social Computing, Behavioral-Cultural Modeling & Prediction and Behavior Representation in Modeling and Simulation (SBP-BRIMs)*

A penalty propagation

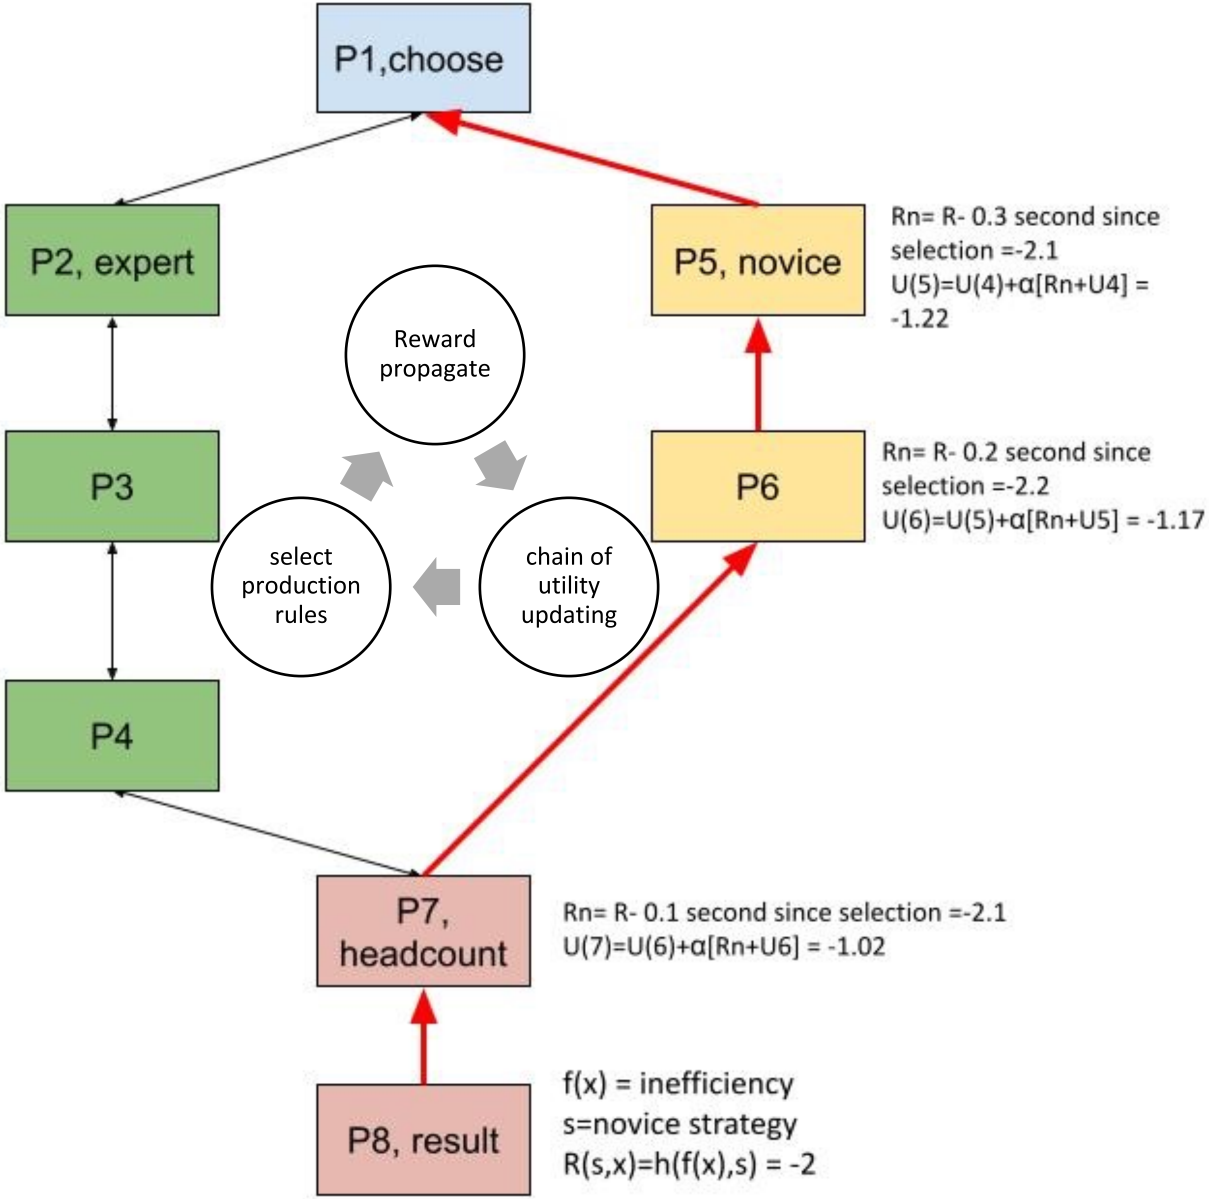

Level of Expertise Mechanism

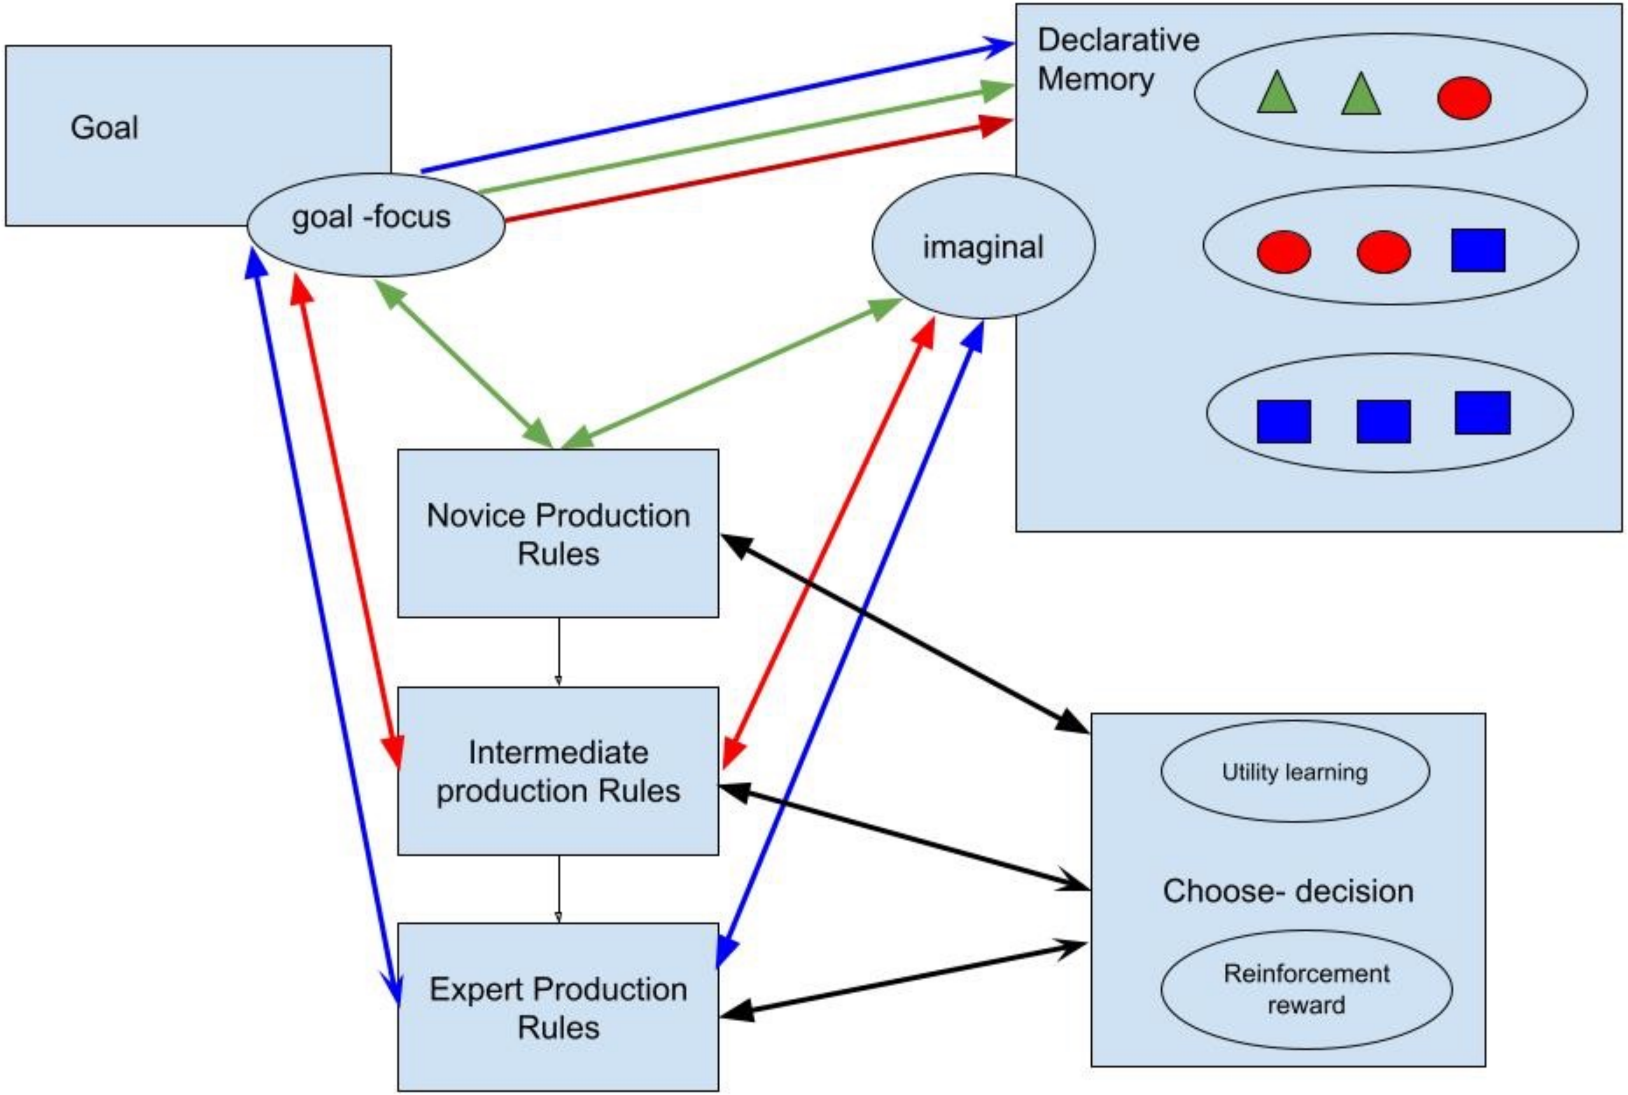

start with beginner strategy

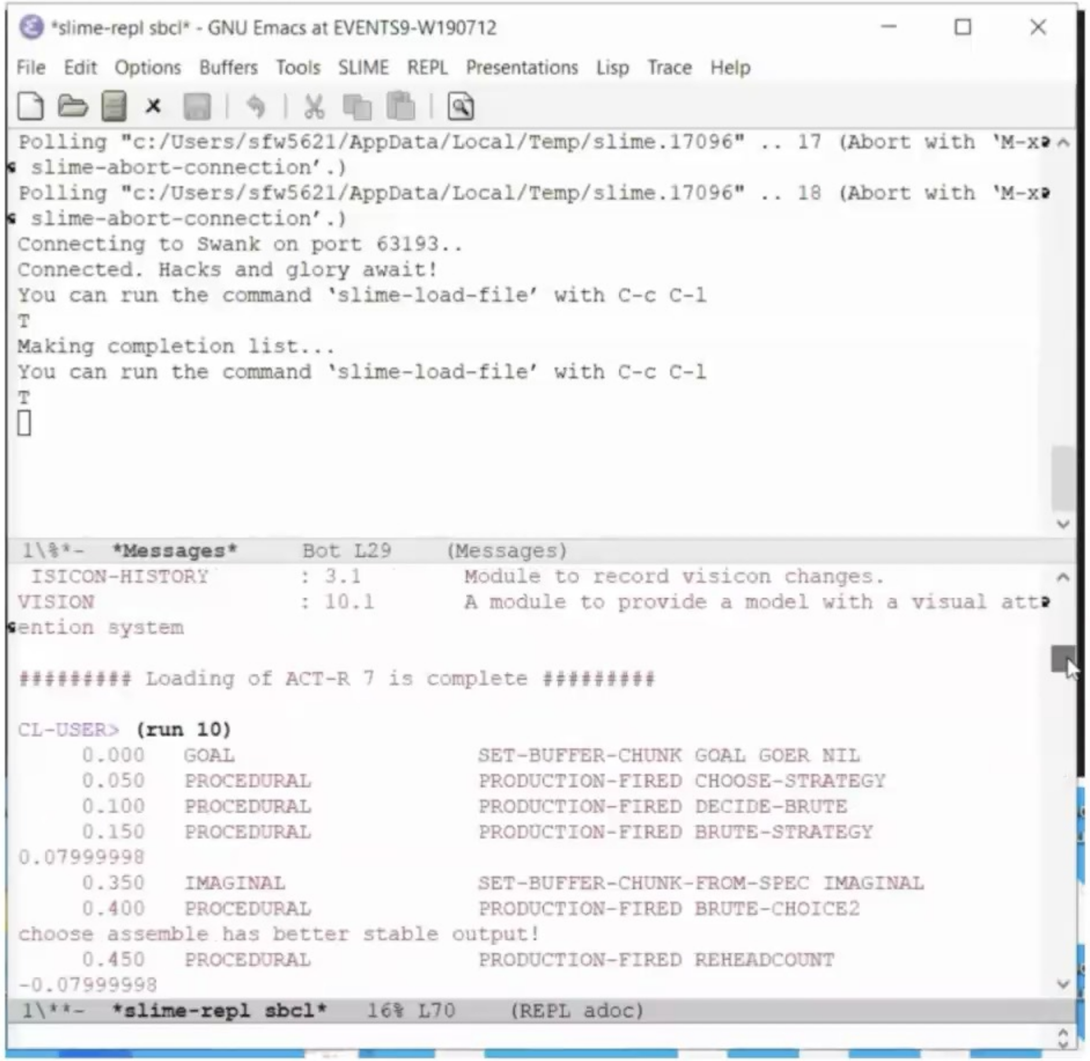

00:00:00:11 / 00:00:33:18

SCH

03

Simulate

BOSCH

Cognitive LLMs
Simulate

- Automate the acquisition of large quantities of ACT-R traces to satisfy data needs in LLM-related research. Previous efforts fine tune LLM conceptual embeddings to embed human behavior (Binz&Schulz,2023) have been limited by data collection costs. This simulation aims to address data scarcity in interdisciplinary psychological research

- 2,000 decision-making trials, obtained by running the developed ACT-R model across 32 problem sets, each consisting of 60-70 trials with 4 individual ACT-R personas.

Binz, M., & Schulz, E.(2023) Turning large language models into cognitive models, ICLR

04

Distill

Distill

Selecting Salient Decision Information

- **Challenges with Reserving Information:** Balancing the need to minimize information loss with reducing computational costs is essential.
- **Rationale for Selective Outputs:** Outputs from specific modules are chosen to distill macro-level cognitive processes related to executive functioning effectively.
- **Rationale for Preserving All Traces:** Preserving all traces involves processes of semantic embedding extraction and dimensional reduction to manage comprehensive data efficiently.

05

Finetuning

Whether LLMs can be informed with executive function knowledge about the reasoning processes of cognitive architecture

Cognitive LLMs

Fine Tuning for Knowledge Transfer

- **Loss Function**: Cross-Entropy.
- **Optimization**: Adam optimization.
- **Data Split:** train test 0.2
- **Batch Size**: batch size of 5 for both training and validation.
- **Learning Rate**: 1e-5.
- **Training Duration**: 10 epochs.
- **Regularization**:
  - Weight decay of 0.01.
  - Dropout rate of 0.5.
  - Gradient accumulation is set to 2.
- **Gradient explosion control**: Gradient clipping of 1.0.

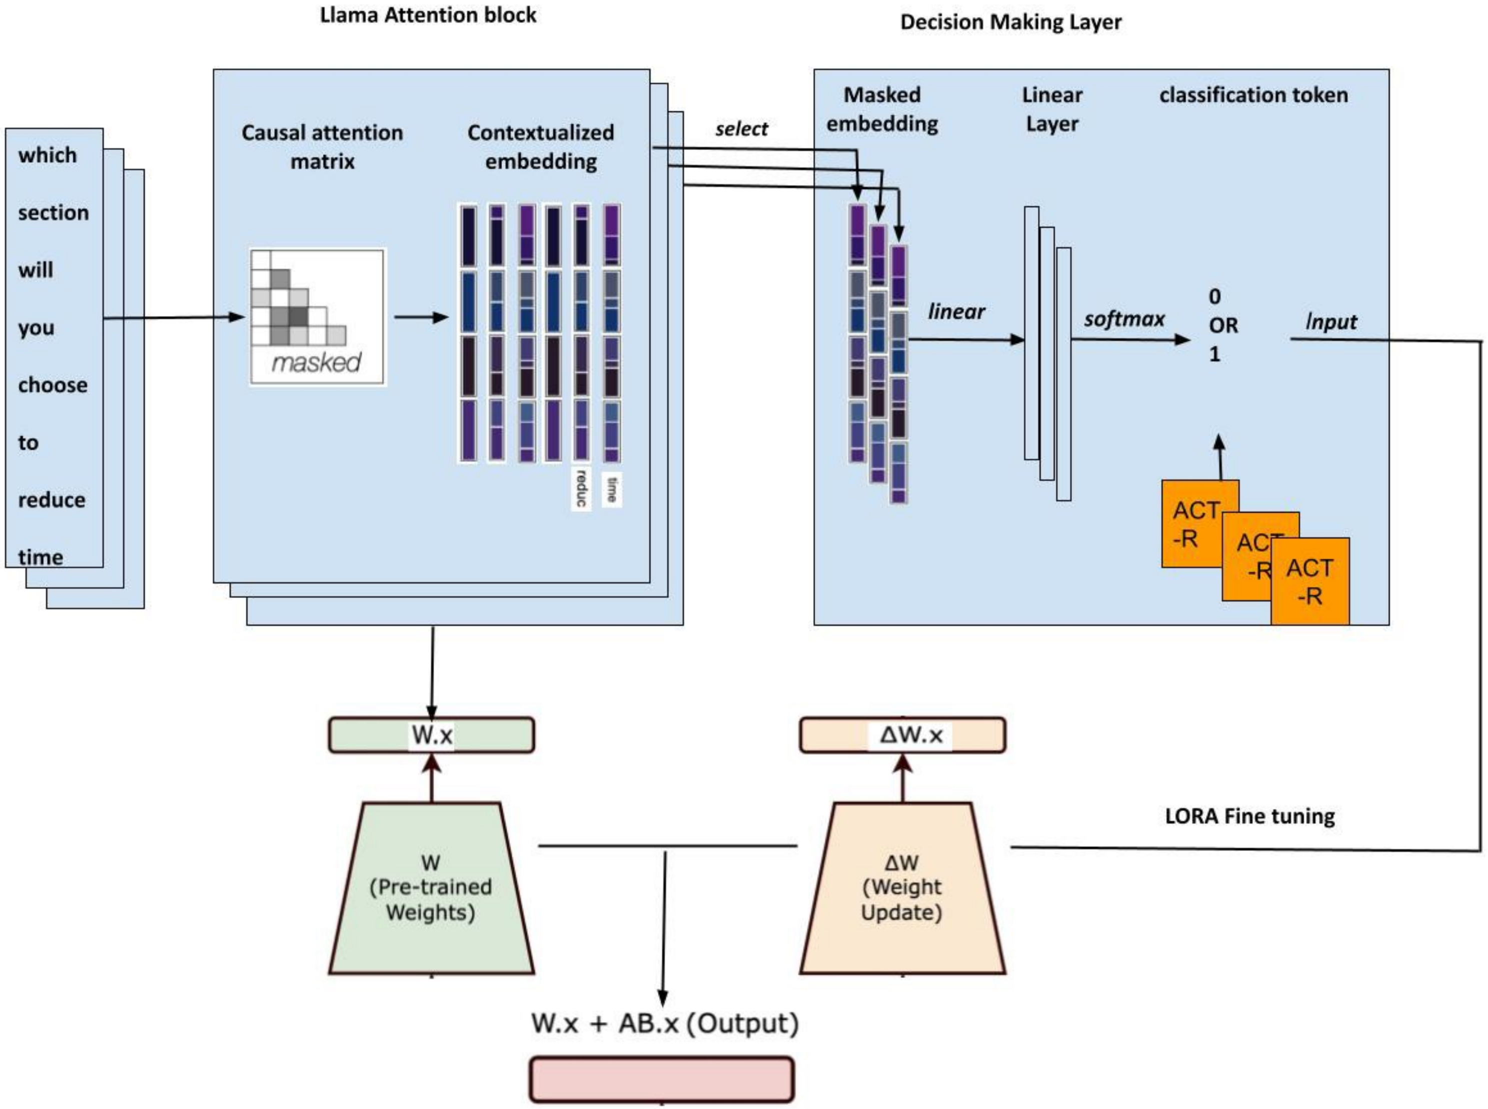

Internal | 2024-06-11

Cognitive LLMs
Results

Model	NLL	Accuracy
LLM-ACTR	0.6534	0.6576
LLAMA	0.7623	0.3564
chance-level	0.69	0.49

BOSCH

Cognitive LLMs
Conclusion

- The VSM-ACTR **cognitive model** exhibits improved performance over time, reflecting a human-like capability to generalize well on unseen data.
- The LLM-ACTR **model** demonstrates significantly better performance across all metrics compared to the LLaMA- only model.
- The LLAMA-only model performs worse than the chance-level model in both accuracy and NLL.
- These results underscore the necessity of fine- tuning pre-trained language models like LLaMA to adapt them to specific human-aligned repeated decision-making tasks.
- The integration of ACT-R-based decision-making processes with learning trajectories associated with LLMs significantly enhances the overall system's performance in capturing complex human-aligned decision-making patterns.

06

How to project the full semantic space of ACT-R to infused knowledge for creating a superior grounded cognitive knowledge base?

Cognitive LLMs

Approach: Automatic Distill for maximum semantic preservation

- A fully automated pipeline of feature extraction and reduction form ACTR cognitive reasoning process

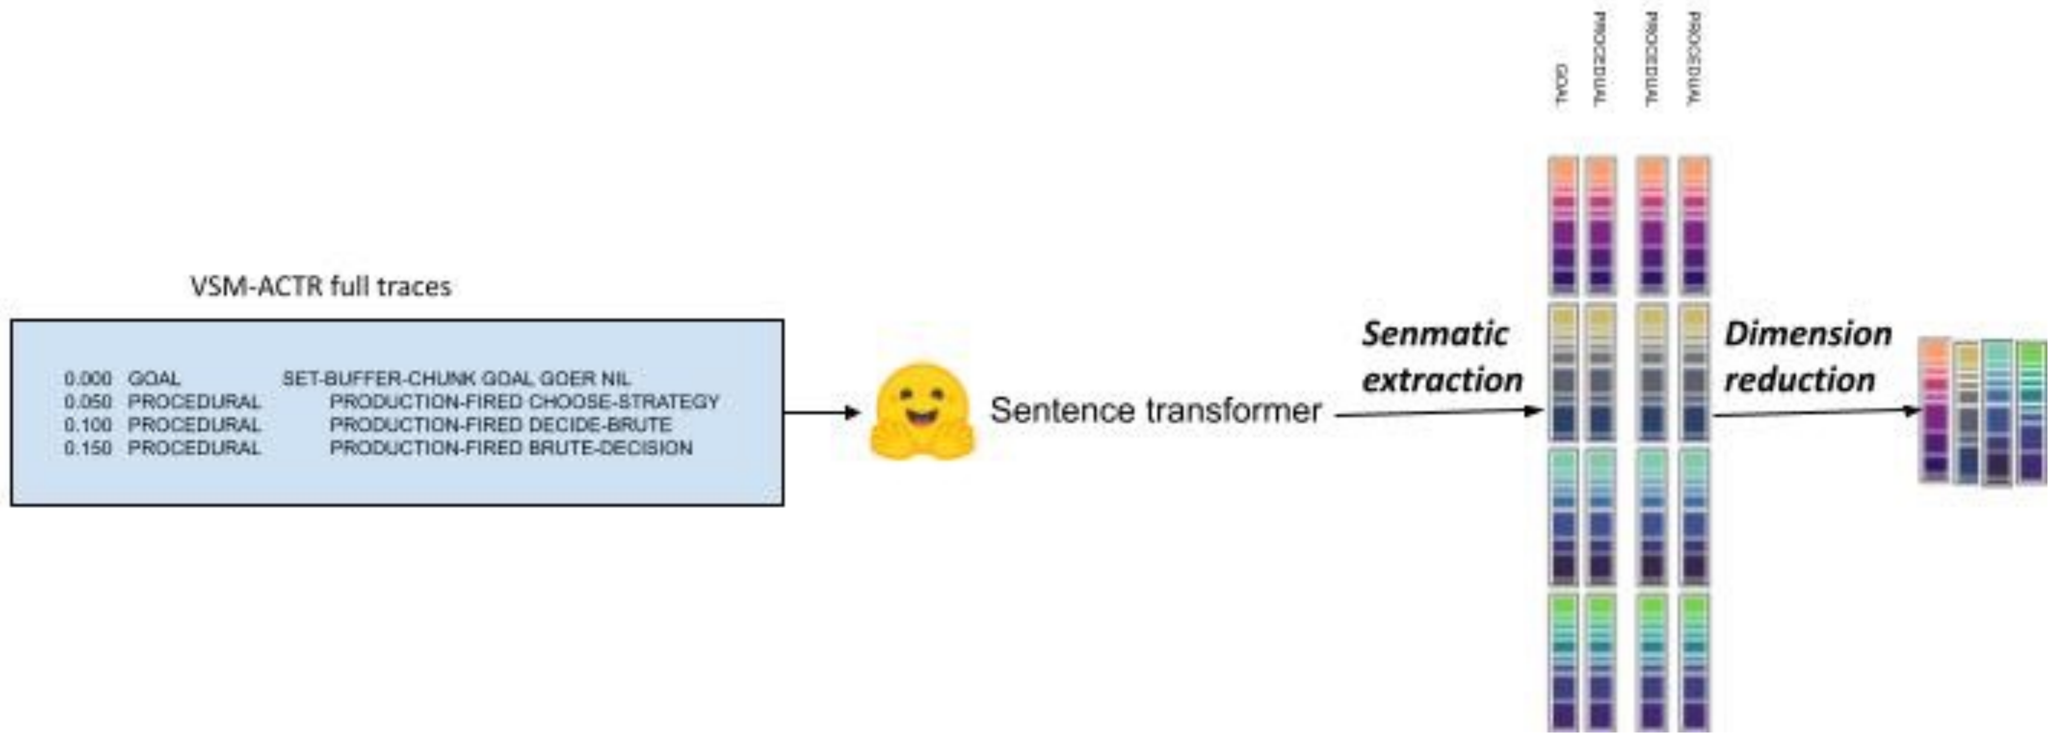

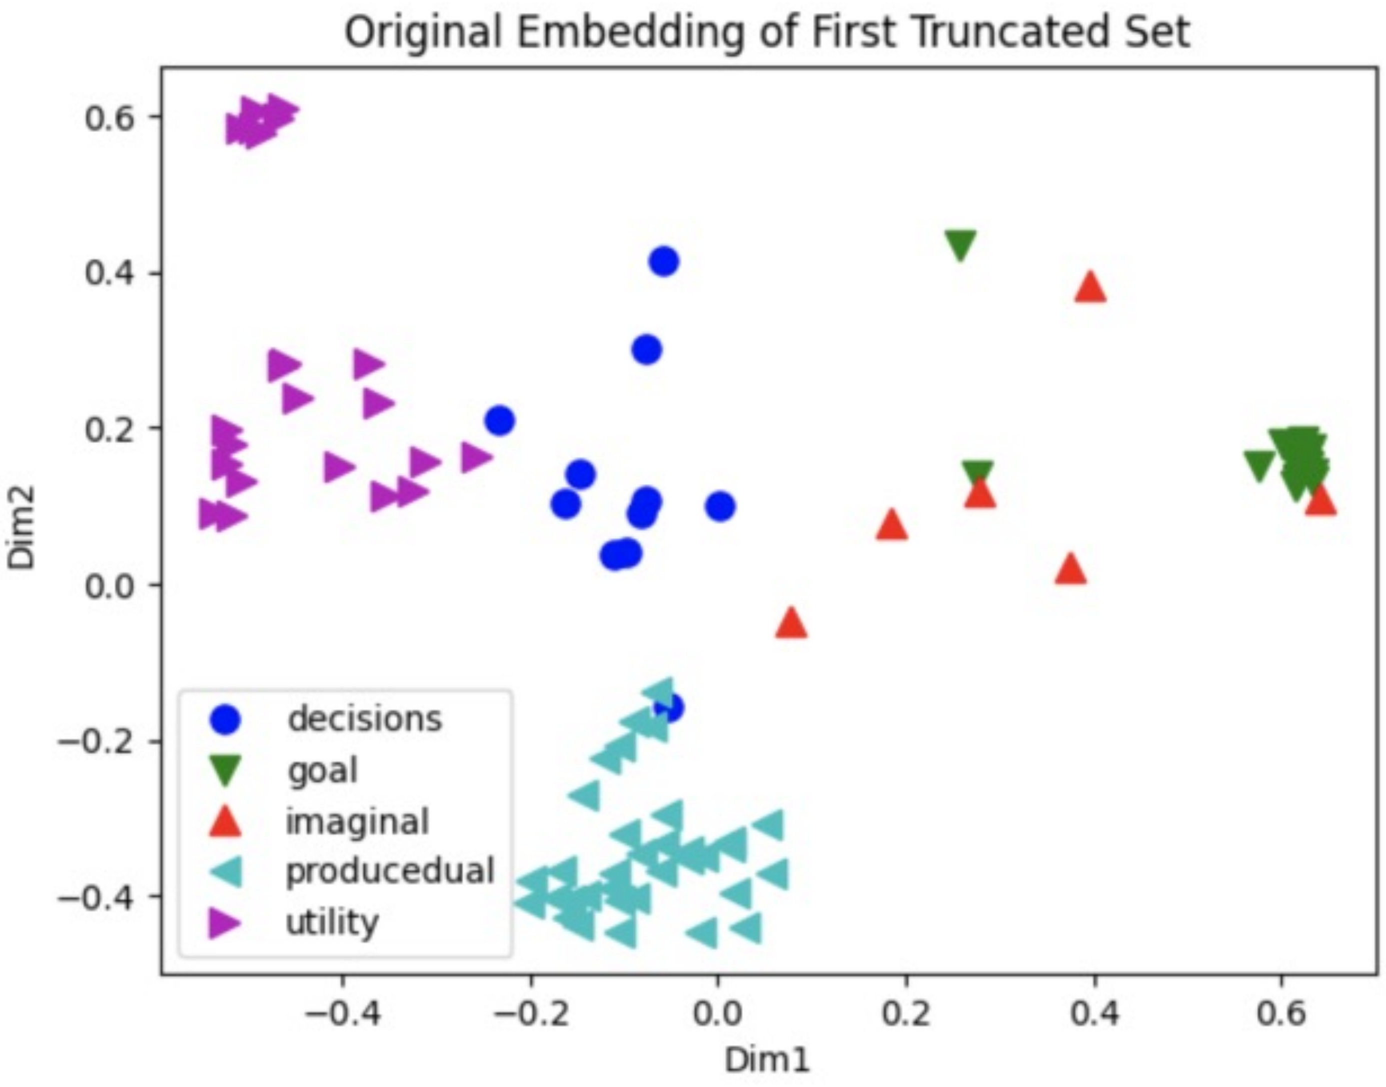
Original Embedding of First Truncated Set
0.6
0.4
0.2
0.0
−0.2
−0.4
Dim2
−0.4
−0.2
0.0
0.2
0.4
0.6
Dim1
decisions
goal
imaginal
producedual
utility

BOSCH

Cognitive LLMs
Implication and Future Direction

- Validates our embedding extraction and dimension reduction strategy of reserving ACT-R cognitive reasoning process.
- Emphasizes the potential of these reduced embeddings to augment foundation models in decision making by grounding them in a robust cognitive framework, while maintaining minimal computational cost.
- This development could open the door to infusing this grounded knowledge at scale into LLMs.

Reviewer’s comments and next step

“This paper describes a cognitive model of control of a manufacturing process at various levels of experience. It is a nice example of applying cognitive architectures to practical problems. ” –Reviewer 1

“Very interesting paper with actual implications which makes the paper more valuable. ” – Reviewer 2

“This work is interesting and tackles an applied problem. The model appears complete and demonstrates behavior that represents level of experience and individual differences.” – Reviewer 3

- In addition to AAAI, the current project report is also a good match for the upcoming NeurIPS. Also open to related conferences and workshops that are following the trend of enabling machine reasoning through neural symbolic approach.
- The future work from this project will be suitable for ICRL and ICML when the reinforcement reasoning process is well integrated through ACT-R feature preservation and injection.
- Journals like Neural Symbolic AI would be a good host of our project report as well.

BOSCH

Internal Acknowledgement

- Thank Kaushik, Yuxin, Marius, and my close intern friends who made the summer fun and enriching. Kaushik helped with the code, Marius encouraged me to automate the semantic data collection process, and Yuxin encouraged me to broaden cross-team collaboration.

08

Leftovers

Flow Chart For Expert Strategy

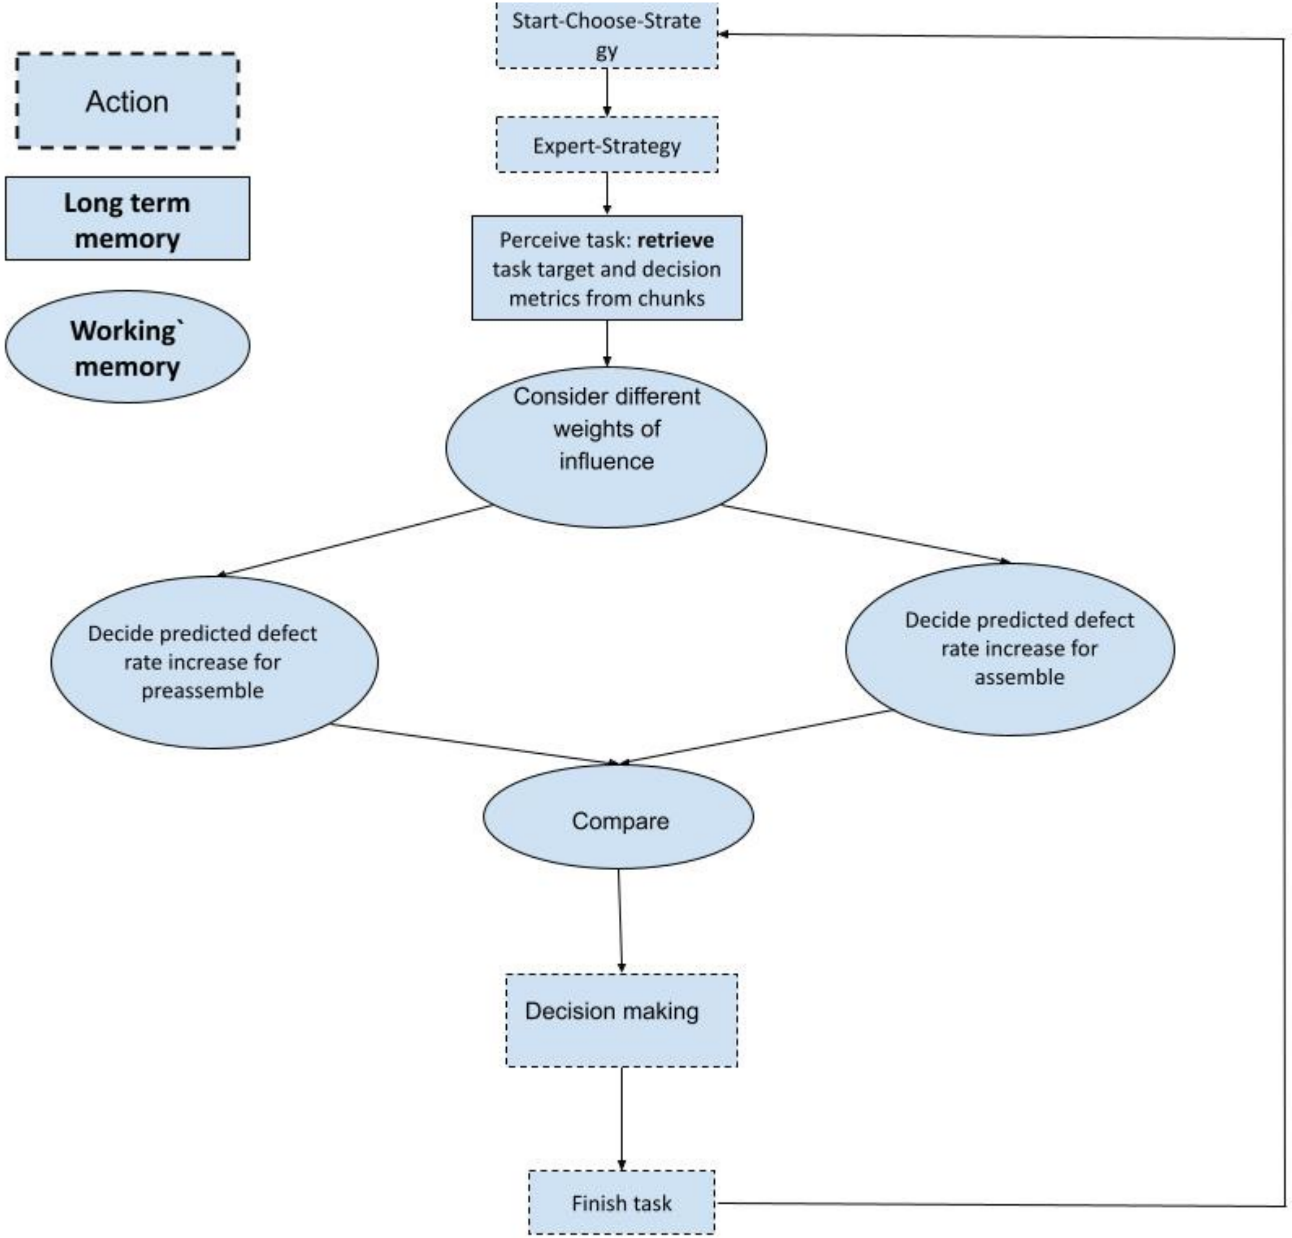

02

VSM-ACT-R 2.0

Reinforcement Learning of Basal Ganglia

- Similar to mapping between CMC components and homologous cortical and subcortical regions, the production system in ACT–R has been associated with the basal ganglia.
- For reinforcement stimuli, dopamine modulates striatal activity based on expected rewards or punishments, influencing decision-making and reinforcing specific neural patterns

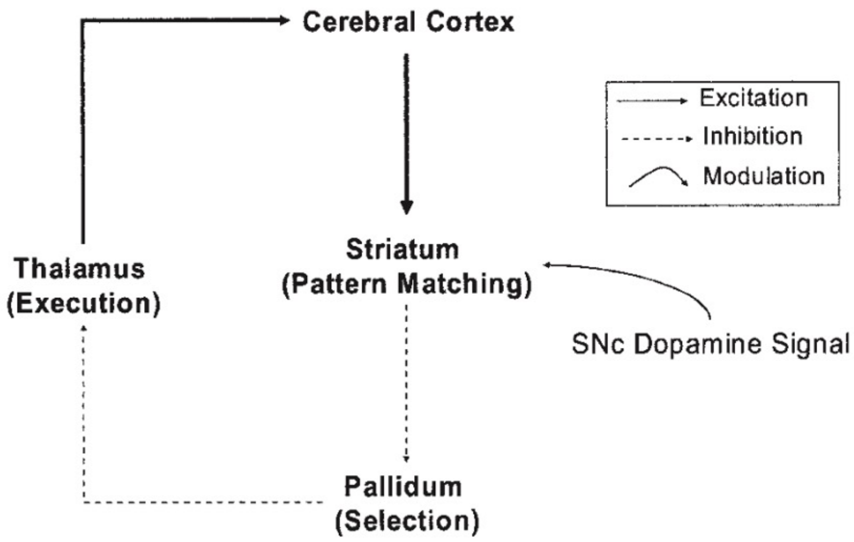

Figure 1. A simplified schematic diagram of the cortical–basal–ganglionic feedback loop. SNc = Substantia Nigra pars compacta.

Figure excerpt from: Fu, W. T., & Anderson,J. R. (2006). From recurrent choice to skill learning: a reinforcement-learning model. *Journal of Experimental Psychology: General*, *135*(2), 184.

Implementing a Reinforcement-Learning Mechanism in a Production-System Framework

- Temporal difference (TD) algorithm in ACT-R

$$U_i(n)=U_i(n-1)+\alpha[R_i(n)+U_i(n-1)]$$

- Each production has a utility (U_i)associated with it
- In which $U_i(n)$ represents the utility of a production after its nth occurrence and $R_i(n)$ represents the actual reinforcement (either a reward or a penalty) received on the nth occurrence.
- The parameter α (0 <a <1) controls the rate of learning.

Reward propagation

- One of the powerful features of the TD algorithm is its ability to propagate credit/reward back to previous productions. Reward or penalty can find its way back to a critical early production through a chain of productions that leads to the reward, it will then influence the utilities of the chain of productions.

```
"update-reward called" Strategy: NOVICE, State: INEFFICIENT
Updated Reward: -2
 Utility updates with Reward = -2.0   alpha = 0.2
  Updating utility of production CHOOSE-STRATEGY
   U(n-1) = -0.9   R(n) = -2.5 [-2.0 - 0.5 seconds since selection]
   U(n) = -1.22
  Updating utility of production DECIDE-BRUTE
   U(n-1) = 2.3179998   R(n) = -2.45 [-2.0 - 0.45 seconds since selection]
   U(n) = 1.3643999
  Updating utility of production BRUTE-STRATEGY
   U(n-1) = -0.8640001   R(n) = -2.4 [-2.0 - 0.4 seconds since selection]
   U(n) = -1.1712
  Updating utility of production BRUTE-CHOICE2
   U(n-1) = -0.77400005   R(n) = -2.15 [-2.0 - 0.15 seconds since selection]
   U(n) = -1.0492
  Updating utility of production REHEADCOUNT
   U(n-1) = -0.756   R(n) = -2.1 [-2.0 - 0.1 seconds since selection]
   U(n) = -1.0248
  Updating utility of production STOP
   U(n-1) = -0.73800004   R(n) = -2.05 [-2.0 - 0.05 seconds since selection]
   U(n) = -1.0004001
     1.550   PROCEDURAL             PRODUCTION-FIRED CHOOSE-STRATEGY
     1.600   PROCEDURAL             PRODUCTION-FIRED EXPERT-STRATEGY
     1.650   PROCEDURAL             PRODUCTION-FIRED PERCEIVE
     1.700   PROCEDURAL             PRODUCTION-FIRED PREASSEMBLE-WEIGHT
0.5
caculate the preassemble defect decision weight
     1.900   IMAGINAL               SET-BUFFER-CHUNK-FROM-SPEC IMAGINAL
     1.950   PROCEDURAL             PRODUCTION-FIRED ASSEMBLE-WEIGHT
0.5
```

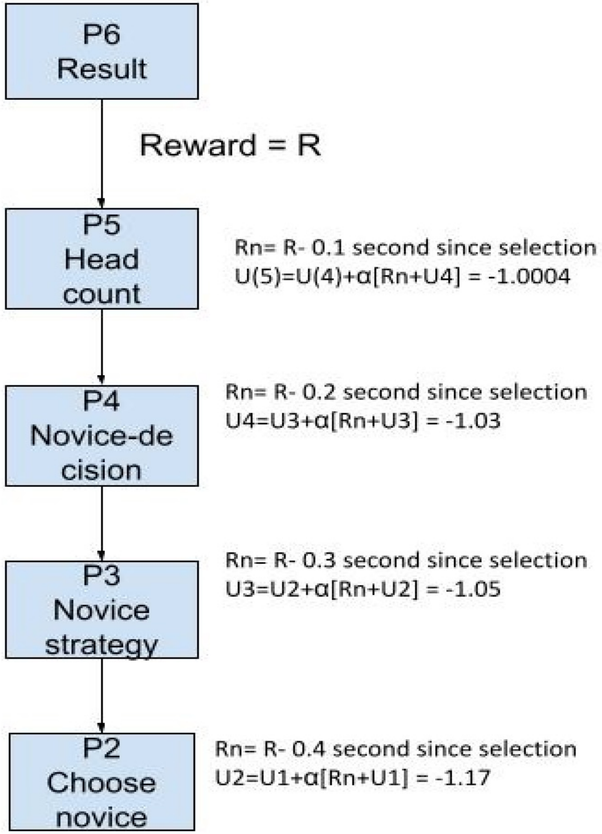

Selecting Productions on the Basis of Their Utilities

- If there are a number of productions competing with expected utility (expert or novice strategy) values *Uj* the probability of choosing production *i* is described by the formula

$$\Pr obability(i) = \frac{e^{U_i / \sqrt{2}s}}{\sum_j e^{U_j / \sqrt{2}s}}$$

- Where the summation *j* is over all the productions which currently have their conditions satisfied
- S is the noise.

A schema of reflecting and evaluating the progress of the selected approach

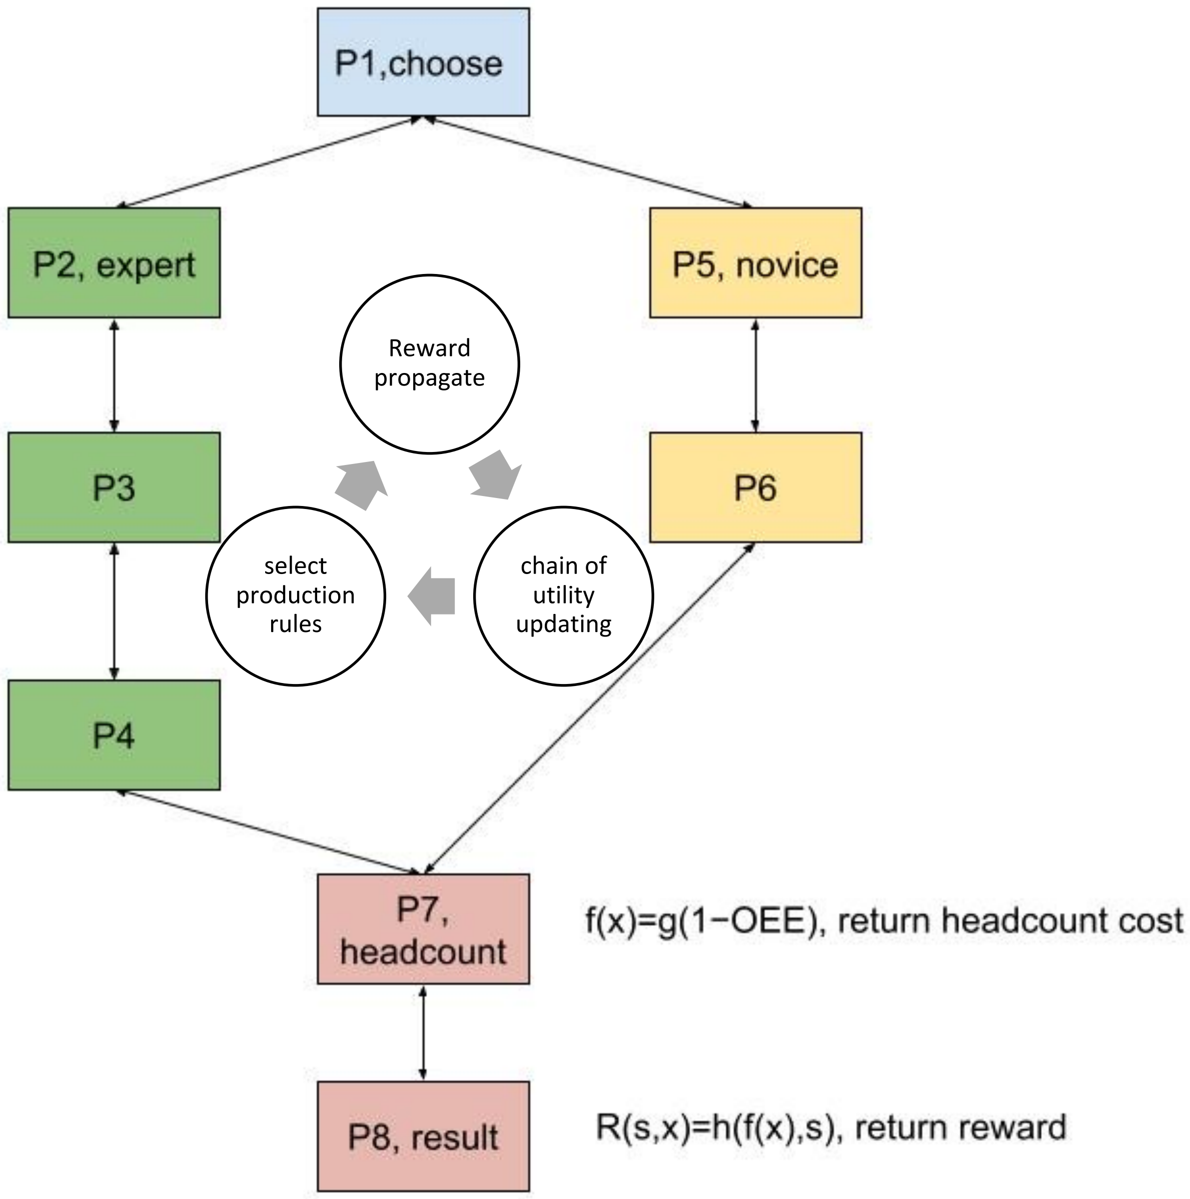

A reward propagate for expert and headcount cost efficient decision

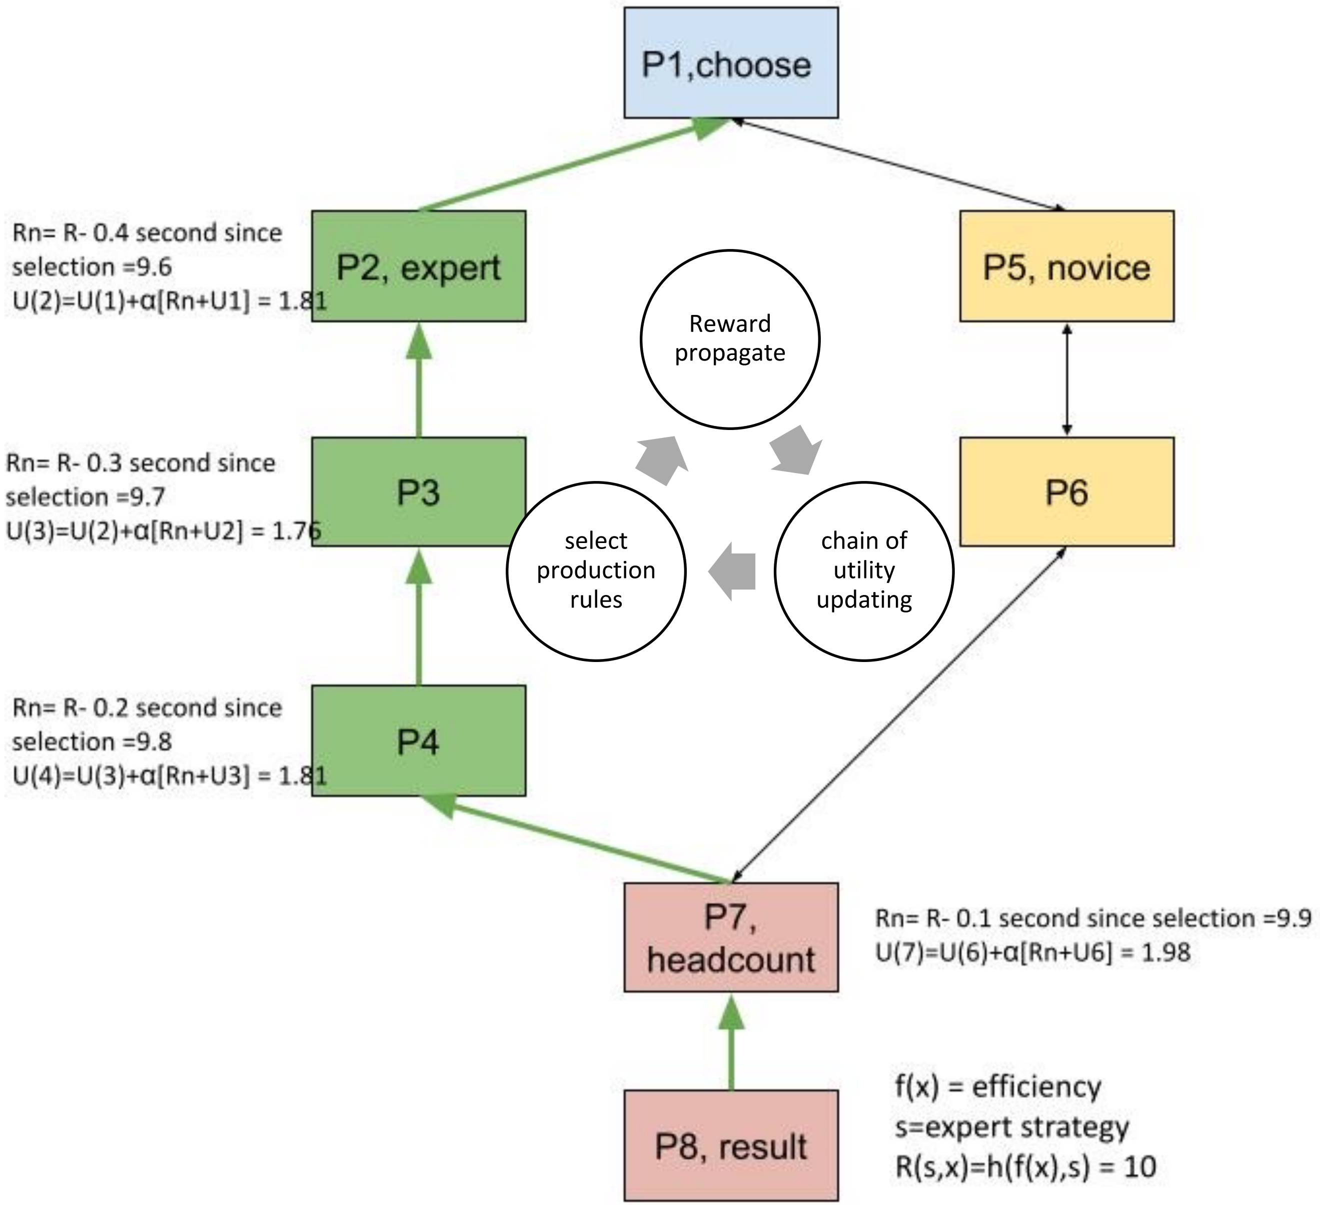

BOSCH

Next Problem

When using a language model as a tool to assist in this domain-specific decision-making, which presents two issues:

- The responses are often incorrect and lack contextual relevance.
- There is a lack of human behavior resonance; the answers to calls across different prompts are not related to each other. The language model is stateless and does not show any traces of learning or improvement across calls.

Solution

- Integrating Llama (open-source large language model) with VSM-ACT-R through downstream fine-tuning will allow the language model to make decisions on the same domain questions as the VSM-ACT-R model.

Fine tuning Llama using LORA

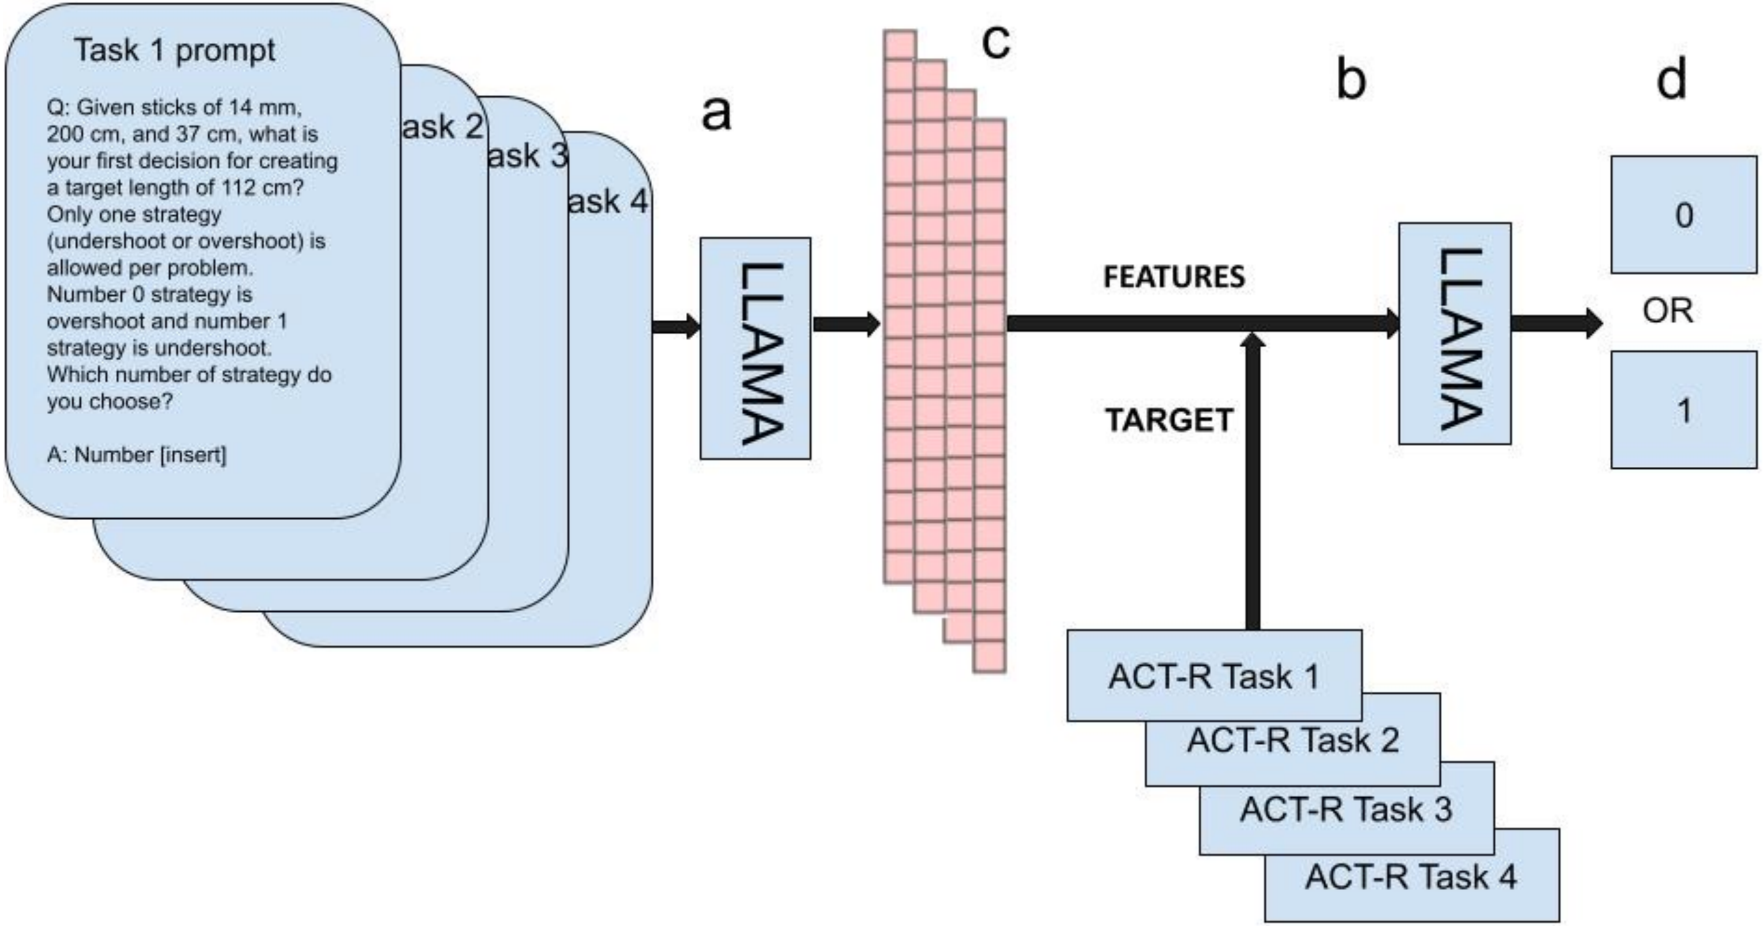

Internal | 2024-06-11