

Peer Review

Review of: "Physical AI Agents: Integrating Cognitive Intelligence with Real-World Action"

Peter Denno¹

1. National Institute of Standards and Technology, Gaithersburg, United States

The paper describes an architecture and is mostly aspirational. Some thoughts on how it could be made better:

- * Figure 1: I'd expect Decisions to affect the Environment. I'd expect Task to be about changing the Environment.

- * Section 2.3 seems too repetitive of Section 1.1.

- * Section 3.3, 4. Manufacturing

Robots should not be making decisions about the weld just by looking at the parts to join.

It should take into account the required strength of the weld and other requirements. This was decided at design time. Manufacturing makes one or more instances of the design.

- * Figure 3: I don't think "Industry-specific LLM" is quite right. Maybe "Environment-specific LLM" But what is either of those, really? I think if you have in mind making an LLM specific by loading its vector store, you ought to say that.

- * Figure 3: In your examples, you have robots acting in the physical world, yet in this diagram, there is simply "Query" and "Response". Maybe that "Query" is the expression of a goal?

- * Some discussion of prior work about robots using LLMs would be helpful.

Declarations

Potential competing interests: No potential competing interests to declare.