

# Cosmos-Reason1: From Physical Common Sense To Embodied Reasoning

NVIDIA<sup>1</sup>

## Abstract

Physical AI systems need to perceive, understand, and perform complex actions in the physical world. In this paper, we present the Cosmos-Reason1 models that can understand the physical world and generate appropriate embodied decisions (e.g., next step action) in natural language through long chain-of-thought reasoning processes. We begin by defining key capabilities for Physical AI reasoning, with a focus on physical common sense and embodied reasoning. To represent physical common sense, we use a hierarchical ontology that captures fundamental knowledge about space, time, and physics. For embodied reasoning, we rely on a two-dimensional ontology that generalizes across different physical embodiments. Building on these capabilities, we develop two multimodal large language models, Cosmos-Reason1-7B and Cosmos-Reason1-56B. We curate data and train our models in two stages: Physical AI supervised fine-tuning (SFT) and Physical AI reinforcement learning (RL). To evaluate our models, we build comprehensive benchmarks for physical common sense and embodied reasoning according to our ontologies. Evaluation results show that Physical AI SFT and RL bring significant improvements. To facilitate the development of Physical AI, we make our code and pre-trained models available under the NVIDIA Open Model License at <https://github.com/nvidia-cosmos/cosmos-reason1>.

## 1. Introduction

Physical AI systems are designed to interact with the physical world. To effectively follow instructions and take appropriate actions to achieve a desired goal, they need first to perceive, understand, and reason about the physical world. Recently, with breakthroughs of post-training and test-time scaling via long chain-of-thought (CoT) processes, large language models (LLMs) have demonstrated remarkable general reasoning capabilities in tackling complex problems across domains such as coding and mathematics (OpenAI, 2024; DeepSeek-AI, 2025). However, a key limitation of these models lies in their ability to ground their knowledge in the physical world. While LLMs trained on vast amounts of textual data from the Internet may acquire essential knowledge to reason about the physical world, they often struggle to establish connections between that knowledge and real-world interactions and dynamics.

In this paper, we begin by defining the fundamental capabilities essential for Physical AI systems. Unlike designing models that excel at solving coding and mathematical problems, our focus is on empowering models with physical common sense knowledge and embodied reasoning capabilities grounded in the real world. To establish a shared framework and help measure progress, we propose two ontologies. First, a hierarchical ontology that organizes physical common sense into three major categories — Space, Time, and Fundamental Physics — which are further divided into 16 fine-grained subcategories. This ontology encapsulates knowledge about how the physical world operates under the laws of physics and how it responds to interactions with embodied agents. Second, we introduce a two-dimensional ontology for embodied reasoning, which encompasses four key reasoning capabilities across five types of embodied agents. Effective embodied reasoning based on physical common sense is crucial for understanding and planning actions that achieve specific objectives in the real world. The details are described in Sec. 2.

---

<sup>1</sup>A detailed list of contributors and acknowledgments can be found in App. C of this paper.

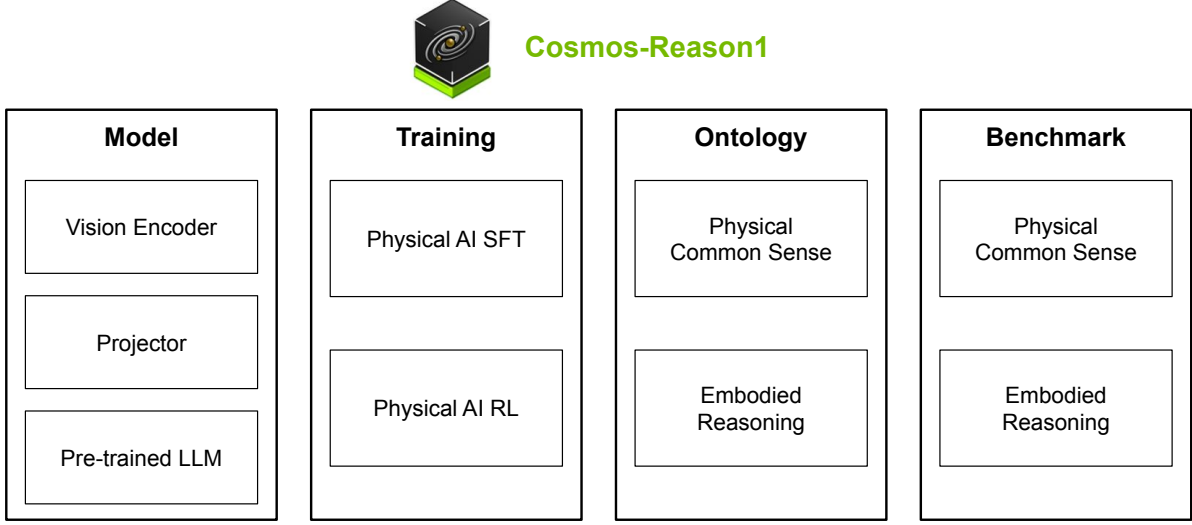


Figure 1: An overview of Cosmos-Reason1. Cosmos-Reason1 contains two multimodal large language models of 7B and 56B, trained in two stages, including Physical AI SFT and Physical AI RL. We also define two ontologies for physical common sense and embodied reasoning, and build two benchmarks to evaluate models’ Physical AI reasoning capabilities.

We introduce *Cosmos-Reason1* as a step toward enabling multimodal LLMs to generate more physically grounded responses. We focus on the visual world, where the observations of the world are represented as videos. Cosmos-Reason1 perceives the physical world through video input, understands it, and reasons about it through long chain-of-thought thinking processes before generating responses. These responses, expressed in natural language, include both explanatory insights and embodied decisions, such as determining the next action to take. We employ a decoder-only multimodal LLM architecture where the input video is processed by a vision encoder followed by a projector to align with the text token embeddings before feeding into the LLM. We experiment with both dense Transformer and hybrid Mamba-MLP-Transformer architecture as the LLM backbone. Cosmos-Reason1 comes with two model sizes: Cosmos-Reason1-7B and Cosmos-Reason1-56B. We describe details about the model architecture in [Sec. 3](#).

Constructing rule-based, verifiable rewards at scale has been critical for the reasoning LLMs’ success in solving math and coding problems. Can we design rule-based, verifiable rewards for training Physical AI reasoning models with reinforcement learning? In this work, we explore two types of rewards based on answering multiple-choice questions (MCQs). The first type of MCQs are designed based on human annotations. With inspiration from video self-supervised learning, we automatically generate a second type of MCQ based on the structure of video data itself, such as solving puzzles with shuffled spatiotemporal video patches or predicting the arrow of time on whether a video is playing forward or backward. All these rewards are rule-based, verifiable, and highly relevant to Physical AI capabilities. We discuss details about the design of RL training algorithm and infrastructure in [Sec. 4](#).

Data determines our model’s ceiling. To strengthen our models’ physical common sense and embodied reasoning capabilities, we carefully curate around 4M annotations of video and text pairs, including captions for physical understanding, multiple choice questions, and long chain-of-thought reasoning traces. We develop two pipelines to curate these physical common sense and embodied reasoning data according to our ontologies. The data are curated based on human annotations and model distillation from DeepSeek-R1 ([DeepSeek-AI, 2025](#)) for Physical AI supervised fine-tuning. Details about data are discussed in [Sec. 5](#).

In order to evaluate our models, we build new benchmarks for evaluating Physical AI capabilities in [Sec. 6](#). For physical common sense in [Sec. 6.1](#), we build three benchmarks (Space, Time, and Fundamental Physics) containing 604 questions from 426 videos. For embodied reasoning in [Sec. 6.2](#), we build six benchmarks

containing 610 questions from 600 videos, covering a wide range of tasks across different physical embodiments, including humans, robot arms, humanoid robots, and autonomous vehicles.

**Sec. 7** presents the evaluation results of Cosmos-Reason1 and comparisons with existing models. In **Sec. 7.1**, we introduce the experiment setup, including training details for Physical AI SFT, and evaluation results of SFT models on our benchmarks. In **Sec. 7.2**, we present the evaluation results for Physical AI RL. Training using our rule-based, verifiable rewards in RL post-training leads to improvements in all of our benchmarks.

**Fig. 1** shows an overview of Cosmos-Reason1. In summary, we introduce two multimodal large language models, Cosmos-Reason1-7B and Cosmos-Reason1-56B. The models are trained in two stages: Physical AI SFT and Physical AI RL. We define ontologies for physical common sense and embodied reasoning and build benchmarks to evaluate models’ Physical AI reasoning capabilities. To facilitate the advancement of Physical AI, we make our code and pre-trained models available under the NVIDIA Open Model License at <https://github.com/nvidia-cosmos/cosmos-reason1>. Building reasoning models for Physical AI is an open problem that is far from being solved, and we hope our paper contributes to the advancement of this field.

## 2. Physical AI Reasoning

We identify two important capabilities for Physical AI reasoning models — physical common sense reasoning and embodied reasoning. First, Physical AI models should possess physical common sense, meaning a general, embodiment-agnostic understanding of the environment and forms the basis for predicting what is plausible and implausible in the real world. Second, Physical AI models should also help embodied agents perceive, reason, and make decisions about planning future interactions with the physical environment. We seek to incorporate both “System 1” and “System 2” in physical common sense reasoning and embodied reasoning. “System 1” enables fast, intuitive responses such as pattern recognition and instinctive judgments, while “System 2” operates more slowly, engaging in deliberate reasoning for complex decision-making (Kahneman, 2011).

### 2.1. Common Sense Reasoning

Humans acquire physical common sense primarily through passive observation of the world. For example, infants can understand basic concepts such as object permanence and gravity in a few months after birth (Riochet et al., 2021). This common sense encompasses a collection of knowledge about what is possible, impossible, or likely to happen in the real world. Training AI systems in real-world environments is expensive and can pose risks to both the system and its surroundings. Utilizing physical common sense, AI systems can quickly learn new skills with minimal trial and error while avoiding making critical mistakes in uncertain scenarios (LeCun, 2022).

To define physical common sense, we introduce an ontology comprising three broad categories: Space, Time, and other Fundamental Physics, further divided into 16 fine-grained subcategories. Inspired by Morris et al. (2024), we focus on capabilities rather than processes. Specifically, our ontology identifies key capabilities that Physical AI models should possess, without specifying the mechanisms or embodiments by which a system accomplishes tasks. For example, we believe the ability to understand the spatial relationship of objects, the temporal order of events, and object permanence are fundamental to Physical AI. However, such systems need not necessarily act human-like, such as grasping with dexterous hands with fingers or walking on two legs.

We show our physical common sense ontology in **Fig. 2**. The Space category encompasses the relationships between objects, their interactions, and the surrounding environment. It includes concepts such as Relationship, Plausibility, Accordance, and Environment. The Time category pertains to actions and events that unfold over a duration, covering Actions, Order, Causality, Camera, and Planning. Lastly, we introduce a Fundamental Physics category to address objects and core physical principles, including Attributes, States, Object Permanence, Mechanics, Electromagnetism, Thermodynamics, and Anti-Physics. Detailed definitions of all subcategories are described in **Tab. 1**.




Figure 2: A pie chart showing our physical common sense ontology. The ontology has three categories (Space, Time, and Fundamental Physics) and 16 fine-grained subcategories.

## 2.2. Embodied Reasoning

Physical AI operates in the real world, where environments are dynamic, uncertain, and governed by complex physical interactions. Unlike *abstract reasoning* in mathematics and programming, which manipulates symbols in a structured and deterministic manner, embodied reasoning requires an AI system to interact with and learn from the physical world. Unlike *passive understanding*, reasoning in embodied AI is often grounded in action, enabling robots not only to comprehend what they currently observe but also plan intelligent behaviors for the future in uncertain and dynamic environments. Specifically, embodied reasoning requires the capability to:

1. **Process Complex Sensory Inputs.** Unlike symbolic reasoning, which works with clean data representations, embodied reasoning must extract meaningful patterns from raw, often incomplete, and ambiguous sensory inputs.
2. **Predict Action Effects.** Actions have physical consequences, and effective reasoning requires an intuitive grasp of cause-and-effect relationships. An AI system must predict how an object will respond to force, how a robot’s body will interact with its surroundings, or how a vehicle’s movement will be influenced by terrain and physics.
3. **Respect Physical Constraints.** Unlike abstract problem-solving, which often involves optimizing discrete choices, embodied reasoning must account for real-world physics, such as inertia, friction, and material properties. It requires AI to generate long-horizon action plans that are feasible given physical constraints, ensuring stability, efficiency, and safety in execution.
4. **Learn from Interaction.** In Physical AI, actions do not occur in isolation; every movement or decision affects the environment and generates feedback. Embodied reasoning must continuously update its understanding based on these interactions, allowing the system to refine its behavior dynamically.

Embodied reasoning is also not confined to a single type of agent — it is applicable to humans, animals, and robots across various forms (*e.g.*, robotic arms, humanoid figures, or autonomous vehicles). They all need to develop similar embodied reasoning skills to navigate, manipulate, and make adaptive decisions under different environmental conditions and task goals. We summarize capabilities and types of physical embodiments with examples into a two-dimensional ontology in [Tab. 2](#).



Table 1: Definition of Physical AI systems’ capabilities for each subcategory in our common sense ontology.

Category: Subcategory	Capability
<b>Space: Relationship</b>	Determine the spatial relationship of objects in a scene. Perspective is important; for example, an object is left to the person or left to the camera view.
<b>Space: Plausibility</b>	Determine if a possible spatial relationship is feasible.
<b>Space: Affordance</b>	Understand object interaction with subjects such as humans, animals, robots, <i>etc.</i>
<b>Space: Environment</b>	Understand the scene or the surrounding environment.
<b>Time: Actions</b>	Understand actions, including the ability to accurately describe the action (movement, direction, intensity, <i>etc.</i> ), determine action objective, subtask or goal decomposition, and determine if a task/objective is successfully completed.
<b>Time: Order</b>	Understand the timestamp and sequential order of events.
<b>Time: Causality</b>	Understand if event A causes B.
<b>Time: Camera</b>	Determine the position and movement of the camera, including camera movement, camera angle/position, and transition of scenes.
<b>Time: Planning</b>	Come up with a future plan based on past observations.
<b>Fundamental Physics: Attributes</b>	Determine physical properties of an object, including semantic description, size, color, material, mass, temperature, solidity (can objects pass through one another?), <i>etc.</i>
<b>Fundamental Physics: States</b>	Determine the object state and understand the state change ( <i>e.g.</i> , ice changed to water, eggs changed from raw to cooked).
<b>Fundamental Physics: Object Permanence</b>	Understand object permanence, which properties can/cannot change in certain conditions (weight, shape, size, color, <i>etc.</i> ).
<b>Fundamental Physics: Mechanics</b>	Understand laws of physics related to Mechanics, including Statics (balance, stability, support, elasticity, deformation, the center of mass, <i>etc.</i> ), Kinematics (velocity, acceleration, linear motion, circular motion, rotational motion, <i>etc.</i> ), and Dynamics (gravity, collision, friction, sliding, inertia, conservation of momentum, fluids and particles, <i>etc.</i> ).
<b>Fundamental Physics: Electromagnetism</b>	Understand laws of physics related to Electromagnetism, including Optics (lighting, shadow, occlusion, reflection, refraction, diffraction, absorption, transmission, <i>etc.</i> ), Electricity, and Magnetism.
<b>Fundamental Physics: Thermodynamics</b>	Understand laws of physics related to Thermodynamics, such as heat, temperature change, evaporation, heat transfer, thermal expansion and contraction, <i>etc.</i>
<b>Fundamental Physics: Anti-Physics</b>	Understand situations that defy the laws of physics, such as anti-gravity, reverse of time, perpetual motion, sudden disappearance, <i>etc.</i>

In this paper, we focus on the first three embodied reasoning capabilities we defined above and leave “Learn from Interactions” as future work. Specifically, we focus on video input as a representative example of “Process Complex Sensory Inputs”. For “Predict Action Effects”, we focus on two tasks, including task-completion verification for determining whether a task has been completed, and the next plausible action prediction for predicting the next plausible next action to achieve a goal. For “Respect Physical Constraints”, we focus on action affordance to assess whether it is possible to perform a specific action toward a goal. We collect videos across different agents, including humans, robot arms, humanoid robots, and autonomous vehicles. By investigating these varied cases, we aim to deepen our understanding of how embodied reasoning enables intelligent interaction with the physical world.

### 3. Cosmos-Reason1

Cosmos-Reason1 is a family of multimodal large language models specialized in Physical AI reasoning. The family comprises two models: Cosmos-Reason1-7B and Cosmos-Reason1-56B. In this section, we introduce the design of our multimodal architectures and the choices of LLM backbones.

Table 2: Embodied reasoning ontology, with an example for each combination of capability and agent type.

	Natural Agents (humans, animals)	Robotics Systems (robot arms, humanoid robots, autonomous vehicles)
<b>Process Complex Sensory Inputs</b>	A person watches videos about a cooking recipe. A bat locates prey using echolocation.	A robot arm recognizes objects using its camera. A robot detects obstacles while walking. A self-driving car recognizes a stop sign and pedestrians.
<b>Predict Action Effects</b>	A carpenter anticipates wood splintering before cutting. A dog estimates a ball’s landing spot to catch it.	A robotic arm compensates for momentum before gripping an object. A robot estimates an object’s weight before lifting it. A self-driving car predicts tire slippage on ice.
<b>Respect Physical Constraints</b>	A pilot maintains altitude within aerodynamic limits. A cheetah limits speed to avoid muscle strain.	A robotic gripper limits its force to prevent breaking objects. A robot adjusts joint torque to prevent falls. A drone avoids exceeding wind resistance thresholds.
<b>Learn from Interactions</b>	A golfer corrects their stance after observing ball trajectory. A dog learns to open doors through repeated attempts.	A factory robot improves alignment after detecting misplacements. A robot learns new handshakes. A self-driving car refines braking distances.

### 3.1. Multimodal Architecture

There are different architecture designs for building multimodal large language models (LLMs) using existing text-only LLM backbones and vision encoders. Commonly used architectures are the decoder-only architecture (e.g., LLaVA (Liu et al., 2023)) and the cross-attention-based architecture (e.g., Flamingo (Alayrac et al., 2022) and Llama 3-V (Grattafiori et al., 2024)). We utilize the decoder-only architecture similar to LLaVA (Liu et al., 2023) and NVLM-D (Dai et al., 2024) for its simplicity and unified handling of all modalities by aligning other modality tokens (image or video) into the text token embedding space. Specifically, the model begins with the vision encoder (Chen et al., 2024), followed by the projector containing a downsampling two-layer MLP, and then the decoder-only LLM backbone (Nvidia et al., 2024; Waleffe et al., 2024; DeepSeek-AI, 2025).

For Cosmos-Reason1-7B, we choose Qwen2.5-VL (Bai et al., 2025) as our pre-trained model and follow the same image and video processing. For Cosmos-Reason1-56B, we leverage InternViT-300M-V2.5 (Chen et al., 2024) as our vision encoder and Nemotron-H (NVIDIA, 2025) as our LLM backbone. We perform the following processing for the Cosmos-Reason1-56B hybrid model. For each input image, we dynamically adjust the image to a predefined aspect ratio and segment it into 1 to 12 tiles, each measuring  $448 \times 448$  pixels, depending on the image’s resolution. Additionally, we generate a thumbnail tile; a scaled-down version of the full image to preserve the global context. More details can be found in Dai et al. (2024). For each input video, we uniformly sample up to 32 frames at a maximum rate of 2 frames per second, resizing each frame to  $448 \times 448$  pixels. For each  $448 \times 448$  video frame input, the vision encoder generates 1,024 visual tokens with the patch size of  $14 \times 14$ , which are then downsampled by a factor of  $2 \times 2$  using PixelShuffle (Shi et al., 2016), reducing them to 256 tokens by transforming spatial dimensions into channel dimensions. The image tokens from multiple tiles are concatenated with interleaved tile ID tags, as described in Dai et al. (2024), while the video tokens from multiple frames are concatenated directly. More discussions about the hybrid LLM backbone are in Sec. 3.2. We illustrate our hybrid multimodal architecture in Fig. 3 and summarize our model configurations in Tab. 3.

### 3.2. Hybrid Mamba-MLP-Transformer Backbone

Since its introduction, the Transformer architecture (Vaswani et al., 2017) has revolutionized the field of language modeling, becoming the de facto standard for building foundation models. However, its self-attention mechanism has a quadratic time complexity with respect to its context length. In contrast, the recently proposed Mamba architecture (Gu and Dao, 2023) introduces linear-time sequence modeling with selective state space

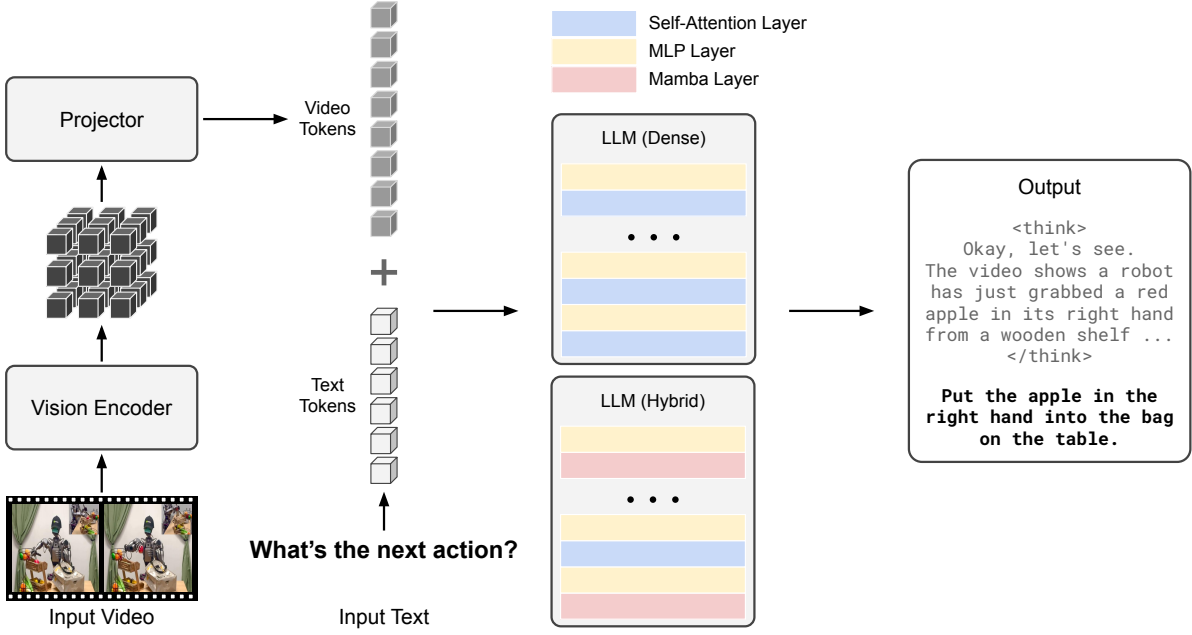


Figure 3: An illustration of our multimodal large language model architecture. Given an input video and an input text prompt, the video is projected into the LLM’s token embedding space as video tokens by a vision encoder followed by a projector. The text tokens are concatenated with the video tokens and fed into the LLM backbone, a dense Transformer or a hybrid Mamba-MLP-Transformer architecture. Our model can output responses with long chain-of-thought reasoning processes.

Table 3: Configuration details of Cosmos-Reason1 models.

Configuration	Cosmos-Reason1-7B	Cosmos-Reason1-56B
Vision Encoder		
Architecture	ViT-676M	ViT-300M
Input Size	Dynamic	$448 \times 448$
Patch Size	$14 \times 14$	$14 \times 14$
Number of Layers	32	24
Model Dimension	1,280	1,024
FFN Hidden Dimension	3,456	4,096
Projector		
Downsampling (HxWxT)	$2 \times 2 \times 2$	$2 \times 2 \times 1$
Number of Layers	2	2
Input Dimension	1,280	4,096
Hidden Dimension	5,120	32,768
Output Dimension	3,584	8,192
LLM Backbone		
Architecture	Transformer	Mamba-MLP-Transformer
Number of Layers	28	118
Model Dimension	3,584	8,192
FFN Hidden Dimension	18,944	32,768
Number of Attention Heads	28	64

models, making it significantly more efficient for handling long sequences. In practice, the selective state spaces of Mamba may not be sufficient to capture every detail within long sequences. To address this, a small portion of Transformer layers is incorporated for long-context modeling, giving rise to the hybrid Mamba-MLP-Transformer architecture (Waleffe et al., 2024).

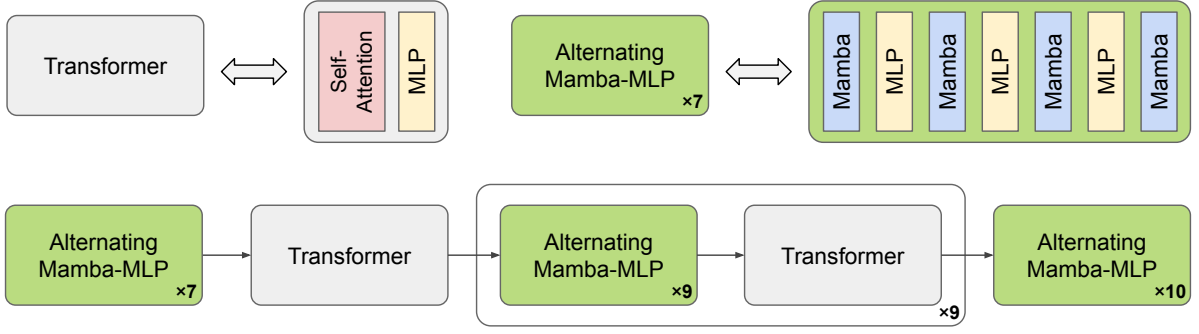


Figure 4: An illustration of our hybrid Mamba-MLP-Transformer backbone architecture used in Cosmos-Reason1-56B. A Transformer block consists of a self-attention layer and an MLP layer. We also show an example of an Alternating Mamba-MLP module on top of the figure.

In Cosmos-Reason1-56B, we use a pre-trained model with a hybrid Mamba-MLP-Transformer architecture (NVIDIA, 2025) as our backbone. An illustration of the 56B LLM architectures can be found in Fig. 4. We train the Cosmos-Reason1-7B dense model with a Tensor Parallelism of 4 (TP=4) (Shoeybi et al., 2019), while the Cosmos-Reason1-56B hybrid model is trained with a Tensor Parallelism of 8 and a Pipeline Parallelism of 2 (TP=8, PP=2).

## 4. Reinforcement Learning

We employ two training stages, Physical AI SFT and Physical AI RL, to adapt a pre-trained visual language model (e.g., Qwen2.5-VL (Bai et al., 2025) or Nemotron-H-VLM (NVIDIA, 2025)) to a Physical AI reasoning model. After the SFT stage, we post-train our models using RL with Physical AI focused tasks to further enhance their physics common sense and embodied reasoning abilities. In this section, we explain the algorithm and the customized training framework, which highly efficient and robust for RL training.

### 4.1. Algorithm

We adopt GRPO (Shao et al., 2024) as our RL algorithm of choice for its simplicity and computational efficiency, as it circumvents the necessity of training and maintaining a separate critic model. GRPO enables a streamlined approach to policy optimization, wherein the advantage function is derived by normalizing rewards within a cohort of responses generated for each prompt. Let  $R(o_i)$  denotes the reward for response  $o_i$ , in a group of responses  $\mathcal{G} = \{o_1, o_2, \dots, o_G\}$ , then the computed advantage can be expressed as:

$$A_i = \frac{R(o_i) - \text{mean}(\mathcal{G})}{\text{std}(\mathcal{G})}$$

### 4.2. Training Framework

To make more efficient use of RL training data, we also propose a novel, fully asynchronous and highly robust RL training framework. The architecture is shown in Fig. 5. First, unlike mainstream colocated frameworks (Hu et al., 2024; Sheng et al., 2024) that suffer from inefficient resource utilization due to synchronization overhead, the proposed framework adopts a heterogeneous deployment strategy for policy training and actor rollout. By leveraging a unified dispatcher to schedule and dispatch training prompts, our design enables asynchronous parallelism across the entire training pipeline. This achieves end-to-end asynchrony while maintaining training effectiveness, resulting in approximately a 160% improvement in training efficiency compared with the colocated frameworks. Second, attributed to our novel designed training mesh management logic, the framework is able to rapidly reconfigure itself and continue the current training step even if any node fails during training, without requiring costly restarts or recovery procedures. Furthermore, the dispatcher redundancy mechanism

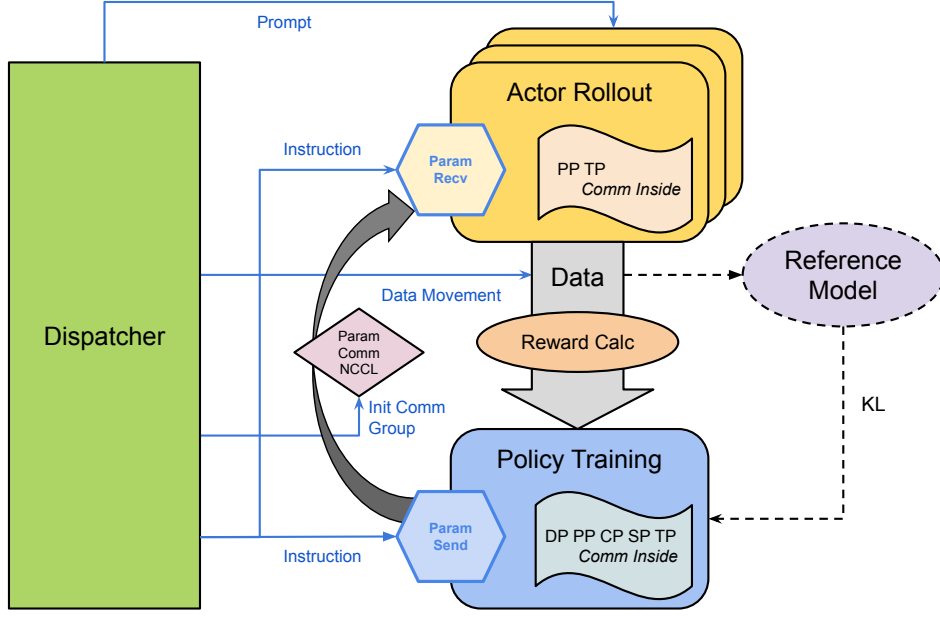


Figure 5: The architecture of the proposed RL training framework. There are mainly three parts of the framework. (1) Dispatcher: schedule and distribute training data, manage the status of the framework. (2) Actor Rollout: generate responses from the prompt, compute rewards and advantages for the policy training. (3) Policy Training: conduct an RL algorithm on the actor and optimize the actor based on the rewards and other constraints. The policy training nodes support 5D-parallelism, *i.e.* data parallel (DP), pipeline parallel (PP), context parallel (CP), fully-sharded data parallel (FSDP), and tensor parallel (TP). The actor rollout nodes support DP, PP, and TP. We adopt a customized NCCL communicator to communicate between the dispatcher and the rollout/policy node.

also improves overall framework robustness, and the fault-tolerant design inherently supports dynamic scaling, both up and down, enabling flexible resource allocation based on workload demands.

## 5. Data

In this section, we explain the data sources and their curation procedure of Physical AI SFT and RL for training Cosmos-Reason1 models.

### 5.1. Physical AI Supervised Fine-Tuning

In this stage, we fine-tune models from the previous phase on domain-specific data to specialize in Physical AI. This process aims to achieve two key outcomes: (1) enhance the model’s vision-language capabilities on Physical AI-specific data and (2) develop two critical reasoning abilities — physical common sense reasoning and embodied reasoning (detailed in Sec. 2.1 and Sec. 2.2). Unlike the previous two training stages, existing data sources cannot be directly utilized for Physical AI SFT. To address this challenge, we develop a specialized pipeline to carefully curate SFT datasets for both physical common sense and embodied reasoning applications. A portion of the Physical AI SFT data, particularly visual question-answering (VQA) pairs, are generated through a model-in-the-loop approach rather than directly from human curation.

For physical common sense, we build VQA datasets to answer free-form and multiple-choice questions from videos. For embodied reasoning, we subsample and convert existing datasets into our SFT datasets, where we cover a wide range of tasks across different embodiments, including humans, robot arms, humanoid robots, and autonomous vehicles. For each dataset, we collect two types of annotations: **understanding** and **reasoning**. An understanding annotation contains the video’s question and answer for common sense and a detailed description of states and actions in the video (a structured video caption). A reasoning annotation contains a

long chain-of-thought (CoT) thinking trace for the given text prompt and an input video. Additionally, we further curate specific reasoning SFT datasets to enhance our models’ ability to understand spatiotemporal visual stimuli (through puzzles and arrow of time in videos) as well as object permanence (through physics-based simulation). These datasets are collectively referred to as intuitive physics. Tab. 4 summarizes the datasets we used for Physical AI SFT and Fig. 6 shows examples of video frames from the Physical AI SFT datasets. Next, we describe the setting-specific curation pipelines in detail.




Figure 6: Example of video frames from our Physical AI supervised fine-tuning datasets.

#### 5.1.1. Physical Common Sense SFT

As stated earlier, for physical common sense, we collect VQA datasets that consist of both free-form and multiple-choice questions (MCQs). Our curation pipeline for physical common sense data consists of five stages:

1. **Human-in-the-loop Video Curation.** We curate a set of high-quality videos based on human preferences. We extract short clips from these videos and use them as training samples.
2. **Detailed Captioning.** We employ either reliable human annotators or pre-trained vision-language models (VLMs) to extract detailed descriptions of videos. These serve as “captions” which we use to construct understanding and reasoning annotations for video clips.
3. **Curating QA Pairs.** We prompt an LLM to construct free-form or multiple-choice questions based on the detailed clip descriptions. We construct two kinds of questions: (1) “understanding” questions that cover the content in the video (as observed through the captions) and (2) hypothetical “reasoning” questions that require information from the caption for constructing the question but cannot be answered directly based on the detailed caption. The “reasoning” questions require more thought beyond just perceiving the events and objects in the clip. Our reasoning questions focus on common sense reasoning, spatial reasoning, and temporal reasoning from videos. Prompt A.1 shows a sample question-construction prompt template used to produce reasoning questions.
4. **Extracting Reasoning Traces.** To obtain complete “reasoning” annotations, we prompt a DeepSeek-R1 (DeepSeek-AI, 2025) to answer the reasoning subset of questions by using the detailed caption as context. Then, we parse DeepSeek-R1’s response into a thinking trace and an answer. We find it is

Table 4: A summary of datasets used for physical AI supervised fine-tuning.

	Physical Common Sense VQA		BridgeData V2	Embodied Reasoning				Puzzle	Intuitive Physics		Total
	Free-form	MCQ		RoboVQA	Agibot	HoloAssist	AV		AoT	Object Permanence	
Understanding	99K	1.2M	129.2K	218.5K	19.4K	136.3K	12.4K	-	-	-	1.81M
Reasoning	59.4K	605.0K	129.1K	920.0K	19.4K	136.3K	12.3K	11.0K	30.0K	10.0k	1.93M



important to ask questions that cannot be directly answered from the caption. Otherwise, R1 can directly retrieve the answer from the provided caption, making the thinking trace invalid for model training. Our “reasoning” annotations consist of the reasoning questions, corresponding clips, thinking traces, and answers. Prompt A.2 shows a sample prompt used to elicit reasoning from DeepSeek-R1.

5. **Cleaning & Rewriting.** Finally, we employ a rule-based cleaning and rewriting stage for the “reasoning” annotations to produce valid SFT samples. Since we compress the visual context of the clip into text, rewriting helps remove unwanted references such as “description” or “caption” in the SFT training samples.

Using the aforementioned pipeline, we curated physical common sense VQA datasets consisting of both free-form and multiple-choice questions, with a few additional considerations as outlined below:

**Free-form Questions:** We used 9.9K videos from the curated set of “high-quality” clips and obtained corresponding detailed descriptions annotated by humans. The average length of human-annotated captions is  $297.4 \pm 46.4$  words. For free-form questions, we obtain  $\sim 99\text{k}$  understanding SFT samples and  $\sim 59.4\text{k}$  reasoning SFT samples through the pipeline described above.

**Multiple Choice Questions (MCQs):** To ensure our model is capable of answering multiple-choice questions (MCQs), we additionally collect a set of “understanding” and “reasoning” MCQs for the high-quality curated clips. Unlike the free-form questions, we first annotate the set of  $\sim 1.2\text{M}$  high-quality clips with detailed captions from a VLM. Using these captions, we construct  $\sim 2.4\text{M}$  “understanding” MCQs. Then, we take a subset of  $\sim 356\text{k}$  clips and use the detailed captions to produce  $\sim 600\text{k}$  “reasoning MCQs”.

### 5.1.2. Embodied Reasoning SFT

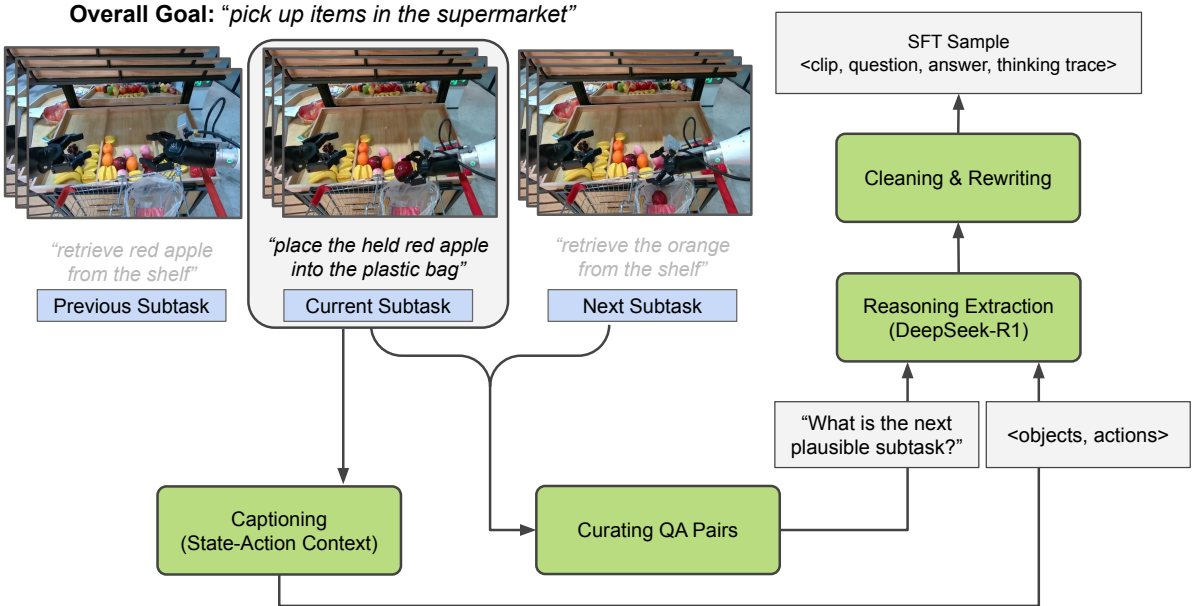


Figure 7: Embodied reasoning SFT data curation pipeline. We demonstrate an illustrative example for AgiBot, where we (1) extract short horizon segments corresponding to the subtask, (2) caption the extracted clip to obtain state-action context, (3) curate QA pairs for “next plausible subtask prediction”, (4) prompt R1 with the question and caption to elicit reasoning, (5) clean and rewrite the reasoning trace to obtain valid SFT samples.

Our embodied reasoning SFT data-curation pipeline focuses on three key properties essential for decision-making in Physical AI agents: (1) “task-completion verification”: the ability to determine whether a task or subtask has been successfully completed; (2) “action affordance”: the ability to assess whether executing a specific action or making progress toward a goal is possible; and (3) “next plausible action prediction”: the ability to identify the most plausible next action or subtask to advance toward a specified goal. These properties

are fundamental for effective decision-making across various embodiments and task configurations. To develop these reasoning capabilities, we curate SFT samples from both public and proprietary datasets. The embodied reasoning SFT dataset we use contains structured entries with four components: visual captions, questions, corresponding answers, and detailed reasoning traces.

Videos demonstrating Physical AI agents pursuing goal-oriented tasks serve as our primary source of embodied reasoning data. We collect SFT data from sources featuring demonstrations of humans, robots, or vehicles executing specific tasks. For embodied reasoning, we focus specifically on short-horizon reasoning related to our key properties of interest — determining whether an agent can reason about an immediate next subtask or action given a goal (affordance and next plausible action prediction), or evaluate the successful completion of short-horizon tasks (verifying task-completion). Since existing Physical AI demonstration datasets may lack the dense annotations needed to extract such localized action or subtask sequences, we use a series of specialized steps to extract such segments. We ensure our curated dataset is rich in terms of diversity, short-horizon granularity (immediate next action or immediate next subtask), embodiments and reasoning prompts. The curation pipeline we use has the following general steps (also illustrated in Fig. 7):

1. **Extracting Short-Horizon Segments.** Since we are interested in short-horizon reasoning tasks, we break down long video demonstrations into concise clips focused on short-horizon reasoning tasks. These segments capture either individual actions (e.g., “move left”) or distinct subtasks (e.g., “open fridge door”). When existing datasets already provide appropriately segmented clips or timestamps, we utilize them directly. Otherwise, we leverage complementary annotations such as action primitives and plans to extract these short-horizon segments.
2. **Annotating State-Action Context.** For every short-horizon clip, we use a VLM to produce structured captions that detail the present objects, their attributes and associated actions. When datasets provide supplementary annotations that can enhance the quality of such structured captions, we incorporate this information into the VLM prompt. These constitute our “understanding” annotations for embodied reasoning SFT. For AV data, we directly use human-annotated captions.
3. **Curating Reasoning QA Pairs.** We develop reasoning question-answer pairs focusing on our key properties of interest, based on available subtask and action annotations. For datasets already containing suitable QA pairs addressing our target properties, we apply minimal rule-based preprocessing before adding them to our pool of embodied reasoning SFT data.
4. **Extracting Reasoning Traces.** We utilize DeepSeek-R1 (DeepSeek-AI, 2025) to generate reasoning traces for our curated QA pairs. Since R1 lacks visual processing capabilities, we construct prompts containing the state-action context, question, and additional information (such as subtask instructions or overall goals) to elicit appropriate reasoning traces. Fig. 7 demonstrates this process, while Prompt A.3 shows an example user prompt that transforms visual information into textual context for a short-horizon question about “next plausible action”.
5. **Cleaning & Rewriting.** Finally, we use rule-based cleaning and rewriting to retain only valid and useful reasoning traces. Since we compress the visual context of the clip into text, rewriting helps remove unwanted references to the “description” or “caption”.

The exact specifics of each step in the curation pipeline vary slightly across datasets, but the overall pipeline remains the same. We now describe how this pipeline is applied across the individual data sources.

**BridgeData V2:** BridgeData V2 (Walke et al., 2023) is designed to advance scalable robot learning by providing a wide array of robotic manipulation behaviors. The dataset emphasizes foundational object manipulation tasks such as pick-and-place, pushing, and sweeping, alongside more complex activities like stacking blocks and folding cloths. It comprises 60,096 trajectories, including 50,365 teleoperated demonstrations and 9,731 scripted pick-and-place rollouts, spanning 13 distinct skills across 24 diverse environments. Each trajectory is annotated with natural language instructions corresponding to the task performed by the robot. The environments are categorized into four groups (toy kitchens, tabletops, toy sinks, and other), with a significant

portion of the data collected from seven unique toy kitchens featuring combinations of sinks, stoves, and microwaves. We first split the videos from the dataset “train” split and obtain 129.5K video clips. Then we use a VLM to caption the video clips as the understanding annotations. In the captioning prompt, we also provide additional information such as detected objects and action sequences from ECoT (Zawalski et al., 2024). We produce only “next plausible action prediction” reasoning QA pairs for BridgeData V2, where the answers correspond to action primitives such as *move left*. The reasoning annotations are generated by feeding the captions and questions to DeepSeek-R1.

**RoboVQA:** RoboVQA (Sermanet et al., 2024) is a large-scale robotics-focused visual question answering dataset. It consists of videos, instructions, and question-answer pairs of agents (robots, humans, humans-with-grasping-tools) executing a task. RoboVQA has 6 different question-types that cover aspects related to planning, verifying task-completion, discriminative affordance, generative affordance, past description and future prediction (these correspond to properties outlined before). We directly use the clips in RoboVQA without any clipping to obtain a dataset of  $\sim 220k$  clips. We caption these clips using the VLM and extract reasoning traces from DeepSeek-R1 by combining the task-context, caption and question into a suitable user prompt. This leads to  $\sim 930k$  QA pairs with reasoning traces. We filter a suitable subset post-cleaning and use the clips & QA pairs in the “train” split of the dataset for SFT. SFT samples from RoboVQA encompass all the 3 desired properties in our embodied reasoning curation pipeline.

**AgiBot:** AgiBot World (AgiBot, 2024) is a high-fidelity robot manipulation dataset. The data is collected using the AgiBot G1 hardware platform, covering a wide range of real-life tasks. It consists of 36 tasks. Each task contains multiple episodes that vary in terms of environment and objects. We subsample a portion of episodes for each task, resulting in a total of 3,300 videos. Each video is annotated with overall task information and multiple subtask annotations, including start and end frames. We utilize these action annotations to split the videos into clips, resulting in a final dataset of 19.8K clips. These clips are captioned by the VLM to convert the visual information to the scene/object descriptions and their movements. We produce only “next plausible subtask prediction” questions for AgiBot, where the answer corresponds to a subtask (“*place cucumber in the bag*”). Then we use DeepSeek-R1 to reason about the next possible subtask required to complete the task, based on the generated captions.

**HoloAssist:** Egocentric datasets capture crucial first-person perspectives that provide natural and immersive understanding of human actions, but present unique challenges including camera motion, subtle movements, occlusions, out-of-view objects, spatial perspective issues, and the need for global scene comprehension. Despite these challenges, they remain valuable for developing embodied decision-making capabilities in Physical AI systems, potentially enabling human-like interpretation and response to real-world environments. We choose to build upon HoloAssist (Wang et al., 2023), which contains 166 hours of egocentric video focused on object-centric manipulation tasks. Notably, HoloAssist uniquely includes human mistakes and the corrective steps taken to fix them. These insights can help Physical AI learn in a way that mirrors how humans learn and refine their understanding with objects in the real world. Using timestamped coarse- and fine-grained action annotations from HoloAssist, we split 1,758 videos into a final dataset of 139,653 clips. We employ a VLM to generate caption annotations. We produce only “next plausible subtask prediction” questions for HoloAssist, where the answer corresponds to a subtask. We use DeepSeek-R1 to produce reasoning traces for predicting the next possible subtasks needed to complete a task based on the generated captions. In each pipeline, we provide the task annotation as the overall goal and the fine-grained annotation as the current subtask to supplement captioning.

**Autonomous Vehicles (AV):** As a key domain in Physical AI, autonomous vehicles (AV) rely on large-scale and high-quality data to enable safe and reliable self-driving experiences, particularly in the era of rapidly scaling end-to-end systems. In this work, to avoid captioning hallucinations — especially in subtle behaviors and complex interactions — we utilize proprietary datasets with high-quality captions annotated by humans. Our dataset consists of  $\sim 12.4K$  videos, each 20 seconds long, totaling around 70 hours. Each caption includes three

categories: (1) *general description*, which details ego behaviors, environmental conditions (e.g., scene type, time of day, weather, road conditions), and critical objects (e.g., vehicles, pedestrians, cyclists, traffic lights, traffic signs); (2) *driving difficulty*, which provides a concise assessment of driving complexity based on the level of driver attention required and the scenario’s uniqueness or risk; and (3) *notice*, which highlights notable events, such as signs and signals, road user interactions, and abnormal behaviors. By leveraging these captions, we transform driving videos into structured descriptions. DeepSeek-R1 is then used to generate reasoning traces to predict the next most likely immediate action the ego vehicle will be taking based on these structured descriptions.

### 5.1.3. Intuitive Physics SFT: Spatial Puzzles, Arrow-of-Time and Object Permanence

While the prior SFT stages enable domain-specific reasoning for Physical AI applications, we incorporate additional SFT stages to develop fundamental reasoning abilities centered around intuitive physics. Although intuitive physics reasoning capabilities encompass a broad taxonomy (see Tab. 1), we focus specifically on three key aspects: reasoning about spatial continuity (through spatial puzzles), reasoning about the arrow of time (through the temporal order of events in videos), and reasoning about objectness (through simulation-based settings that assess object permanence). These tasks are inherently self-supervised by construction, which simplifies the data curation process. Despite significant progress in more sophisticated tasks, current state-of-the-art VLMs still underperform substantially on these simpler fundamental reasoning objectives. To address these limitations, we curate specialized SFT datasets targeting spatial continuity, arrow of time, and object permanence.

**Reasoning about Spatial Continuity: Spatial Puzzles.** In addition to spatial relationships, understanding spatial continuity is crucial for Physical AI tasks. To imbue our models with a fundamental understanding of spatial continuity, we further finetune them on a task of solving spatial puzzles. Specifically, we curate 3000 videos featuring diverse backgrounds, motions, and camera poses. For each video, we extract the first frame and divide it into  $2 \times 2$  patches. These patches are then shuffled to create a new clip, where one frame is one patch. We prompt the model to identify the left, top, bottom, and right positions relative to the original frame. To further increase the complexity of the task, we introduce seven additional distractor images, each also divided into  $2 \times 2$  patches. This results in a total of 32 shuffled frames for a single sample, which are provided to the model to reason about the correct positions. Additionally, we design a few “identity”-driven supplementary tasks — determining which two or three frames originate from the same image as the first frame. Akin to contrastive learning, this task requires models to develop strong spatial reasoning capabilities while distinguishing between relevant and irrelevant samples, except now with reasoning.

To generate high-quality reasoning data for spatial continuity, we first caption each of the 32 patches using a VLM and input these descriptions into DeepSeek-R1 to solve one of the three tasks. We only retain samples where R1 makes the correct prediction. Each of the 3,000 images is processed multiple times with varying distractors and shuffle orders. After filtering, our final dataset consists of 11k videos. Prompt A.4 shows a sample prompt to elicit reasoning traces.

**Reasoning about Time: Arrow-of-Time (AoT).** Similar to space, we also imbue our models with the ability to reason about time, specifically the temporal order of events in a macroscopic scale. We want our model to understand that time is inherently irreversible at a macroscopic scale, and the same can be perceived through motion and activity patterns in videos. Being able to reason about the one-way arrow of time is crucial for Physical AI, as it is closely linked to fundamental physical phenomena such as entropy, gravity, and causality. Additionally, it can serve as a proxy supervisory signal for learning intuitive physics. In particular, temporally altered or reversed videos contain physics-defying artifacts that a Physical AI capable reasoning model should be able to evaluate.

We construct an SFT dataset containing 30,000 short video clips and their reversed versions using a subset of videos from the training dataset of Cosmos-Predict1 (NVIDIA, 2025). This subset features videos containing

diverse and complex tasks that frequently involve significant motion. We prioritize videos with large motion, as they serve as the most representative examples for distinguishing the arrow of time. Unlike the prior stages, we use a VLM to directly extract a reasoning trace for both forward and reverse playback clips. We found that applying the same procedure as [Sec. 5.1.2](#) led to suboptimal results for thinking trace extraction using R1. To improve curation, we explicitly indicate whether the video is played forward or in reverse in the user prompt, helping the VLM generate more reasonable justifications. Additionally, we carefully design the prompts to ensure that reasoning traces maintain a consistent style and length for both playback directions. [Prompt. A.5](#) shows an example prompt template used for this purpose. To encourage diversity in the reasoning, we curate two distinct reasoning traces for each forward and backward video.

**Reasoning about Objectness: Object Permanence.** While spatial continuity and temporal order provide foundations for understanding physical relationships and sequences, object permanence — the understanding that objects continue to exist even when they cannot be directly observed — represents a critical reasoning capability that is fundamental for Physical AI agents. Without strong object permanence reasoning, VLMs would struggle with even basic real-world scenarios where objects frequently move in and out of view or become occluded, severely limiting their utility in Physical AI applications that require consistent object tracking and prediction capabilities.

For object permanence, we construct an SFT dataset containing 10K clips synthesized by a robot simulation platform, Libero ([Liu et al., 2023](#)). Libero offers 130 robot arm object manipulation tasks across diverse environments, desktop objects and pre-recorded arm action sequences. To enhance scene diversity, we randomly sample setups from these tasks and apply object permutation and perturbation. The camera is positioned to face the table center and orbits around the scene by selecting a random start and end point on a sphere. During playback of the pre-recorded arm actions, the camera smoothly interpolates from the start to the end point before returning near its starting position. Throughout this transition, some objects may be temporarily occluded, and once fully occluded, certain objects may be randomly removed from the scene. We prompt the model with appropriate context and ask it to analyze each clip and determine whether any objects disappear unexpectedly, violating object permanence. To ensure the model generates reasoning traces that consistently lead to the correct answer, we include hints in the prompt to indicate which objects disappear and do not reappear. However, in the final SFT dataset, these hints are removed from the prompts. For object permanence, we find that the standard pipeline of compressing visual context into captions is suboptimal for extracting useful reasoning traces for SFT. To address this, we extract thinking traces from an intermediate version of Cosmos-Reason1-7B. [Prompt A.6](#) shows a sample prompt used to elicit a reasoning trace from Cosmos-Reason1-7B.

## 5.2. Physical AI Reinforcement Learning

Table 5: Datasets for Physical AI reinforcement learning post-training.

	Common Sense	BridgeData V2	Embodied Reasoning				AV	Intuitive Physics			Total
	MCQ		RoboVQA	Agibot	HoloAssist			Puzzle	AoT	Object Permanence	
Reasoning	5,133	240	252	200	200	200		3,998	9,994	10,087	30,304

While fine-tuning establishes foundational physical common-sense and embodied reasoning capabilities, we further enhance these abilities through reinforcement learning post-training. This approach requires effective reward mechanisms, which we implement using tried-and-tested rule-based and verifiable rewards following [DeepSeek-AI \(2025\)](#). Unlike LLM domains such as mathematics and coding — where correct answers and formats are precisely defined — physical common sense and embodied reasoning typically involve free-form, open-ended responses that complicate reward assignment.

To address this challenge, we convert samples from our reasoning SFT data sources (excluding thinking traces) into multiple-choice questions with single correct answers (see [Table. 5](#) for the number of samples). This



transformation inherently enables simple, rule-based verification of responses. Our RL post-training dataset incorporates samples from all Physical AI SFT data sources, with certain subsets — specifically Spatial Puzzles, AoT, and Object Permanence data — already existing in binary question format, making them directly suitable as MCQs without modification. We manually verify the quality of samples used for RL post-training. We detail the specific characteristic of individual data sources below:

**Physical Common Sense RL Data.** We collect 5133 human annotated binary and multiple-choice questions from 1989 videos. To help control the difficulty of the questions, we use the annotated questions to evaluate four models including GPT-4o, Gemini Flason 2.0, Qwen2.5-VL-7B, and our Cosmos-Reason1-7B model. Based on the evaluation results, we further divided the collected data into two subsets: (1) the easy subset with questions that all models got them correct; (2) the hard subset with questions that at least one model got them wrong.

**Embodied Reasoning RL Data.** We select 200-250 SFT samples from each embodied reasoning data source and convert them to multiple-choice questions (MCQs). To ensure high-quality RL post-training, we carefully verify that these samples are free from answer and instruction ambiguity while maintaining balanced distribution across MCQ options to prevent potential reward hacking. This process requires some manual intervention, particularly for non-binary questions where we must select appropriate distractor options that are plausible yet clearly incorrect. The human-in-the-loop is required to ensure question quality, making it difficult to generate large-scale MCQ data for training.

**Intuitive Physics RL Data.** As previously described, our self-supervised intuitive physics SFT data naturally exists in MCQ format by design, making it scalable to generate diverse questions. For these tasks, we implement additional quality assurance measures to ensure balanced option distributions across all samples. We carefully avoid any overlap with clips used during SFT to prevent early saturation during RL post-training. For the RL post-training phase, we curate 24079 high-quality samples across spatial continuity, arrow of time, and object permanence tasks.

## 6. Benchmark

We compare our trained models with other counterparts on a benchmark geared specifically towards measuring reasoning capabilities about physical common sense, embodied decision-making. In this section, we discuss the procedure of building our common sense and embodied reasoning benchmarks (Tab. 6). We evaluate models by asking either binary yes / no questions or multiple-choice questions (MCQs) based on video context. We note that our benchmarks necessitate reasoning by construction (to arrive at the correct answer) and we only measure the accuracy of the final answer. We leave quantitatively assessing the quality of the thinking trace for future work.

Table 6: Statistics of our curated benchmarks.

Common Sense MCQ	Embodied Reasoning						AV	Total
	BridgeData V2	RoboVQA	RoboFail	Agibot	HoloAssist			
604	100	110	100	100	100		100	1214

### 6.1. Physical Common Sense Reasoning

We construct a physical common sense reasoning benchmark by manually curating questions about internet video clips according to the ontology defined in Sec. 2.1. We initially collected a pool of 5737 questions, including 2828 binary questions and 2909 multiple-choice questions. Fig. 8 shows the distribution of question categories according to our ontology. After that, we went through a manual process to carefully select a subset of 604 questions from 426 video clips used as our physical common sense benchmark, from which 336 are



binary questions and 268 are multiple choice questions. Among the 604 questions, 80 (13.25%) are about Space, 298 (49.33%) are about Time, and 226 (37.4%) are about Fundamental Physics.

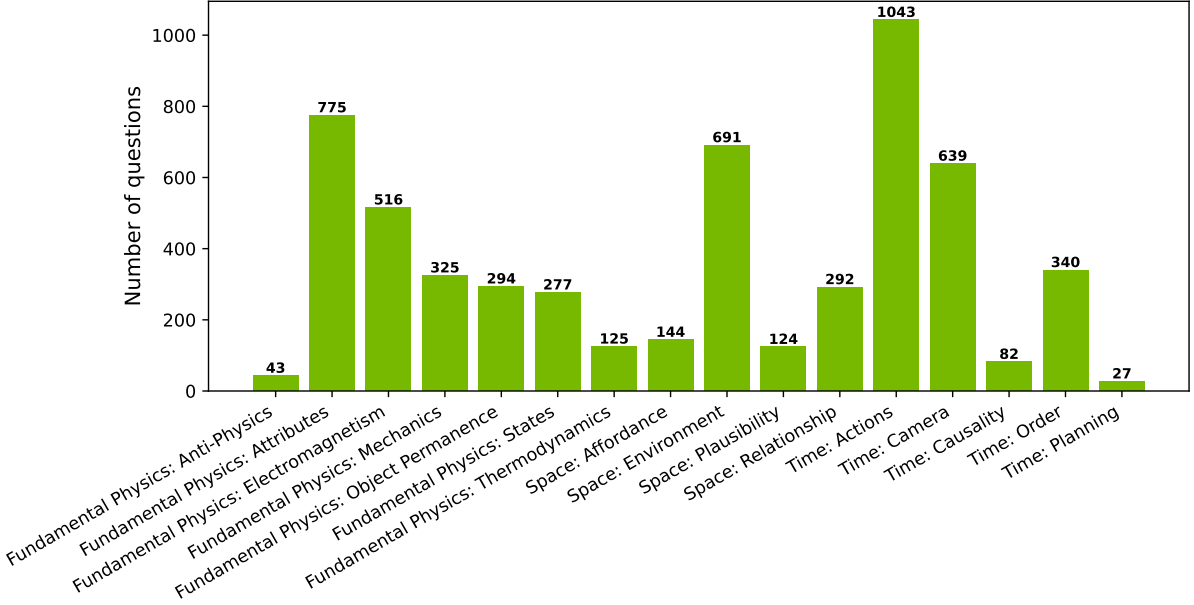


Figure 8: An illustration showing the categorical distribution of initial questions collected for physical common sense benchmark according to the ontology defined in Tab. 1. We select a subset of 604 questions as our evaluation benchmark.

## 6.2. Embodied Reasoning

Similar to our SFT data-curation pipeline for embodied reasoning, we constrain our embodied reasoning benchmark to focus on the previously outlined properties before — (1) “task-completion verification”: the ability to determine whether a task or subtask has been successfully completed; (2) “action affordance”: the ability to assess whether executing a specific action or making progress toward a goal is possible; and (3) “next plausible action prediction”: the ability to identify the most plausible next action or subtask to advance toward a specified goal. We present our embodied reasoning benchmark samples as multiple-choice questions (MCQs) to enable automatic evaluation across models. We adopt a few key steps to ensure our benchmark is useful for measuring embodied reasoning abilities.

1. **Unified Question Templates.** We adopt a unified format for question formulation to ensure that reasoning is conditioned on visual input rather than textual cues. This approach also helps align action granularity and reduce ambiguity across different datasets.
2. **Unified Action Granularity.** We pay special attention to action granularity. In predicting the next immediate action, multiple choices could be potentially correct. For example, the action “water the plant” may involve steps such as “grab the watering can”, “move the watering can”, and “pour the watering can”. However, these steps can also be broken down into finer sub-actions like “move left” or “tilt down”. To address this complexity, we use a hierarchy of actions (Belkhale et al., 2024): we define atomic-level actions as “actions”, more coarse-grained actions as “subtasks”, and dataset-specific tasks as “goals”.
3. **Manual Refinement.** Apart from these systemic approaches to addressing ambiguity, we also manually refine the MCQ choices. These modifications help resolve overly similar options, enforce visual reasoning by preventing answers from being inferred solely from text, and emphasize the full context of the entire clip rather than just the earlier frames.

**RoboVQA:** We sample 101 clips from the “val” split (excluded from SFT training) of the processed RoboVQA samples for our benchmark. For benchmarking purposes, we only consider clips and question-answer pairs

that correspond to either verifying task-completion (whether a provided instruction was followed successfully) or affordance (whether it is possible to follow an instruction / complete a task). These are included as multiple-choice yes / no questions.

**RoboFail:** In addition to RoboVQA, we also manually curate and annotate 100 examples from the RoboFail (Liu et al., 2023) dataset to create a harder “action affordance” and “task completion verification” evaluation split, where the hardness of samples is dictated by (1) the necessity of highly observant perception or comprehensive temporal context processing, (2) identifying physical constraints blocking follow-through for an action (unlike RoboVQA where it’s impossible to complete an action due to perception mismatches or irrelevant instructions) and (3) being able to reason about nuanced questions.

**BridgeData V2:** We split the videos from “val” split of BridgeData V2 into clips following the same way as the training set. Then we sample 100 clips to create 100 multiple-choice QA pairs as the benchmark. In each question, we provide the task that the robot gripper is instructed to perform in the clip, and ask what the most plausible next immediate action is given what the robot has done in the video.

**AgiBot:** We sample 100 clips from the processed AgiBot SFT data to generate 100 multiple-choice QA pairs. For each clip, we additionally provide the task information and ask which of the given subtasks is the most likely next subtask the robot should work towards. We randomly sample the choices from the subtask sequence of the clip’s entire trajectory. Notably, these trajectories are excluded from the training set.

**HoloAssist:** We sample 100 clips from the processed HoloAssist SFT data to generate 100 multiple-choice QA pairs. For each clip, we additionally provide the coarse-grained action annotation as the overall goal and ask which of the given subtasks is the most likely next subtask. We randomly sample the choices from other fine-grained action annotations under that coarse-grained action. All 34 videos (out of a total of 1758) that contain these clips are excluded from the training set to prevent episode leakage.

**AV:** We curate 100 videos from the proprietary data to construct 100 multiple-choice QA pairs. These videos exhibit diverse lateral and longitudinal behaviors, along with rich interactions. The questions are designed to (1) predict the next immediate action the ego vehicle will most likely be taking, (2) verify the completion of a previously executed action, and (3) assess the possibility of a specific action in a given scenario. Note that these samples in the benchmark are excluded from the training set.

## 7. Experiments

In this section, we illustrate the experiment setup for physical AI supervised fine-tuning and physical AI reinforcement learning of Cosmos-Reason1 and discuss the evaluation results on our benchmarks.

### 7.1. Physical AI Supervised Fine-Tuning

We train Cosmos-Reason1-7B for  $12.5K$  iterations with a cosine annealing learning rate decaying from  $1 \times 10^{-5}$  to  $1 \times 10^{-6}$ . We train Cosmos-Reason1-56B for  $30K$  iterations with a learning rate of  $1 \times 10^{-5}$ , followed by  $20K$  iterations with a decayed learning rate of  $1 \times 10^{-6}$ . We use a global batch size of 256 for Cosmos-Reason1-7B and 32 for Cosmos-Reason1-56B. The fused Adam optimizer with  $\beta_1, \beta_2 = (0.9, 0.95)$  and weight decay of 0.1. We follow a balanced data sampling strategy during training so that no specific domain of interest is overrepresented during SFT. Unless stated otherwise, for our models, we report the average accuracy of 5 inferences (temperature 0.6 and top-p 0.95) with different random seeds. For evaluating other models, we employ a zero-shot chain-of-thought prompting (Kojima et al., 2022) by either calling their APIs (GPT-4o, OpenAI o1, Gemini 2.0 Flash) or using their open-sourced model checkpoints (Qwen2.5-VL).

#### 7.1.1. Physical Common Sense Results

Tab. 7 shows the evaluation results on physical common sense benchmark. Cosmos-Reason1-7B and Cosmos-Reason1-56B demonstrate significantly improved capabilities on the physical common sense benchmark com-

pared to their respective backbones, with the 56B variant achieving the best accuracy, slightly outperforming OpenAI o1. These results highlight the effectiveness of our curated common sense dataset, laying a strong foundation for further RL improvements.

Table 7: Evaluation on physical common sense benchmark.

Methods	Space	Time	Other Physics	Avg.
Gemini 2.0 Flash	53.8	50.0	46.9	50.2
GPT-4o	61.3	54.7	50.9	55.6
OpenAI o1	63.8	58.1	58.0	59.9
Qwen2.5-VL-7B	48.8	56.4	37.2	47.4
Nemotron-H-56B	61.3	68.1	45.1	58.2
Cosmos-Reason1-7B	54.2	58.7	50.0	54.3 (+6.9)
Cosmos-Reason1-56B	61.3	65.5	53.9	60.2 (+2.0)

### 7.1.2. Embodied Reasoning Results

Tab. 8 shows the evaluation results on the embodied reasoning benchmark. Cosmos-Reason1 models achieve significantly stronger results than the baseline models on this benchmark, with both the 7B and 56B variants demonstrating over a 10% improvement compared to their respective backbone VLMs. The results demonstrate that our physical AI SFT is highly effective in boosting models’ Physical AI embodied reasoning capabilities.

Table 8: Evaluation on embodied reasoning benchmark.

Models	BridgeData V2	RoboVQA	Agibot	HoloAssist	AV	RoboFail	Avg.
Gemini 2.0 Flash	25.0	78.2	29.0	44.0	37.0	67.0	46.7
GPT-4o	42.0	71.8	32.0	65.0	46.0	63.0	53.3
OpenAI o1	42.0	80.0	44.0	63.0	37.0	61.0	54.5
Qwen2.5-VL-7B	38.0	82.5	40.4	50.0	36.0	57.6	50.8
Nemotron-H-56B	37.0	77.2	37.0	65.0	41.0	64.0	53.5
Cosmos-Reason1-7B	58.8	83.8	49.4	63.0	55.6	60.0	61.8 (+11.0)
Cosmos-Reason1-56B	65.0	80.0	47.6	57.8	65.8	66.2	63.7 (+10.2)

### 7.1.3. Intuitive Physics Results

Although VLMs are often perceived as specialists achieving superhuman performance, our study reveals that many struggle with basic physics reasoning. To test the model’s capacity to understand intuitive physics, we curate 100 videos for each of the three tasks: the arrow of time, spatial puzzle, and object permanence, and generate 100 questions following pipeline in Sec. 5.1.3. We conduct data decontamination to ensure no overlap with the training data. We evaluate model performance on the curated test set.

Tab. 10 shows that the existing models struggle to perform above chance level on arrow of time and object permanence tasks. Notably, GPT-4o and OpenAI o1 handle spatial puzzles much better than random guessing. This observation suggests that current multimodal models are more proficient at reasoning about spatial relationships than temporal dynamics. Given that these models generally perform well on standard benchmarks like MMMU, this suggests that existing evaluations fail to capture their understanding of the physical world. However, our curated intuitive physics dataset enables the 7B model to improve significantly across all three tasks, demonstrating Cosmos-Reason1’s fundamental capability to reason in intuitive physics.

## 7.2. Physical AI Reinforcement Learning

We post-train Cosmos-Reason1 with simple, rule-based verifiable rewards to further enhance their Physical AI reasoning abilities. We summarize our experimental findings on physical common sense, embodied and intuitive physics reasoning tasks.

### 7.2.1. Experiment Setup

We adopt two types of rule-based rewards to optimize our model towards accurate reasoning of Physical AI:

1. An **accuracy reward** evaluates whether the model’s response, enclosed within the `<answer></answer>` tags, matches the ground truth. Since we exclusively use MCQs for RL, this verification can be performed simply through string matching.
2. A **format reward** encourages the model to encapsulate the thinking process in the `<think></think>` tag, and the answer in the `<answer></answer>` tag. This is implemented as a regular expression matching.

During training, we sample from each RL dataset with equal probability, ensuring balanced representation across different domains. We also dynamically shuffle the MCQ choices on-the-fly to encourage generalization. We use a global batch size of 128 questions, for each question we sample 9 outputs, each with a maximum length cutoff of 6144 tokens. We set the learning rate to  $4 \times 10^{-6}$ , the coefficient of the KL penalty term to 0.005, and train the model for 500 iterations.

Table 9: Evaluation on physical common sense and embodied reasoning benchmark.

Models	Common Sense	BridgeData V2	RoboVQA	Agibot	HoloAssist	AV	RoboFail	Avg.
Cosmos-Reason1-7B	54.3	58.8	83.8	49.4	63.0	55.6	60.0	60.7
+ <b>Physical AI RL</b>	56.2	73.5	86.8	54.2	60.0	67.0	62.0	65.7 (+5.0)

Table 10: Evaluation on intuitive physics benchmark.

Models	Arrow of Time	Spatial Puzzle	Object Permanence	Avg.
Random Guess	50.0	25.0	50.0	41.7
Gemini 2.0 Flash	50.0	31.0	48.0	43.0
GPT-4o	50.0	77.0	48.0	58.3
OpenAI o1	51.0	64.0	49.0	54.7
Qwen2.5-VL-7B	50.2	27.2	48.8	42.1
Cosmos-Reason1-7B	56.0	85.4	82.0	74.5 (+32.4)
+ <b>Physical AI RL</b>	64.5	94.0	86.0	81.5 (+7.0)

### 7.2.2. Physical Common Sense and Embodied Reasoning Results

For physical common sense and embodied reasoning, we find that Physical AI RL post-training improves performance across most benchmark components, with the notable exception of RoboFail. The results are summarized in [Tab. 9](#). Performance on RoboFail remains consistently challenging through both SFT and RL stages. This is unsurprising given RoboFail’s deliberate design as a hand-curated benchmark featuring difficult real-world scenarios testing “action affordance” and “task-completion verification”. The benchmark’s difficulty stems from several factors: (1) samples requiring highly observant perception or comprehensive temporal context processing, and (2) affordance questions involving complex physical constraints in action execution, unlike those in RoboVQA.

We attribute the stagnant performance on RoboFail primarily to insufficient representative training data. This hypothesis is supported by examining the specific error patterns in both finetuned and post-trained models, which include: inadequate reasoning about complex affordance scenarios, failure to be sufficiently observant of crucial visual details, and overthinking when encountering nuanced out-of-distribution questions. We believe targeted training on similar representative samples would significantly improve performance on this challenging benchmark. We deliberately maintain RoboFail in our evaluation suite as a valuable indicator of areas requiring further development in embodied reasoning capabilities.

Interestingly, we found that through RL, our model learns to carefully evaluate the provided choices and reject

## EMBODED REASONING BENCHMARK



Based on the ego vehicle's actions in the video, what is the most likely immediate action it will take next?

A: turn right, B: turn left, C: change to right lane, D: change to left lane

## RESPONSE BEFORE RL

Okay, let's break this down. The video shows the ego vehicle is driving on a single-lane road with a 25 mph speed limit. The road has parked cars on the right, so the driver is maintaining a safe distance there. The left side is a two-way divider with oncoming traffic, which is coming from the left in the video.

Since the road is single-lane, changing lanes (right or left) might not be possible. But the video shows a single-lane road, which usually means one lane in each direction, so there's no shoulder to the right for parking, but parked cars are already on the right. Maybe the road is a single lane each way, but parked cars are on the right side of that lane. So the driver is in the correct lane, just keeping a safe distance from the parked cars. The oncoming traffic is in the opposite lane, so the driver is staying in their lane. Changing to the left lane would be into oncoming traffic. So that's not possible. Therefore, the answer must be D.

## RESPONSE AFTER RL

Okay, let's break this down. The user provided a driving scenario and asked for the most likely immediate action based on the ego vehicle's current actions. First, I need to recall the details given. The ego vehicle is on a narrow, two-way residential street with parked cars on both sides. The driver is maintaining a steady speed, staying in their lane, and being cautious because there's oncoming traffic in the opposite lane. The speed limit is 25 mph. The parked cars limit visibility, so they have to be extra careful. The road markings are double yellow lines, meaning no passing, so lane changes aren't allowed. The driver is keeping a consistent pace and staying vigilant.

Now, the question is about predicting the next immediate action. The options are turning right, left, changing to the right lane, or left lane. Let's consider each option.

Option A: Turn right. The video shows a two-way street. The road is narrow with parked cars on both sides, so turning right might not be possible.

Option B: Turn left. Similar reasoning as A. Unless there's a left turn ahead, but the video doesn't indicate any upcoming intersections or turns.

Option C: Change to right lane. The current lane is the right one since it's a two-way street. Changing to the right lane again doesn't make sense because there's already parked cars on both sides. So this option is probably not correct.

Option D: Change to left lane. Wait, but in a two-way street with double yellow lines, there's no left lane to change into. The driver is in their respective lane (right for their direction), and the opposite lane is for oncoming traffic. So changing to the left lane would be into oncoming traffic, which is dangerous and not allowed. So this option doesn't fit either.

Wait, maybe the question is a trick. The video shows the road markings are double yellow lines, which means no passing, so changing lanes is prohibited. Therefore, the driver cannot change lanes, so options C and D are impossible. The only remaining options are turning right or left. But since the video doesn't show any turns, and the driver is maintaining their path, perhaps the answer is that none of the options are correct, but since we have to choose, I have no choice but to select none. Therefore, the answer is none.

Figure 9: Intriguing results before and after RL. When presented with an ambiguous question, we observe that after RL, our model learns to reject all provided options based on its knowledge.

them all if the question is ambiguous. As shown in Fig. 9, the model assesses the feasibility of each choice and adopts a conservative action that is not within the choices when ambiguity arises.

### 7.2.3. Intuitive Physics Results

Similar to self-supervised learning, a key advantage of intuitive physics is the ease of data scaling. Generating training data for the arrow of time requires only simple video reversal, spatial puzzles can be applied to any image to create challenging spatial continuity problems, and object permanence can be easily implemented in any simulated environment. Moreover, these data sources can be seamlessly adapted for RL training with verifiable rewards. With this recipe, we manage to create a larger RL dataset compared to common sense and embodied reasoning tasks with minimal human effort.

Tab. 10 shows that with careful SFT data curation and targeted training, Cosmos-Reason1-7B achieves significant improvements across all tasks, while Physical AI RL is able to further enhance spatial puzzle and object permanence capabilities. However, reasoning about the arrow of time still remains a challenge.

Through Physical AI RL, we found that the model further develops reasoning in space, time and object permanence. Fig. 10 illustrates the differences in how Cosmos-Reason1 reasons about time before and after RL. The model can identify anti-physics motions — e.g. powder defying gravity to rise into the bowl — while remaining unaffected by stationary distractors in the video. This demonstrates that its reasoning extends beyond mere perception. Similarly, in Fig. 11, models without RL tend to conflate spatial questions with temporal reasoning. While they can perceive that the second frame lacks similarity to the first, their inherent biases lead them to default to video order, suggesting they rely more on temporal cues than genuine spatial understanding. RL with spatial puzzles enables the model to extract key features from the first frame and systematically compare them across multiple frames, allowing it to accurately determine spatial relationships. Lastly, Fig. 12 shows that even with long CoT, the model without Physical AI RL struggles with object permanence, often confusing itself when reasoning about an object’s appearance and disappearance. In contrast, the RL model quickly reaches conclusions through direct and concise reasoning.

## 8. Related Work

### 8.1. Foundation Models for Physical AI


Early approaches to integrating large pre-trained models into embodied agents primarily relied on using existing large language models (LLMs) and vision-language models (VLMs) as static modules. For example, the line of work using LLM as zero-shot task planner (Li et al., 2022; Song et al., 2023; Huang et al., 2022; Ahn et al., 2022) extracts high-level action plans in natural language directly from pre-trained LLMs without additional fine-tuning, demonstrating the ability of LLMs to generalize across robotic tasks. Similarly, Code as Policies (Liang et al., 2022) extends the use of LLMs to generate structured code for robotic control. However, these pre-trained models, especially the VLMs, are not explicitly designed for physical interactions and, therefore, often lack the physical common sense understanding and result in suboptimal plans and execution.

Another line of work focuses on directly training vision-language-action (VLA) models end-to-end for embodied agents (Brohan et al., 2023; Driess et al., 2023; Kim et al., 2024; Yang et al., 2025; Gemini Robotics Team, 2025). They often start with a pre-trained vision-language model and fine-tune the model with robotics data to generate embodied actions. This approach has shown significant promise in improving generalization and robustness across robotic tasks. A notable subcategory of these methods introduces hierarchical policy representations that separate high-level language-driven planning from low-level action execution (Shi et al., 2025; Li et al., 2025). These system often use a high-level VLM interpret natural language instructions and generate plans, while low-level VLA execute fine-grained motor control. These hierarchical architectures improve task and decomposition.

Recent efforts have sought to enable embodied AI with stronger reasoning abilities. One such approach is



## INTUITIVE PHYSICS: ARROW OF TIME



Detect whether the video plays forward or backward with confidence.

A: forward, B: backward

## RESPONSE BEFORE RL

Okay, let's see. The user is asking whether the video plays forward or backward. The video shows a step-by-step process of mixing pink powder into a white liquid in a green bowl. The key points shown are the hand holding a blue container, pouring the pink powder into the bowl, and the camera being static.

Now, the question is about the direction of playback. Since the video shows a process where steps are performed in a logical sequence-like pouring the powder into the liquid-it's likely that the video is showing the correct order of events. If it were played backward, the powder would be disappearing from the bowl into the container, which doesn't make sense in this context. The actions shown, such as pouring and mixing, are typical of a forward process. The video does not show any reverse actions or objects moving in the opposite direction. The camera angle is static, so there's no indication of reverse motion effects. Therefore, the video should be playing forward.

## RESPONSE AFTER RL

Okay, let's break this down. In the sequence of frames provided, the hand holding the container appears to be pouring a pink powder into a green bucket containing a white powder.

As the sequence progresses, the amount of pink powder in the container increases while the amount in the bucket decreases. This suggests that the action is being undone, which is characteristic of a video playing in reverse. Additionally, the hand's motion and the position of the container change in a way that indicates the action is being retraced, further supporting the idea that the video is playing backward.

The text on the screen remains constant throughout the frames, indicating that the content of the video is not changing, but rather the action being performed is being undone.

The text on the screen remains constant throughout the frames, indicating that the content of the video is not changing, but rather the action being performed is being undone.

Figure 10: While the model before RL struggles to understand and associate perception with reverse actions, RL enables the model to reason through time while avoiding distractions, such as stationary text.

the embodied chain-of-thought (CoT) framework (Zawalski et al., 2024), which enables robotic agents to reason through sequential decisions before executing actions. This aligns with broader trends in AI where explicit reasoning mechanisms improve interpretability and adaptability. Similarly, Liu et al. (2023); Elhafsi et al. (2023) introduce methods that perform step-by-step reasoning for explaining and correcting failures. While these methods advance the cognitive capabilities of embodied AI, most still rely on manual prompting to structure their reasoning processes, limiting their autonomous adaptation and generalization. CoT-VLA (Zhao et al., 2025) introduces a CoT process which first predicts a future image and then predicts the action which should interact with the environment to achieve the predicted future image.

Beyond robotics, VLA models have been applied to other physical embodiments, such as autonomous driving. For example, CoVLA (Arai et al., 2024) introduces a large-scale vision-language-action dataset specifically for self-driving applications, facilitating research into multimodal decision-making in autonomous systems.

## 8.2. Vision Language Models

The community has made significant strides in building vision language models. Prominent families of models include Flamingo (Alayrac et al., 2022), LLaVA (Liu et al., 2023), InternVL (Chen et al., 2024), QwenVL (Bai et al., 2025), NVLM (Dai et al., 2024), Llama-3.2-Vision (Grattafiori et al., 2024). These vision language models typically adopt one of two common architectures: the decoder-only architecture, as seen in models like LLaVA (Liu et al., 2023) and InternVL (Chen et al., 2024), which integrates image tokens within the LLM’s self-attention layers, and the cross-attention-based architecture, exemplified by Flamingo (Alayrac et al., 2022) and Llama-3.2-Vision (Grattafiori et al., 2024), where image tokens are processed through the LLM’s cross-attention layers. Dai et al. (2024) compares both architectures in a state-of-the-art setting and finds that the decoder-only architecture exhibits stronger reasoning capabilities in college-level multidisciplinary knowledge and mathematical reasoning tasks within a visual context. Building on this, we adopt the decoder-only architecture to develop the reasoning model for Physical AI.

## 8.3. LLMs and VLMs with Reasoning Capabilities

Early studies have shown that large language models (LLMs) exhibit basic reasoning capabilities in mathematics (Cobbe et al., 2021), coding (Chen et al., 2021), and general reasoning tasks. These capabilities can be further enhanced through chain-of-thought prompting (Wei et al., 2022). Recently, OpenAI o1 (OpenAI, 2024; Jaech et al., 2024) demonstrated that LLMs’ reasoning capabilities in coding and mathematics can be significantly enhanced through large-scale reinforcement learning. Notably, the open-sourced DeepSeek-R1 (DeepSeek-AI, 2025) has shared its training methodology with the community, providing valuable insights into building high-performance reasoning models. However, existing studies primarily focus on reasoning tasks related to coding, mathematics, and STEM fields (Liu et al., 2024), even within multimodal reasoning settings (Qwen-Team, 2024). Recently, there has been a surge of efforts aimed at integrating R1’s reasoning capabilities into VLMs (Liu et al., 2025; Zhou et al., 2025; Zhao et al., 2025; Huang et al., 2025; Haonan Wang, 2025). In this work, we explore reasoning capabilities in the context of Physical AI.

## 9. Conclusion

In this work, we present Cosmos-Reason1, a family of multimodal large-language models specialized in physical world understanding and reasoning. To specialize the model in Physical AI, we define ontologies to encapsulate foundational capabilities for Physical AI models and construct supervised fine-tuning data and benchmarks for common sense and embodied reasoning accordingly. We further explore the idea of Physical AI RL crafting rule-based, verifiable rewards and using reinforcement learning to improve the model’s capabilities with reasoning about space, time, and intuitive physics. Our experiments show that Physical AI SFT improves the backbone VLM’s performance by more than 10% on the proposed physical common sense and embodied reasoning benchmarks. Physical AI RL further boosts accuracy by over 5%. With Physical AI SFT and RL, Cosmos-Reason1 can learn intuitive physics, such as the arrow of time and object permanence, which existing models struggle

with. We will make our code open-source and models open-weight to expedite the progress of building Physical AI systems that understand and perform complex tasks in the physical world.

## A. Prompts Used for Data Curation

### A.1. Physical Common Sense Question-Construction Prompt

#### PHYSICAL COMMON SENSE QUESTION-CONSTRUCTION PROMPT

You will be given a detailed caption describing the video. Your task is to generate 6 extremely challenging questions to evaluate the reasoning ability of a state-of-the-art model that require multi-step deep reasoning from the caption.

Try your very best to use your creativity to generate extremely challenging questions!

Here is a list of categories of questions you should generate:

1. Common sense reasoning, including but not limited to:
  - Physical common sense, such as gravity, balance, stability, support, elasticity, deformation, lighting, heat, motion, acceleration, etc.
  - Physical attributes that are not directly mentioned in the caption, such as mass, temperature, etc.
  - Object state changes, such as egg changed from liquid to solid, steak changed from raw to cooked, etc.
  - Object permanence, such as object visibility, occlusion, etc.
2. Spatial reasoning, including but not limited to:
  - Spatial plausibility, such as whether the object can be placed in a certain location, in a certain orientation, etc.
  - Affordance, such as whether the object can be used for a certain purpose, etc.
  - Scene or surrounding environment that is not directly mentioned in the caption, such as in a tunnel, underwater, weather (sunny, rainy, etc.), etc.
3. Temporal reasoning, including but not limited to:
  - Complex action understanding, such as subtask or goal decomposition, whether a task is completed, etc.
  - Temporal order of events, such as before/after/simultaneously, etc.
  - Planning, such as whether the object can be used for a certain purpose, come up with a plan based on the video, what are the next steps, etc.

Below are some additional rules you must follow:

1. You must create questions that require both the information in the caption and the external knowledge to challenge the model's reasoning ability with your creativity.
2. You must NOT create any questions with answers that are directly given in the caption.
3. You must NOT create any questions that can be answered by external knowledge only without the information from the video caption.
4. When asking questions, you should give as little information as possible. The model you are evaluating on is expected to get any information needed to answer the question from the video itself.

In your 6 questions, 2 of them should be about common sense reasoning and planning using world knowledge, 2 of them should be about spatial reasoning, and 2 of them should be about temporal reasoning.

Your question should be concise with a maximum of 10 words.

This is the caption:

{caption}

You should treat video caption as the video. Focus on the video itself and do not mention anything related to the captions. For example, you should not mention "the caption", "the description", "what is mentioned", etc.

Instead, you can use wordings like "in the video", "the video shows", etc.

## A.2. Physical Common Sense Reasoning Extraction Prompt

### PHYSICAL COMMON SENSE REASONING EXTRACTION PROMPT

This is the video you see:

This video showcases noodles being prepared.

The video begins with someone standing by a large boiling pot and putting noodles into one of six small circular compartments dipped into boiling water in a large circular stainless steel pot. The compartments are also stainless steel. The person using their right hand, stirs the noodles in a circular motion using light brown wooden chopsticks, and their left holds the compartment into place using a dark metallic stick-like object. The pot is placed on a black burner. Steam rises from the pot, and bubbles form, indicating active boiling. The person cooking is wearing a grey top with the sleeves reaching his elbow and has a rope tied to their waist.

The setup includes a stainless steel surface stained with white drops, a silver pot filler tap above the boiling pot, and on its left there are two deep silver filters. A stainless steel pot with a stainless steel lid on the cooking pot's left, and a white cloth placed on the lid. The background also includes a big spoon next to the burner used to cook the noodles. The lighting is bright, causing reflections and shadows over the silver surface. The camera keeps alternating a push-in and a push-out motion in a medium shot over a high angle, providing a clear, unobstructed view of the cooking process.

Answer the following question:

What would happen if the pot was not boiling?

### A.3. AgiBot Reasoning Extraction Prompt

#### AGIBOT REASONING EXTRACTION PROMPT

The state of the environment is as follows:

Inside a supermarket near a fruit stand. Well-lit indoor environment with artificial lighting. The floor appears to be tiled, and the overall ambiance suggests a clean and organized retail space.

Camera - Positioned at a medium height, slightly above the fruit stand, capturing the robot arms and the fruit stand from a frontal perspective. The camera provides a clear, steady shot focusing on the interaction between the robot arms and the fruits. It captures the details of the objects in the foreground while keeping the background slightly out of focus. The camera remains stationary, providing a stable view of the robot arms and the fruit stand.

Fruit Stand - Center of the frame, consisting of wooden trays filled with various fruits. The fruit stand has multiple compartments made of light-colored wood. It contains a variety of fruits such as bananas, oranges, apples, pears, and others, arranged neatly in rows.

Shopping Cart - In the foreground, partially visible under the robot arms. A standard metal shopping cart with red handles and a basket containing some items, including a plastic bag where the apple is being placed.

Plastic Bag - Hanging from the shopping cart, below the robot arms. A transparent plastic bag with some red items inside, likely other fruits or vegetables. Hanging from the shopping cart, partially filled with items. Receives the red apple dropped by Robot Arm 2.

Grapes - Held by Robot Arm 1, positioned over the fruit stand. A cluster of dark purple grapes, appearing fresh and ripe.

Apple - Held by Robot Arm 2, positioned over the plastic bag in the shopping cart. A single red apple, shiny and smooth, indicating freshness.

A robot is performing the following task (may not be finished in the video): Pickup items in the supermarket

This is how the robot's actions impact the environment:

Robot Arm 1 - On the left side of the frame, extending towards the fruit stand. A mechanical arm with a black and white color scheme. It has a gripper mechanism at the end holding a bunch of grapes. The arm is articulated with joints allowing for precise movement. Holds a bunch of grapes with its grippers, positioned over the fruit stand.

Robot Arm 2 - On the right side of the frame, extending towards the fruit stand. Another mechanical arm similar in design to Robot Arm 1 but with a different configuration of the gripper mechanism. It is holding an apple and appears to be placing it into a plastic bag. Holds a red apple with its grippers, positioned above the plastic bag in the shopping cart. Releases the red apple into the plastic bag in the shopping cart.

**\*\*Output Requirements\*\*:**

- Predict the next immediate action of the robot.

**\*\*Response Format\*\*:**

- <action> your predicted action </action>.



#### A.4. Spatial Puzzle Reasoning Extraction Prompt

##### SPATIAL PUZZLE REASONING EXTRACTION PROMPT

You will be given a puzzle, which has 8 images, each broken into a 2x2 grid: top-left, top-right, bottom-left, bottom-right.

As a result, there are 32 frames and the first frame shows that "{first\_frame\_caption}".

{frames\_and\_captions}

For the remaining 31 frames, which three are most likely to be from the same images as the first frame? Then among the three, which one is most likely to be at {direction} of the first frame?

Give your answer in the following format:

<answer>

Same image: Frame a, Frame b, Frame c

{direction}: Frame d

</answer>

#### A.5. Arrow of Time Reasoning Extraction Prompt

##### ARROW OF TIME REASONING EXTRACTION PROMPT FOR FORWARD VIDEO, BACKWARD VIDEO PROMPT IS SIMILAR

###### HUMAN TURN:

Here is a sequence of frames. Is it playing forward or backward?

###### ASSISTANT TURN:

The video is playing forward.

###### HUMAN TURN:

So here is a forward-playing video. Describe in details what visual clues do you notice that indicate the time is going forward i.e. the video is not reversed? Please avoid repeating the question and refrain from using bullet points. Present the evidence but no need to give the verdict.

#### A.6. Object Permanence Reasoning Extraction Prompt

##### OBJECT PERMANENCE REASONING EXTRACTION PROMPT

This is a video of a robotic simulation environment. The robotic arm moves and may occlude objects. The camera moves around the scene and then returns to a position near its initial location. Occlusion can also occur due to the camera's movement.

The objects in the videos are akita black bowl and akita black bowl and cookies and glazed rim porcelain ramekin and plate and wooden cabinet and flat stove.

akita black bowl is occluded in the middle of video and then vanishes unexpectedly.

Is there an object that becomes temporarily occluded but does not reappear in the end, contradicting object permanence?

Please answer the question in the following format: <think> your reasoning </think> <answer> your answer </answer>.

## B. Additional results

### INTUITIVE PHYSICS: SPATIAL PUZZLE



You are given 32 frames from different contexts. Looking at the first frame, which other frame is most likely to be at right of the first frame?

A: Frame 2, B: Frame 30, C: Frame 9, D: Frame 11

### RESPONSE BEFORE RL

Okay, let's tackle this question. The user is asking which frame is most likely to be on the right of the first frame based on the given video description. First, I need to recall the details of each frame. The first frame shown is a close-up of a shirtless man lifting weights with a barbell overhead. The question is about the spatial relationship between frames, specifically which frame would be positioned to the right of this first frame.

Looking at the options provided: A is Frame 2, B is Frame 30, C is Frame 9, D is Frame 11. To figure out the correct answer, I need to check the order of the frames as listed in the video description. The frames are numbered sequentially from 1 to 32, and each frame's content gives a clue about its position relative to others.

The video starts with Frame 1 (the weightlifter), then moves through various scenes. Let's scan through the order to see which frame comes after the first one. For example, after Frame 1, the next frame might be Frame 2, which could be the two women walking in an airport. But the question is about the right side of the first frame, not necessarily the next frame in the sequence. However, since the frames are shown in order, and the answer options are based on their numerical sequence, it's likely that the next frame (Frame 2) would be adjacent or in the same general area, especially if the video is structured as a continuous sequence with adjacent scenes. Looking at the given frames, Frame 2 is the two women in an airport. If the first frame is the weightlifter, and the next logical scene is the women walking, then Frame 2 would be to the right. Alternatively, if the sequence jumps around, but given the structure of the video (as per the video), the frames are presented in order. Therefore, the right adjacent frame would be the next one in sequence, which is Frame 2. Other options like Frame 30 or 9 or 11 don't immediately follow Frame 1 in the video order.

So the most logical choice based on the sequence provided is Frame 2.

### RESPONSE AFTER RL

Okay, let's break this down. To determine which frame is most likely to be to the right of the first frame, we need to identify frames that share similar characteristics with the first frame. The first frame depicts a shirtless man with long hair and a beard doing pull-ups in a minimalist room. Key elements include the man's appearance, the pull-up bar, and the room's simplicity.

We examine the other frames for matching details. Frame 11 shows a shirtless man with curly hair and a beard doing pull-ups in a room with beige walls and a closed door. This closely matches the first frame, particularly the man's appearance and the room's features. Both frames involve pull-ups and share a minimalist aesthetic.

Other frames do not align as closely. Frame 4 shows a child on a chair, which is unrelated. Frame 17 features a kitchen, which is a different setting. Frame 14 shows a person in a kitchen, also unrelated. Frame 24 depicts a snowy landscape, which is entirely different. Given the similarities between Frame 1 and Frame 11, it is reasonable to conclude that Frame 11 is the most likely candidate to be to the right of the first frame. Both frames depict the same man performing pull-ups in a similar setting, making them strong candidates for being part of the same image.

Therefore, the correct answer is Frame 11.

Figure 11: While the model before RL mistakenly associates spatial questions with temporal reasoning, RL enables the model to identify key features of the first frame and compare them with subsequent frames.

## INTUITIVE PHYSICS: OBJECT PERMANENCE



This is a video of a robotic simulation environment. The robotic arm moves and may occlude objects. The camera moves around the scene and then returns to a position near its initial location. Occlusion can also occur due to the camera's movement. Is there an object that becomes temporarily occluded but does not reappear in the end, contradicting object permanence?

A: Yes, B: No

## RESPONSE BEFORE RL

Okay, let's break this down. The user is asking whether there's an object that gets temporarily hidden (occluded) but doesn't come back, which would challenge the idea of object permanence. First, I need to recall what the video shows. The setup has a robotic arm moving around three silver plates and a red plate. The camera moves around the scene and then goes back to its original spot. The key points here are about occlusion and object permanence. Object permanence is the concept that objects continue to exist even when they're not visible. So if something is occluded (temporarily hidden) but reappears later, that's consistent with object permanence. But if it's occluded and doesn't come back, that would contradict it.

Looking at the actions: the robotic arm picks up a silver plate and moves it towards the red plate. The camera moves around, possibly showing parts of the scene from different angles. The video shows that the camera returns to its original position. So when the arm moves the plate, the camera's movement might cause temporary occlusion (like when the arm is between the camera and the plate), but once the camera moves back, the plate should still be there. However, the question is about whether the object comes back after being occluded.

The video ends with the arm moving the plate towards the red plate. The camera returning to its initial position might mean that the arm's action of moving the plate is completed, but does the plate reappear? Since the arm is moving the plate, once it's moved, the plate is no longer in its original position. But object permanence isn't about the object's location changing; it's about whether the object is lost from view entirely and not found again. If the camera's movement causes temporary occlusion (like the arm blocking the plate from view), but when the camera moves back, the plate is still there (just not in view), then that's okay. The plate doesn't vanish and reappear; it's just not visible at a moment.

The key here is that the plate was moved by the arm. If during the camera's movement, the arm's action caused the plate to be occluded (like moving it out of view), but once the camera returns, the plate is no longer in the camera's frame but still exists elsewhere. However, since the camera returns to its original position, and the arm has moved the plate, the plate might not reappear in the camera's view again. But does that count as the object coming back? If the arm has moved it out of the scene entirely, then it's gone, which would contradict object permanence. But if it's just moved to another part of the table and the camera's original position doesn't show it again, but the plate is still there somewhere else, that's fine.

Wait, the question says "occluded but does not reappear in the end" So if the plate is moved by the arm and then the camera's movement doesn't show it again, but the plate is still in the video (like on the table but not in the camera's current view), then it's still there. So the plate didn't vanish, just not in view. So object permanence holds because the object still exists, even if not seen. Therefore, the answer would be "No", because even if occluded temporarily, the object remains present and would reappear if the camera or arm moves back. The plate is moved but still exists in the video, just not in the camera's current view. So temporary occlusion doesn't contradict object permanence. Therefore, the answer is B.

## RESPONSE AFTER RL

Okay, let's break this down. The question is asking if there's an object that becomes temporarily occluded ( meaning hidden from view for a bit ) but then doesn't reappear, which would contradict the idea of object permanence. Object permanence is the concept that objects exist even when not seen, so if something disappears and doesn't come back, that's a problem. Looking at the video: The akita black bowl is shown as being occluded in the middle of the video and then vanishes unexpectedly. The key points here are that it's occluded (so temporarily hidden) and then it disappears. The question is whether the disappearance contradicts object permanence.

Object permanence would mean that even if something is hidden, it should still exist. But in the video, the bowl vanishes unexpectedly. So if it's no longer there, that would mean it's gone, which would contradict object permanence. But wait, maybe the bowl was moved out of the camera's view, but the video shows vanishes unexpectedly, implying it's not just moving out of the camera but actually disappearing.

So the answer would be A, yes, because the bowl disappears and doesn't reappear, which would contradict object permanence. The other objects like the plates and cabinet are shown as remaining, so only the bowl is the one that disappears. Therefore, the correct answer is A.

Figure 12: While the model before RL fails to reason about object permanence through long CoT, RL enables the model to correctly infer that the object's disappearance is not due to camera movement, using concise and direct reasoning.

## C. Contributors and Acknowledgments

### C.1. Core Contributors

Alisson Azzolini, Junjie Bai, Jiaxin Cao, Prithvijit Chattopadhyay, Huayu Chen, Yin Cui, Yifan Ding, Liang Feng, Siddharth Gururani, Imad El Hanafi, Zekun Hao, Jacob Huffman, Jingyi Jin, George Kurian, Nayeon Lee, Zhaoshuo Li, Xuan Li, Maosheng Liao, Tsung-Yi Lin, Ming-Yu Liu, Xiangyu Lu, Wei Ping, David W. Romero, Shuran Song, Lyne Tchapmi, Andrew Z. Wang, Boxin Wang, Haoxiang Wang, Fangyin Wei, Jiashu Xu, Dinghao Yang, Xiaodong Yang, Zhuolin Yang, Jingxu Zhang, Xiaohui Zeng, Zhe Zhang

**Contributions:** YC, SS, MYL, TYL defined physical common sense and embodied reasoning ontologies. YC, PC, JX, AZW, TYL curated physical common sense data and benchmarks. XY, PC, FW, XL, AZW, SG, TYL curated embodied reasoning data and benchmark. ZL, JJ, TYL designed the captioning method. JX, ZH, LT, JJ, XL, ZL, SG, TYL curated self-supervised data and benchmark. YC, PC, SG post-processed extracted reasoning traces. BW, NL, ZY, WP trained the base VLM models. DWR, HW, XZ built the infrastructure for supervised fine-tuning. XZ, GK, AA, HW, SG, ZL, DWR, FW, TYL trained the Physical AI supervised fine-tuned models. ZL, NL, FW, YD evaluated the reasoning models. JH, IH, ZZ, HW, ZH, DWR, AA, JC, LF, ML, XLu, DY, JZ, JB built the reinforcement learning infrastructure. ZH, JH, HC, JX, LT trained Physical AI reinforcement learning models. YC, TYL organized paper writing. MYL, TYL designed the overall system.

### C.2. Contributors

Hannah Brandon, Jinju Chu, Jenna Diamond, Francesco Ferroni, Rama Govindaraju, Jinwei Gu, Brendan Johnson, Rizwan Khan, Elena Lantz, Yen-Chen Lin, Alice Luo, Andrew Mathau, Yun Ni, Lindsey Pavao, Misha Smelyanskiy, Yao Xu

**Contributions:** YX, LP, AM, RK, JC, BJ, EL, HB, JD helped collect human annotations. JG helped with the vision encoder. FF curated training data. YN, RG, MS supported scaling model training infrastructure. YCL provided insights on physical common sense ontology. AL helped with overall data curation.

### C.3. Acknowledgments

We’d like to thank Wenliang Dai, Guo Chen, Guilin Liu, Zhiding Yu, Mohammad Shoeybi, Andrew Tao, Bryan Catanzaro for discussion and data curation of general VLM training. Xinshuo Weng, Boris Ivanovic for data curation of AV. Moo Jin Kim for setting up a simulation environment of object permanence data collection. Heng Wang for human annotation pipeline. Pooya Jannaty for AoT idea discussion.

## References

- [1] AgiBot. Agibot world colosseum. <https://github.com/OpenDriveLab/AgiBot-World>, 2024. 13
- [2] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. In *CoRL*, 2022. 22
- [3] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. In *NeurIPS*, 2022. 6, 24
- [4] H Arai, K Miwa, K Sasaki, Y Yamaguchi, et al. Covla: Comprehensive vision-language-action dataset for autonomous driving. In *ICRA*, 2024. 24
- [5] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibong Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 6, 8, 24
- [6] Suneel Belkhale, Tianli Ding, Ted Xiao, Pierre Sermanet, Quon Vuong, Jonathan Tompson, Yevgen Chebotar, Debidatta Dwibedi, and Dorsa Sadigh. Rt-h: Action hierarchies using language. *arXiv preprint arXiv:2403.01823*, 2024. 17
- [7] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023. 22
- [8] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021. 24
- [9] Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*, 2024. 6, 24
- [10] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *CVPR*, 2024. 6, 24
- [11] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021. 24
- [12] Wenliang Dai, Nayeon Lee, Boxin Wang, Zhuolin Yang, Zihan Liu, Jon Barker, Tuomas Rintamaki, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. NVLM: Open frontier-class multimodal LLMs. *arXiv preprint arXiv:2409.11402*, 2024. 6, 24
- [13] DeepSeek-AI. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 1, 2, 6, 10, 12, 15, 24
- [14] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, et al. Palm-e: An embodied multimodal language model. In *ICML*, 2023. 22
- [15] Amine Elhafi, Rohan Sinha, Christopher Agia, Edward Schmerling, Issa AD Nesnas, and Marco Pavone. Semantic anomaly detection with large language models. *Autonomous Robots*, 2023. 24

- 
- [16] Google DeepMind Gemini Robotics Team. Gemini robotics: Bringing ai into the physical world, 2025. URL <https://deepmind.google/discover/blog/gemini-robotics-brings-ai-into-the-physical-world/>. 22
  - [17] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. 6, 24
  - [18] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. In *COLM*, 2023. 6
  - [19] Tianyu Pang Haonan Wang, Chao Du. V1: Toward multimodal reasoning by designing auxiliary task, 2025. URL <https://faint-basin-c34.notion.site/1b2db15ae55b800da077e70aafe40212>. 24
  - [20] Jian Hu, Xibin Wu, Zilin Zhu, Xianyu, Weixun Wang, Dehao Zhang, and Yu Cao. Openrlhf: An easy-to-use, scalable and high-performance rlhf framework. *arXiv preprint arXiv:2405.11143*, 2024. 8
  - [21] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *ICML*, 2022. 22
  - [22] Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Yao Hu, and Shaohui Lin. Vision-r1: Incentivizing reasoning capability in multimodal large language models. *arXiv preprint arXiv:2503.06749*, 2025. 24
  - [23] Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. OpenAI o1 system card. *arXiv preprint arXiv:2412.16720*, 2024. 24
  - [24] Daniel Kahneman. *Thinking, Fast and Slow*. Macmillan, 2011. 3
  - [25] MJ Kim, K Pertsch, S Karamcheti, T Xiao, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024. 22
  - [26] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. In *NeurIPS*, 2022. 18
  - [27] Yann LeCun. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review*, 2022. 3
  - [28] Shuang Li, Xavier Puig, Chris Paxton, Yilun Du, Clinton Wang, Linxi Fan, Tao Chen, De-An Huang, Ekin Akyürek, Anima Anandkumar, et al. Pre-trained language models for interactive decision-making. In *NeurIPS*, 2022. 22
  - [29] Y Li, Y Deng, J Zhang, J Jang, M Memme, and R Yu. Hamster: Hierarchical action models for open-world robot manipulation. In *ICLR*, 2025. 22
  - [30] Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. Code as policies: Language model programs for embodied control. In *ICRA*, 2022. 22
  - [31] Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. Libero: Benchmarking knowledge transfer for lifelong robot learning. In *arXiv preprint arXiv:2306.03310*, 2023. 15
  - [32] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 6, 24
-



- 
- [33] Z Liu, A Bahety, and S Song. Reflect: Summarizing robot experiences for failure explanation and correction. In *CoRL*, 2023. 18, 24
  - [34] Zihan Liu, Yang Chen, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. Acemath: Advancing frontier math reasoning with post-training and reward modeling. *arXiv preprint arXiv:2412.15084*, 2024. 24
  - [35] Ziyu Liu, Zeyi Sun, Yuhang Zang, Xiaoyi Dong, Yuhang Cao, Haodong Duan, Dahua Lin, and Jiaqi Wang. Visual-rft: Visual reinforcement fine-tuning. *arXiv preprint arXiv:2503.01785*, 2025. 24
  - [36] Meredith Ringel Morris, Jascha Sohl-Dickstein, Noah Fiedel, Tris Warkentin, Allan Dafoe, Aleksandra Faust, Clement Farabet, and Shane Legg. Position: Levels of agi for operationalizing progress on the path to agi. In *ICML*, 2024. 3
  - [37] NVIDIA. Cosmos world foundation model platform for physical ai. *arXiv preprint arXiv:2501.03575*, 2025. 14
  - [38] NVIDIA. Nemotron-h: A family of accurate and efficient hybrid mamba-transformer models. *arXiv preprint arXiv:2504.03624*, 2025. 6, 8
  - [39] Nvidia, Bo Adler, Niket Agarwal, Ashwath Aithal, Dong H Anh, Pallab Bhattacharya, Annika Brundyn, Jared Casper, Bryan Catanzaro, Sharon Clay, Jonathan Cohen, et al. Nemotron-4 340b technical report. *arXiv preprint arXiv:2406.11704*, 2024. 6
  - [40] OpenAI. Learning to reason with llms, 2024. URL <https://openai.com/index/learning-to-reason-with-llms/>. 1, 24
  - [41] Qwen-Team. Qvq: To see the world with wisdom, 2024. URL <https://qwenlm.github.io/blog/qvq-72b-preview/>. 24
  - [42] Ronan Riochet, Mario Ynocene Castro, Mathieu Bernard, Adam Lerer, Rob Fergus, Véronique Izard, and Emmanuel Dupoux. Intphys 2019: A benchmark for visual intuitive physics understanding. *TPAMI*, 2021. 3
  - [43] Pierre Sermanet, Tianli Ding, Jeffrey Zhao, Fei Xia, Debidatta Dwibedi, Keerthana Gopalakrishnan, Christine Chan, Gabriel Dulac-Arnold, Sharath Maddineni, Nikhil J Joshi, et al. Robovqa: Multimodal long-horizon reasoning for robotics. In *ICRA*, 2024. 13
  - [44] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024. 8
  - [45] Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint arXiv: 2409.19256*, 2024. 8
  - [46] Lucy Xiaoyang Shi, Brian Ichter, Michael Equi, Liyiming Ke, Karl Pertsch, Quan Vuong, James Tanner, Anna Walling, Haohuan Wang, Niccolo Fusai, et al. Hi robot: Open-ended instruction following with hierarchical vision-language-action models. *arXiv preprint arXiv:2502.19417*, 2025. 22
  - [47] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *CVPR*, 2016. 6
-

- [48] Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*, 2019. 8
- [49] Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M. Sadler, Wei-Lun Chao, and Yu Su. Llm-planner: Few-shot grounded planning for embodied agents with large language models. In *ICCV*, 2023. 22
- [50] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 6
- [51] Roger Waleffe, Wonmin Byeon, Duncan Riach, Brandon Norick, Vijay Korthikanti, Tri Dao, Albert Gu, Ali Hatamizadeh, Sudhakar Singh, Deepak Narayanan, et al. An empirical study of mamba-based language models. *arXiv preprint arXiv:2406.07887*, 2024. 6, 7
- [52] Homer Rich Walke, Kevin Black, Tony Z Zhao, Quan Vuong, Chongyi Zheng, Philippe Hansen-Estruch, Andre Wang He, Vivek Myers, Moo Jin Kim, Max Du, et al. Bridgedata v2: A dataset for robot learning at scale. In *CoRL*, 2023. 12
- [53] Xin Wang, Taein Kwon, Mahdi Rad, Bowen Pan, Ishani Chakraborty, Sean Andrist, Dan Bohus, Ashley Feniello, Bugra Tekin, Felipe Vieira Frujeri, Neel Joshi, and Marc Pollefeys. Holoassist: an egocentric human interaction dataset for interactive ai assistants in the real world. In *ICCV*, 2023. 13
- [54] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*, 2022. 24
- [55] J Yang, R Tan, Q Wu, R Zheng, B Peng, and Y Liang. Magma: A foundation model for multimodal ai agents. In *CVPR*, 2025. 22
- [56] Michał Zawalski, William Chen, Karl Pertsch, Oier Mees, Chelsea Finn, and Sergey Levine. Robotic control via embodied chain-of-thought reasoning. In *CoRL*, 2024. 13, 24
- [57] Jiaxing Zhao, Xihan Wei, and Liefeng Bo. R1-omni: Explainable omni-multimodal emotion recognition with reinforcement learning. *arXiv e-prints*, pages arXiv-2503, 2025. 24
- [58] Qingqing Zhao, Yao Lu, Moo Jin Kim, Zipeng Fu, Zhuoyang Zhang, Yecheng Wu, Zhaoshuo Li, Qianli Ma, Song Han, Chelsea Finn, et al. Cot-vla: Visual chain-of-thought reasoning for vision-language-action models. *arXiv preprint arXiv:2503.22020*, 2025. 24
- [59] Hengguang Zhou, Xirui Li, Ruochen Wang, Minhao Cheng, Tianyi Zhou, and Cho-Jui Hsieh. R1-zero's "aha moment" in visual reasoning on a 2b non-sft model. *arXiv preprint arXiv:2503.05132*, 2025. 24