

Open-endedness and intrinsic motivation in embodied virtual agents: A Systematic Literature Review

Mircea Lică¹,
Ojas Shirekar¹,
Chirag Raman¹

Motivations

Humans are able to develop a wide repertoire of skills and knowledge in all sorts of environments.

Current generation of virtual agents like DreamerV3 [1] are **incapable of generalization** and require **human reward modelling**.

By lacking autonomy and generalization capabilities, these agents **fail to adapt to real world scenarios** and perform **out-of-distribution tasks**.

Scope of the survey

Conduct a **review of the literature** on both open-ended learning and intrinsic motivation and how these have been applied in the context of virtual agents.

Question 1. How has open-ended learning and intrinsic motivation been applied in the context of embodied virtual agents?

Question 2. What kind of benchmarks are used to assess an agent's open-endedness and intrinsic motivation capabilities?

Question 3. What are the limitations of the current methods in the field and what are the possible directions for future research?

We focus our attention towards the ability of open-ended learning and intrinsic motivation to overcome the limitations of current systems.

Methodology

To ensure the survey contains the most significant and relevant research in the field, while maintaining clarity in the selection process, we follow the **Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA)** framework.

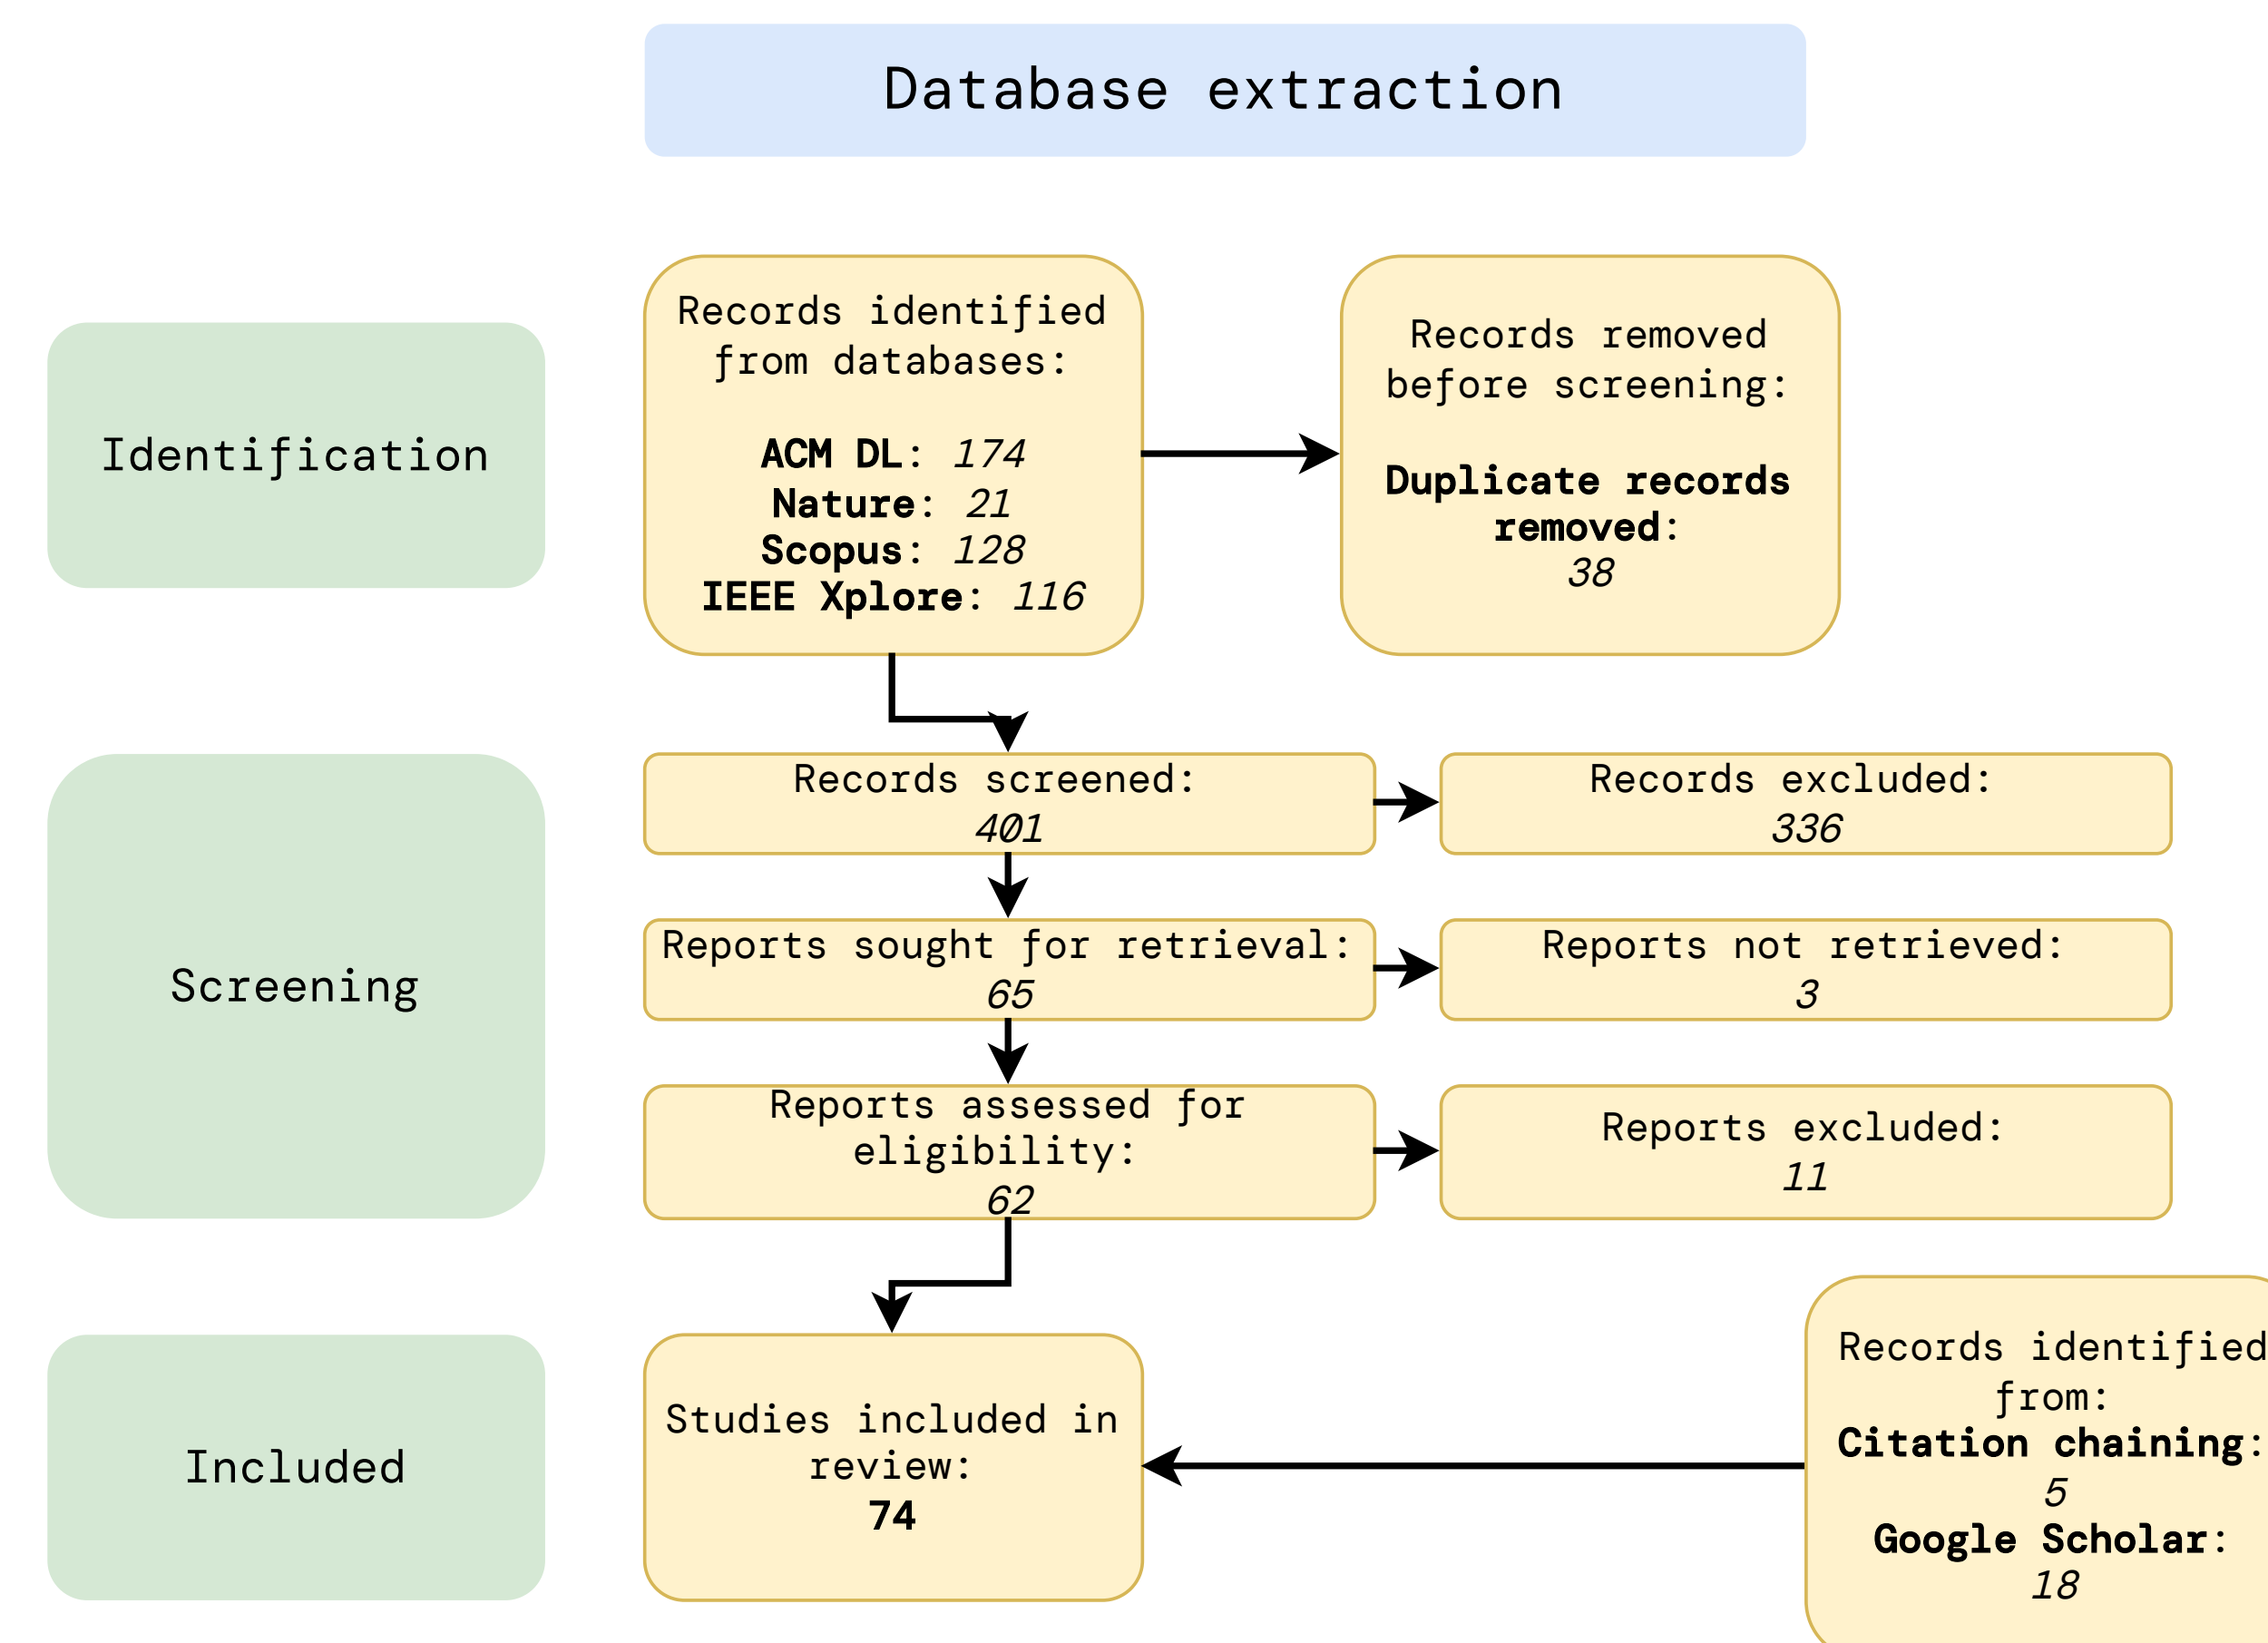


Figure 1: Lorem ipsum dolor sit amet, consectetur adipiscing elit. Praesent porttitor arcu luctus, imperdiet urna iaculis, mattis eros. Pellentesque iaculis odio vel nisl ullamcorper, nec faucibus ipsum molestie.

In addition, we enhance the results with articles found through Google Scholar and citation chaining.

Findings

Finding 1. Social interaction and cooperation provide a mechanism of recursive collaborative improvement, enhancing adaptability and efficiency.

Finding 2. Agent's embodiment choice directly influences its synergy with the environment, and possibly hampering its adaptability.

Finding 3. Deploying open-ended agents in the real world, requires the development of more sample efficient learning algorithms.

Finding 4. Both open-ended learning and intrinsic motivation research suffer from the lack of standardized benchmarks.

Findings 1 - 4 highlight the most important conclusions drawn from the reviewed literature. **Finding 5** underlines a more general insight, which highlights the importance of open-ended learning and intrinsic motivation.

Finding 5. Open-endedness and intrinsic motivation are necessary for emergence of generally capable agents.

- Shrink the gap between simulation and reality by leveraging large vision models, together with other modalities like language. [3]
- Tackle the data bottleneck by leveraging world simulators [4].
- Designing a set of standardized benchmarks and environments.

Open-endedness and intrinsic motivations are necessary for emergence of generally capable agents.

Affiliations

¹ Department of Computer Science, Delft University of Technology

References

- ¹ Hafner et al (2020, October). Mastering Atari with Discrete World Models. ICLR.
- ² Wang et al. (2024). RL-VLM-F: Reinforcement Learning from Vision Language Foundation Model Feedback. arXiv preprint arXiv:2402.03681
- ³ Kim et al (2024). OpenVLA: An Open-Source Vision-Language-Action Model. arXiv preprint arXiv:2406.09246.
- ⁴ Bruce et al (2024). Genie: Generative Interactive Environments. arXiv preprint arXiv:2402.15391.

Promising Research Directions

- Using cultural transmission and social learning for enhancing overall capabilities of the agents.
- Use the general knowledge of large pretrained models to guide the agent through preferences [2].