

# SWITCH: BENCHMARKING INTERACTION AND VERIFICATION ON REAL-WORLD INTERFACES IN LIFE-LONG EMBODIED AGENTS

Jieru Lin<sup>1,2\*</sup>, Zhiwei Yu<sup>1</sup>, Börje F. Karlsson<sup>1†</sup>

<sup>1</sup> Beijing Academy of Artificial Intelligence (BAAI), <sup>2</sup> Harbin Institute of Technology

## ABSTRACT

Autonomous agents operating in the real world must interact continuously with existing physical and semantic infrastructure, track delayed consequences, and verify outcomes over time. Everyday environments are rich in tangible control interfaces (TCIs)—*e.g.*, light switches, appliance panels, and embedded GUIs—posing core challenges for lifelong embodied agents, including partial observability, causal reasoning across time, and failure-aware verification under real-world constraints. Yet, current benchmarks rarely consider such long-horizon interaction and causality requirements. We introduce **SWITCH** (Semantic World Interface Tasks for Control & Handling), an embodied, task-driven benchmark created through iterative releases to probe these gaps. Its first iteration, SWITCH-Basic<sup>1</sup>, evaluates five complementary abilities—*task-aware VQA*, *semantic UI grounding*, *action generation*, *state-transition prediction*, and *result verification*—under egocentric RGB video input and device diversity across 351 tasks spanning 98 real devices/appliances. Results from commercial and open LMMMs reveal systematic failures, highlighting critical gaps for lifelong agent deployment. SWITCH provides data, code, and held-out splits to enable reproducible non-contaminated evaluation and community contributions toward more challenging future iterations of the benchmark and the creation of relevant training data.

## 1 INTRODUCTION

Lifelong autonomous agents operating in the real world must interact persistently with existing physical and semantic infrastructure, track delayed consequences, and verify outcomes over time. Despite significant progress in computer vision (CV), large multi-modality models (LMMMs), and interactive agents, research efforts have largely overlooked understanding and operating tangible control interfaces (TCIs)<sup>2</sup>, from light switches and appliances to on-device GUI panels — the primary medium of **everyday human-device interaction**.

Recent benchmark efforts (*e.g.*, Foss et al. (2025); Chen et al. (2025); Chow et al. (2025)) have probed models’ ability to understand common-sense causality or physics-based outcomes and consistency. However, TCI interfaces require modeling causality beyond these dimensions only, as the effects of TCI interactions can depend on specific devices (*e.g.*, same device type has different behaviours), present temporal delays (*e.g.*, pre-heat an oven), and even may require spatial verification (*e.g.*, pressing a switch in one room to turn a light in another location).

These challenges are further amplified by real-world deployment constraints. Despite decades of work on smart homes and IoT, most buildings and devices remain **non-API-enabled** or expose frag-

\*Work performed during the first author’s internship at BAAI.

†Correspondence to <borje@baai.ac.cn>.

<sup>1</sup><https://github.com/BAAI-Agents/SWITCH>

<sup>2</sup>We use the term TCI (Tangible Control Interfaces) to clearly specify this type of interaction interface in tangible form. As well as avoid ambiguity with usage of UI (usually associated directly with GUIs), TUI (which refers to physically interacting with the digital world through tangible assets — Ishii & Ullmer (1997), not our scenario), and HMI (Human-Machine-Interface, terminology used in industry for visual control interfaces for settings as complex machinery or facilities — Stouffer et al. (2023), *e.g.*, control of a power plant).

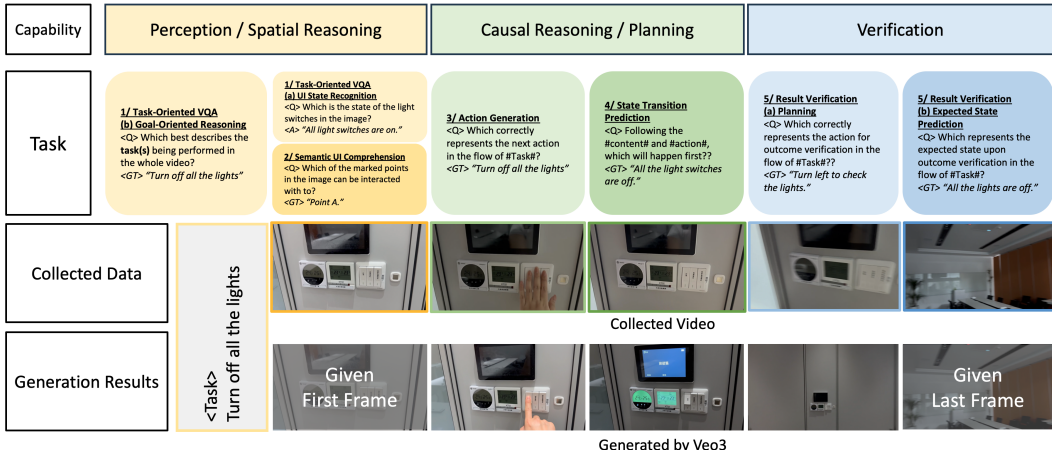


Figure 1: An overview of the SWITCH benchmark, using the case “Turn off all the lights” as a running example. SWITCH covers the **collection and annotation** of real-world TCI interaction data (“Collected Data”), which we systematically structure into **five distinct tasks**. These tasks are designed to evaluate models across **three crucial capability dimensions: Perception/Spatial Reasoning, Causal Reasoning/Planning, and Verification**. Furthermore, we leverage the benchmark to evaluate advanced generative models, like Veo3 Google (2025). By comparing generated videos against ground truth, we illustrate how current models still exhibit significant critical gaps in logical consistency and fine-grained interaction for real-world use, thus underscoring the importance of SWITCH’s target scenarios.

mented, proprietary control mechanisms. Consequently, agents operating in everyday environments cannot assume seamless programmatic access, but must rely on vision, physical action, and outcome verification through tangible interfaces. Evaluating whether current foundation models can function in such settings therefore requires moving beyond text-image benchmarks toward embodied tasks grounded in realistic human–device interaction.

Similarly to the work on manipulating GUIs, grounding is critical for proper **situated interaction** with TCIs. However, current benchmarks and simulators either don’t cover such interfaces or don’t model them in enough detail transferable to real-world settings. These issues highlight the critical need for models that can understand and operate within such interaction-centric environments, bridging not only perception, reasoning, and action; but also **verification of outcomes and adaptation**. As well as ways to evaluate them in both higher-level causality understanding and planning of actions, but also in grounded realistic execution of interactions.

To address these gaps, we introduce **SWITCH** (Semantic World Interface Tasks for Control and Handling), a new effort towards a unified task-driven benchmark designed to evaluate interaction-centric capabilities essential for lifelong embodied agents. Its initial release, **SWITCH-Basic**, emphasizes five complementary capabilities: *1/ Task-Aware Visual Question Answering*: answering questions conditioned on multimodal observations and task goals; *2/ Semantic UI Comprehension*: grounding and interpreting actionable UI elements in context; *3/ Action Generation*: planning and executing context-aware actions aligned with user intent; *4/ State Transition Prediction*: reasoning about the causal consequences of UI-actions; *5/ Result Verification*: post-hoc evaluation to determine task success. Figure 1 illustrates the benchmark design dimensions.

SWITCH adopts an iterative benchmark design to enable timely assessment of current model capabilities, while supporting progressively more challenging future iterations. Its first iteration emphasizes single-step interactions and scalable automatic evaluation through multimodal multiple-choice questions (MCQ), complemented by case studies with generative video models unconstrained by predefined answers. Across a diverse set of real consumer devices, evaluations of state-of-the-art commercial and open large multimodal models reveal inconsistent performance, with strong reliance on textual cues and limited use of visual and temporal evidence—limitations that are likely to compound under continuous deployment.

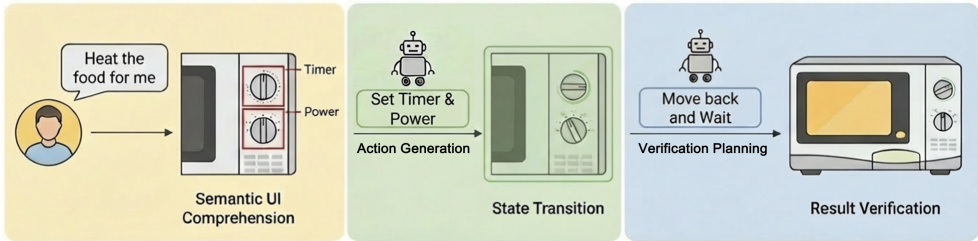


Figure 2: Example of a one-step task. The agent interprets the user ask and identifies related TCI elements (e.g., timer, power), generates corresponding actions (e.g., only a timer change is needed). The state transition occurs as the switch settings change. Since the agent is too close to the device during operation, result verification planning involves moving back to capture the full view of the relevant microwave part for reliable result verification and waiting.

In summary, our work makes two main contributions. First, we introduce **SWITCH**, an open and reproducible benchmark that starts to address previously underexplored challenges in TCI-centric interactive intelligence. Second, we conduct a comprehensive evaluation of state-of-the-art large multi-modality models (LMMs) on this benchmark - SWITCH-Basic, providing detailed analysis of their strengths and limitations, thereby offering insights to guide future model development in real-world interactive tasks. SWITCH also provides a leaderboard<sup>3</sup> for each iteration and difficulty level, with held out answers for fair comparison between different approaches.

## 2 RELATED WORK

From the perspective of lifelong agents, evaluating large multimodal models and world models is not merely about measuring predictive accuracy or short-horizon task success. Agents are expected to operate continuously in dynamic environments fundamentally shaped by physical tangible interfaces—from switches to smart appliances—that encode human intent. While a wide range of benchmarks has emerged, spanning **multimodal physical perception and reasoning** (e.g., Chow et al. (2025); Chen et al. (2025); Wang et al. (2025); Puyin et al. (2025); Zhang et al. (2025b); Li et al. (2023)), **world model evaluation** (e.g., Li et al. (2025); Duan et al. (2025); Cai et al. (2025); Wiedemer et al. (2025); Zhang et al. (2025a); Motamed et al. (2025)), and **digital interfaces manipulation** (e.g., Xie et al. (2024); Rawles et al. (2024); Tan et al. (2025); Yao et al. (2022)), they remain largely fragmented and exhibit incomplete scenario coverage. Specifically, physical benchmarks often overlook the semantic logic of control panels; world model evaluations frequently prioritize visual fidelity or passive video generation over interactive task success, as in Zhang et al. (2025a); and digital agents master virtual GUIs, but lack grounding and real-world causality. SWITCH addresses these gaps by focusing on Tangible Control Interfaces (TCIs), uniquely combining the semantic understanding of user interfaces with the physical verification and causal reasoning required for embodied world modeling. An expanded discussion is presented in Appendix A.

## 3 SEMANTIC WORLD INTERFACE TASKS

By bridging the gap between virtual, physical, and control-centric interactions, the SWITCH benchmark enables a comprehensive assessment of models’ embodied intelligence and world modeling capabilities in realistic settings. Section 3.1 introduces the design principles behind the benchmark. Next, Section 3.2 details the data curation process.

### 3.1 BENCHMARK DESIGN

Lifelong embodied agents must continually acquire, retain, and reuse interaction knowledge as they encounter new devices, layouts, and contexts. In modern human environments, switches and control interfaces provide a ubiquitous and low-level interaction primitive through which agents can ground perception into action and observe causal effects over time.

<sup>3</sup><https://huggingface.co/spaces/BAAI-Agents/SWITCH-Basic-Leaderboard>

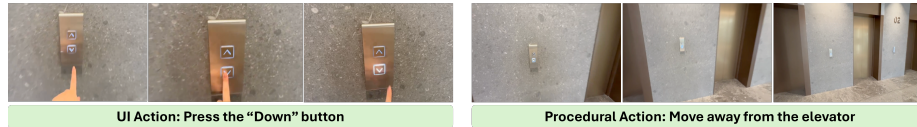


Figure 3: Examples of the two action categories. Left: UI Action (*i.e.*, grounded interaction). Right: Procedural Action (*i.e.*, transitional or verification steps).

To illustrate the breadth and depth of our evaluation framework, we present a canonical example in Figure 2, which breaks down a simple, single-step task (heating food) into the fundamental stages of our benchmark: from semantic comprehension and action generation to state transition and result verification. For a more complex scenario (illustrated in Figure 6 in the appendices), we present a multi-step task (printing a document) that involves sequential reasoning, adaptation to dynamic UI changes, and completing a full perception-action-verification loop. These scenarios are enabled by our five core tasks, which we define as follows:

**1/ Task-Aware Visual Question Answering (VQA):** Most VQA datasets focus on static images or generic reasoning, often ignoring task relevance, user intent, and interaction context. In contrast, our benchmark defines a task-aware VQA setting composed of two complementary sub-tasks: **(a) UI State Recognition:** assessing whether the model can recognize and describe the current state of TCI elements within the scene; and **(b) Goal-Oriented Reasoning:** testing whether the model can interpret the purpose and outcome of actions observed in dynamic video sequences, reasoning about whether these interactions successfully achieve the intended task goals. Together, these sub-tasks measure a model’s ability to ground semantics in visual context, comprehend high-level task objectives, and reason about the operational consistency between actions and outcomes.

**2/ Semantic UI Comprehension:** Existing visual recognition and interactive datasets focus on object detection or purely digital virtual elements, largely ignoring TCI components and their functional roles. However, in real-world human-computer and human-robot interactions, understanding interactive UI elements and their spatial-functional relationships is essential. This task tests whether a model can accurately localize and interpret actionable TCI elements in cluttered or dynamic settings, reasoning about their spatial and functional relationships while inferring human intent.

**3/ Action Generation:** Previous action prediction datasets typically emphasize isolated motion sequences or simulations, lacking the need to comprehend TCI functions and layouts, user goals, or environmental context. In realistic interactive scenarios, models must integrate these elements to select meaningful actions. This task evaluates a model’s ability to infer intent and plan executable, context-aware action sequences that achieve user objectives. **(a) UI Action Identification:** Detect the relevant interaction region, recognize its affordances, and predict the appropriate mode of interaction; **(b) Action Execution Planning:** Generate the necessary physical actions to perform the intended TCI operation, such as moving toward the interface or manipulating its components.

**4/ State Transition Prediction:** Previous state prediction benchmarks often assume deterministic, object-centric transitions, overlooking the diverse causal effects of TCI controls and their dependency on environmental context. This task focuses on predicting the immediate consequences of an executed action—capturing how TCI states, physical surroundings, and observable views change in response. It evaluates a model’s capacity for causal reasoning, short-term prediction, and fine-grained understanding of action–state dynamics in interactive environments. **(a) UI-State Transition:** Predict changes in the visual or functional state of TCI elements after an action (*e.g.*, switch toggled, button activated). **(b) Environment-State Transition:** Predict corresponding updates in the surrounding physical or visual environment (*e.g.*, lights turning on, revealing a new scene or perspective). **(c) Coupled Transition:** Reason about interdependent updates where TCI and environment states jointly change as part of the same causal event (*e.g.*, pressing a control panel triggers a screen update and changes ambient light level).

**5/ Result Verification:** Existing benchmarks rarely assess post-hoc verification or causal reasoning in interactive contexts. However, verifying whether an intended action has successfully achieved its goal is essential for reliable real-world deployment—especially when diverse UI controls and device responses introduce uncertainty. **(a) Verification Planning:** Test whether the model can infer what actions or checks are required to verify the outcome of a previous operation, such as observing key

Table 1: Overview of the SWITCH-Basic benchmark. (a) Categories and devices covered in the collected data. (b) Detailed statistics with overall video duration of 74m12s.

(a) Data scenarios		(b) Data statistics		
Categories	Devices	Categories	Statistics	Num
Environment Controls	Light Switch, Air Conditioner Panel, Air Purifier, Fan, Humidifier	Collected	Video Samples	193
		Annotated	High-level Actions	193
Fine-grained Actions	508			
- UI Action	350			
- Procedural Action	81			
- Verification	77			
Overall States	772			
Office Equipments	Printer, Adjustable Standing Desk, Robot Lifter, PC Power Switch, Power Strip Switch, Water Dispenser	- UI-State Trans.	509	
		- Env-State Trans.	263	
Kitchen & Home	Microwave, Washing Machine, Kitchen Hood, Kettle, Air Fryer, Stove, Electric Cooker, Coffee Machines	Processed	Unique Pairs	1489
		- 1/ Task-Aware VQA	516	
Building Infrastructure	Elevator	- 2/ Semantic UI Comp.	246	
		- 3/ Action Generation	309	
		- 4/ State Trans.	316	
		- 5/ Result Verification	102	

indicators, querying system feedback, or performing follow-up interactions. **(b) Expected State Prediction:** Assess whether the model can predict what the expected state should look like after a successful interaction, providing a causal grounding for evaluating success or failure.

### 3.2 DATA CURATION

The data curation process is organized into three key parts: Video Collection, Data Annotation, and Question Processing, ensuring a structured and systematic approach. Detailed statistics for the curated dataset can be found in Table 1b.

**1/ Video Collection** To build a comprehensive dataset, we designed a collection methodology focused on capturing rich and diverse human interactions with physical user interfaces in realistic settings. The process spanned a wide array of common living and working environments, including devices from four main categories: (1) environmental controls (*e.g.*, light switches), (2) office equipment (*e.g.*, printers), (3) kitchen and home appliances (*e.g.*, coffee machines), and (4) building infrastructure (*e.g.*, elevator). A detailed list of the specific scenarios can be found in Table 1a.

Beyond ensuring fundamental video quality (clarity, stability, and a first-person perspective), the key principle of our methodology was the capture of complete, task-oriented interaction sequences over isolated, atomic actions. This approach ensures our data reflects the full spectrum of user behavior, which includes not only direct manipulations, but also crucial reasoning steps. Specifically, we recorded: **(1) UI Action:** Direct TCI element interactions like pressing, swiping, or rotating elements. A UI action example ("Press the *Down* button", when calling the elevator) is shown on the left side of Figure 3. **(2) Procedural Action:** Behaviors and camera view changes when a user is uncertain of action outcome, such as moving closer for inspection, visually scanning the interface, or pausing to wait for system feedback. An procedural action example ("Move away from the elevator", to see which one arrives) can be seen in the right side of Figure 3.

To guarantee the diversity of interaction styles and a minimum quality of collection, we utilized a team of 10 pre-trained data collectors. Their work produced a final (post-filtering) corpus of 193 high-quality video sequences, covering 508 fine-grained actions, forming the foundation of the SWITCH benchmark.

We also performed a survey of related datasets for data reuse. Unfortunately, only two such datasets — COIN, Tang et al. (2019) and HP-EPIC, Perrett et al. (2025) — included possibly reusable videos covering TCI scenarios. However, even those still presented issues and would require additional filtering and pre-processing for use in SWITCH. In the COIN dataset, angles change too

Table 2: Accuracy (%) results on SWITCH-Basic for different baselines, using image in questions and text or image answers. IT: image in question, text choices as answers, and II: image in question, image choices as answers.

Models	1/ Task-Aware VQA	2/ UI Compr.	3/ Action		4/ State Transition Prediction		5/ Verification		
	(a) UI State Recognition		IT	IT	IT	II	(a) Planning	(b) Expected State Prediction	
Question Format		IT	IT	IT	IT	II	IT	IT	II
Claude Sonnet 4, Anthropic (2025)	54.59	57.72	68.61	<b>75.00</b>	<b>58.54</b>	72.92	76.92	44.90	
Gemini 2.5 Flash, Comanici et al. (2025)	68.77	72.36	<b>73.14</b>	74.68	41.14	<b>81.25</b>	75.00	32.65	
Gemini 2.5 Pro, Comanici et al. (2025)	70.60	<b>74.80</b>	71.84	74.05	42.41	72.92	76.92	36.73	
Gemma3-4B, Gemma Team et al. (2025)	40.68	41.87	50.49	60.44	26.58	50.00	53.85	22.45	
Qwen3-VL-8B-Instruct, Yang et al. (2025a)	62.73	65.45	65.37	<b>68.67</b>	29.75	60.42	69.23	16.33	
Qwen3-VL-235B-Instruct, Yang et al. (2025a)	<b>70.87</b>	71.95	67.96	70.57	32.91	75.00	<b>80.77</b>	<b>48.98</b>	

Table 3: Accuracy (%) results on SWITCH-Basic for different baselines using video in either questions or answers. VT: video in question, text choices as answers, IV: image in question, video choices as answers, VV: video in question, video choices as answers, and VI: video in question, image choices as answers. Failure to return a response was counted as an incorrect prediction.

Models	1/ Task-Aware VQA	3/ Action			4/ State Transition Prediction		5/ Verification				
	(b) Goal-Oriented Reasoning	VT	IV	VV	VT	VI	(a) Planning			(b) Expected State Prediction	
Question Format		VT	IV	VV	VT	VI	VT	IV	VV	VT	VI
Gemini 2.5 Flash, Comanici et al. (2025)	73.33	66.45	59.80	44.33	<b>67.41</b>	50.32	76.00	<b>54.17</b>	39.58	69.23	36.73
Gemini 2.5 Pro, Comanici et al. (2025)	<b>74.81</b>	<b>69.08</b>	<b>60.78</b>	<b>48.33</b>	<b>67.41</b>	54.11	<b>78.00</b>	50.00	<b>56.25</b>	73.08	46.94
Qwen3-VL-8B-Instruct, Yang et al. (2025a)	62.22	55.92	40.52	18.67	63.61	51.58	52.00	33.33	18.75	67.31	44.90
Qwen3-VL-235B-Instruct, Yang et al. (2025a)	69.63	60.20	50.65	23.33	65.51	<b>59.49</b>	62.00	31.25	35.42	<b>75.00</b>	<b>48.98</b>

frequently (often cutting procedural actions), speed varies during videos, and many videos include textual instructions across video frames. In HP-EPIC, which focuses on kitchen tasks, multiple cases including appliances suffer from strong occlusion and all potential videos after filtering would still require segmentation by TCI-related action. Thus, in this iteration of the benchmark, we utilize only newly collected videos.

**2/ Data Annotation** Inspired by Ego4D’s goal-step in Song et al. (2023), our annotation process captures the fine-grained causal relationship between an agent’s actions and resulting changes in the world state. To achieve this, we model each interaction as a sequence of (pre-condition state, action, post-condition state) tuples, which explicitly links every action to its immediate outcome. **(a) State Annotation:** We provide a detailed description for each state, capturing both the **UI-State Transition** (e.g., indicator lights, on-screen text, panel settings) and the **Environment-State Transition** (e.g., brightness, or an open door). Crucially, we also annotate the evidential source for each state component, specifying how a human infers the state (e.g., from visual cues on the UI or observations of the physical world). **(b) Action Annotation:** Our hierarchical action annotation first distinguishes actions by their primary intent. They are categorized as either Task-Execution Actions or a distinct Verification category. Task-Execution Actions are the primary operations performed to make progress on the task. These are further subdivided into the two categories mentioned previously: **UI Actions** and **Procedural Actions** (examples of both types are shown in Figure 3). To account for real-world complexities, we further annotate actions for attributes such as error occurrence, corrective measures, and whether an action is optional for the task. Verification Actions are those aimed specifically at confirming a system’s status or the outcome of a prior action.

**3/ Question Processing** We then transformed our rich annotations into a structured question-answer (QA) benchmark encompassing the five previously mentioned tasks. For standardized and automated evaluation, we uniformly adopted a MCQ format. Each question consists of one correct answer, derived from our annotations, and several plausible distractors, randomly generated from available similar annotated cases.

Furthermore, to push models to perform genuine reasoning rather than relying on superficial environmental cues, we implemented a ground-truth substitution strategy, inspired by the methodology in WorldPrediction, Chen et al. (2025). For specific Tasks (3, 4, and 5), we alter the visual context of ground-truth by using sentence similarity. Additionally, for Tasks 3, 4, and 5, we designed questions that present context in diverse formats, including text, images, videos, and more complex combinations of images and videos. Finally, following all automated processing, the entire dataset was manually reviewed to ensure its final quality and correctness.

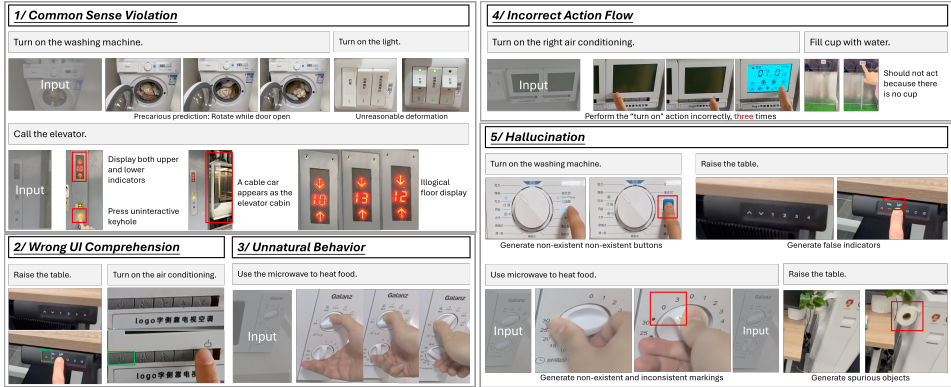


Figure 4: Case studies of representative errors in video generation. Examples using Veo3. Cropped image frames to focus on specific issues.

## 4 EXPERIMENTS

To validate the utility of our benchmark and systematically assess the capabilities of existing LMMs, we conduct an evaluation of open and closed models on SWITCH. Our primary goal is to provide a detailed analysis of existing LMMs, shedding light on their strengths and limitations in tasks requiring interactive intelligence. Beyond this, we also conduct some experiments with a leading video generation model to probe its world modeling ability.

### 4.1 EXPERIMENTAL SETUP

**Models Evaluated:** We selected a representative set of state-of-the-art LMMs, categorized into two groups:

1/ General-purpose proprietary models: These are leading closed-source models known for their powerful general reasoning abilities. We evaluated **Claude Sonnet 4**, Anthropic (2025); and both **Gemini 2.5 Flash** and **Pro**, Comanici et al. (2025). Claude Sonnet 4 accepts text and static images as input, while Gemini 2.5 accepts video inputs in addition to text and images.

2/ General-purpose open models: To represent the frontier of publicly available research, we evaluated **Gemma3-4B**, Gemma Team et al. (2025); and both **Qwen3-VL-8B-Instruct** and **235B-Instruct**, Yang et al. (2025a). Gemma3 is a free-weight model with the same structure as Gemini, but but supports only text and image inputs. While a key advantage of Qwen3-VL is its native support for text, image, and video inputs, making it well-suited for full range of tasks in SWITCH.

**Evaluation Protocol:** Each model was evaluated on all tasks compatible with its input modalities. Consequently, Claude Sonnet 4 and Gemma3-4B were benchmarked on image-based tasks, while Gemini 2.5 series and Qwen3-VL series were evaluated across the full suite of tasks, including those with video. For each question, we provided the model with the multimodal context (text, image, and/or video) and the multiple-choice questions in a zero-shot setting. The model’s task was to select the correct answer from the provided options. Illustrative examples of our question-answer pairs and their prompts are provided in Figure 1.

**Implementation Details:** We use accuracy as the primary evaluation metric and we evaluate each baseline official implementation or API. All evaluation scripts and annotations are to be publicly released for transparency.

### 4.2 OBSERVATIONS AND CHALLENGES

The overall performance results of the evaluated models across tasks are summarized in Tables 2 and 3<sup>4</sup>. To investigate the root causes of the observed performance degradation, we conducted a

<sup>4</sup>Regarding Task 5 (limited to  $N = 102$  samples), we note two specific patterns: (1) Claude Sonnet 4 and Gemini 2.5 Pro achieve identical accuracy (Table 2 for Tasks 5.a and 5.b), but fail differently (dynamic details

detailed qualitative analysis of failure cases<sup>5</sup>. We structure our findings into a progressive hierarchy of cognitive deficiencies, ranging from fundamental state estimation to complex situated adjustment.

**1/ Weakness in Fine-grained State Estimation.** The primary source of error (52.97%) lies in the precise identification of physical states. While models generally succeed in recognizing target objects (*e.g.*, identifying a device or switch), they exhibit significant instability in determining their fine-grained statuses. This instability contributes to performance degradation when adding additional visual context (*e.g.*, Task 3 Qwen3-VL-8b-Instruct drops from 65.37% (IT) to 40.52% (IV), Task 4 Gemini 2.5 Pro drops from 74.05% (IT) to 42.41% (II)).

**(1) General State Misclassification:** We observe frequent errors where the model correctly locates the object but fails to distinguish its current physical mode. As shown in Figure 7 and 8, models frequently misjudge basic states, such as determining whether water is *flowing* versus *static*.

**(2) Task-Induced Hallucination:** A more critical pattern involves models being biased by textual instructions. As shown in Figure 9, when prompted to "*Turn on the light*," the model erroneously categorizes a "Turn off" action as "Turn on," ignoring the visual contradiction in the switch's position. This highlights a deficiency in grounding binary states under the effect of instruction bias.

**2/ Deficiencies in Causal Reasoning and Device Logic.** A distinct subset of failures (11.39%) arises when models struggle with the *causal logic* of specific devices, often confusing visual correlations with functional causality. In elevator navigation scenarios (Figure 11), models frequently interpret a "lit display arrow"—a feature merely correlated with movement—as a signal that the elevator has been "called." Conversely, they ignore the "lit button," which is the true indicator of a registered request. Consequently, models incorrectly predict "Wait" rather than "Press button", demonstrating a superficial understanding of device mechanics.

**3/ Failure in Adaptive Adjustment and Situated Grounding.** A significant portion of failures (35.64%) occurs when models possess both correct perception and general common sense but fail to *adaptively adjust* their policy based on specific environmental constraints or viewpoints.

**(1) Physical Affordance Mismatch:** Models often apply generic solutions without considering physical accessibility. In Figure 12, the model suggests "pressing the power button" for an air conditioner—a valid action for a standing unit—despite visual input showing a wall-mounted unit installed out of reach. The model fails to switch to the feasible action: "finding a remote control."

**(2) Viewpoint Adaptation:** Models lack the meta-cognitive ability to recognize when their current viewpoint is insufficient. In elevator waiting scenarios (Figure 13), where the field of view is narrow or obstructed, the optimal policy is to "move back" to check availability. However, models typically default to a passive "wait", failing to actively adjust their position to a better visual information.

#### 4.3 PROBING WORLD MODELING IN GENERATIVE MODELS

To push beyond discriminative QA and its limitations, and to directly probe the internal world models of generative systems (*e.g.*, Fung et al. (2025); Ali et al. (2025); Rivard et al. (2025)), we conducted further experiments on a state-of-the-art video generation model. We selected Google's Veo3, see Google (2025), as a representative model and tested its ability to simulate physical-temporal TCI-related dynamics in two settings:

**State Transition Prediction:** Analogous to Task 4, the model is provided with an initial frame and a textual description of an action, tasking it to generate the subsequent frames depicting the outcome. For instance, in the top-left example in Figure 4, when given an image of a washing machine and the instruction "Turn on the washing machine", Veo3 generates a coherent sequence of future frames.

**Action Generation:** Analogous to Task 3, the model is given the initial and final frames of the state change with a high-level action description. The task is to generate the intermediate video showing the specific action that connects them. As an example, for the microwave shown in the bottom-right

vs. spatial reasoning); (2) Gemini 2.5 Flash outperforms Pro (Table 2 for Task 5.a) by avoiding Pro's over-complication of verification logic and bias toward physical actions (*e.g.*, "move") over passive checks (*e.g.*, "wait"), instead relying on more direct reasoning.

<sup>5</sup>Due to space constraints, detailed examples are provided in Appendix D.

corner of Figure 4, the model receives the first and last frames along with the action "Use microwave to heat food" and is expected to generate the intermediate frames.

Analysis of the generated results reveals several fundamental error patterns, indicating limitations in the model’s physical and causal understanding (see Figure 4 for examples):

**1/ Common Sense Violation:** A frequent failure mode is the model’s inability to maintain physical realism and logical consistency, leading to the generation of implausible and nonsensical scenarios. As shown in the top-left of Figure 4, these violations can be categorized as follows: **(1) Violations of Physical Constraints:** The model disregards fundamental real-world principles. For example, it generates a video of a washing machine’s drum rotating while the door is wide open—a clear breach of physical safety interlocks. **(2) Violations of Physical Inconsistency:** The model fails to maintain the physical consistency of objects. In the light switch example, it generates an "Unreasonable deformation" of the solid switch panel. **(3) Violations of Functional Logic:** The model produces functional impossible outcomes. The elevator case exhibits multiple such failures: illuminating both the up and down indicators simultaneously, generating a nonsensical scene where a cable car appears inside the elevator cabin, and displaying floor numbers in a non-sequential illogical order.

**2/ Incorrect UI Comprehension:** The model fails to correctly interact with the appropriate UI elements. This includes attempting to raise a table or turn on an air conditioner by pressing an incorrect area on the respective control panels (shown in the bottom-left cases of Figure 4).

**3/ Unnatural Behavior:** The model performs actions in a manner that is inefficient, or unlike how a human would. The bottom-left example of "3/ Unnatural Behavior" in Figure 4 shows the use of a microwave by trying to rotate the two knobs at the same time with one hand.

**4/ Incorrect Action Flow:** The model executes a logically incoherent sequence of actions or violates necessary preconditions, such as performing the "Turn on" action for an air conditioner incorrectly *three* times or dispensing water from a machine with no cup present (even if explicitly mentioned in the instruction) in the top-right cases of Figure 4.

**5/ Hallucination:** The model exhibits a tendency to fabricate visual content that does not exist in the input and contradicts physical reality. As shown in the bottom-right cases of Figure 4, these hallucinations can be categorized as follows: **(1) Fabrication of UI Elements:** Generating non-existent markings or indicators. For example, it generates inconsistent and non-existent markings on a washing machine’s dial, creates a false indicator "up" on a standing desk’s control panel, and adds phantom numbers to a microwave’s timer knob (which are also inconsistent). **(2) Generation of Spurious Objects:** Introducing contextually irrelevant objects into the scene. A clear example is the generation of a roll of toilet paper attached to an office air purifier when the task prompt was to "Raise the table", or transforming an air-purifier into a printer as the camera view pans.

## 5 CONCLUSION

In this work, we introduce SWITCH, a benchmark effort designed to evaluate how models perceive, reason about, and act upon semantic TCI interactions and their effects in real-world contexts. Unlike prior datasets and benchmarks that focus either on fully virtual simulations or object-centric physical interactions only, SWITCH targets a critical but overlooked scenario in embodied AI — the interactive layer between humans and smart environments and their appliances/devices, mediated through TCIs (*i.e.*, switches, buttons, and control panels). Through five systematically designed tasks, SWITCH examines complementary dimensions of embodied interactive intelligence — from perceiving UI semantics and generating purposeful actions, to predicting causal state transitions and verifying task outcomes. Our evaluation across multiple mainstream LMMs reveals consistent weaknesses in understanding affordances, anticipating environmental effects, and verifying causal consequences, highlighting key challenges toward achieving well grounded world models.

By open-sourcing SWITCH-Basic and its evaluation toolkit, we aim to provide a standardized, extensible platform for studying perception-action reasoning in multimodal embodied agents. The current release of SWITCH represents only the first phase, and future iterations will expand dataset scale and diversity, incorporate richer TCI interactions, multi-step and verification scenarios, and broader evaluation metrics to assess models’ world modeling, reasoning, and adaptability. For further discussion on evolving the benchmark, please see Appendix E.

## REFERENCES

- Arslan Ali, Junjie Bai, Maciej Bala, Yogesh Balaji, Aaron Blakeman, Tiffany Cai, Jiaxin Cao, Tianshi Cao, Elizabeth Cha, Yu-Wei Chao, et al. World simulation with video foundation models for physical ai. *arXiv preprint arXiv:2511.00062*, 2025.
- Anthropic. Introducing Claude4, 2025. URL <https://www.anthropic.com/news/claude-4>.
- Zefan Cai, Haoyi Qiu, Tianyi Ma, Haozhe Zhao, Gengze Zhou, Kung-Hsiang Huang, Parisa Kordjamshidi, Minjia Zhang, Xiao Wen, Jiuxiang Gu, et al. Mmgr: Multi-modal generative reasoning. *arXiv preprint arXiv:2512.14691*, 2025.
- DeLong Chen, Willy Chung, Yejin Bang, Ziwei Ji, and Pascale Fung. Worldprediction: A benchmark for high-level world modeling and long-horizon procedural planning. *arXiv preprint arXiv:2506.04363*, 2025.
- Wei Chow, Jiageng Mao, Boyi Li, Daniel Seita, Vitor Guizilini, and Yue Wang. Physbench: Benchmarking and enhancing vision-language models for physical world understanding. *arXiv preprint arXiv:2501.16411*, 2025.
- Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.
- Mengfei Du, Binhao Wu, Zejun Li, Xuan-Jing Huang, and Zhongyu Wei. Embspatial-bench: Benchmarking spatial understanding for embodied tasks with large vision-language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 346–355, 2024.
- Haoyi Duan, Hong-Xing Yu, Sirui Chen, Li Fei-Fei, and Jiajun Wu. Worldscore: A unified evaluation benchmark for world generation. *arXiv preprint arXiv:2504.00983*, 2025.
- Aaron Foss, Chloe Evans, Sasha Mitts, Koustuv Sinha, Ammar Rizvi, and Justine T. Kao. CausalVQA: A physically grounded causal reasoning benchmark for video models. *arXiv preprint arXiv:2506.09943*, 2025. URL <https://arxiv.org/abs/2506.09943>.
- Pascale Fung, Yoram Bachrach, Asli Celikyilmaz, Kamalika Chaudhuri, DeLong Chen, Willy Chung, Emmanuel Dupoux, Hongyu Gong, Hervé Jégou, Alessandro Lazaric, et al. Embodied ai agents: Modeling the world. *arXiv preprint arXiv:2506.22355*, 2025.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025.
- Google. Veo: a text-to-video generation system, 2025. URL <https://storage.googleapis.com/deepmind-media/veo/Veo-3-Tech-Report.pdf>.
- Hiroshi Ishii and Brygg Ullmer. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI'97)*. ACM, 1997. ISBN 0897918029. doi: 10.1145/258549.258715. URL <https://doi.org/10.1145/258549.258715>.
- Chengshu Li, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-Martín, Chen Wang, Gabrael Levine, Michael Lingelbach, Jiankai Sun, et al. Behavior-1k: A benchmark for embodied ai with 1,000 everyday activities and realistic simulation. In *Conference on Robot Learning*, pp. 80–93. PMLR, 2023.
- Dacheng Li, Yunhao Fang, Yukang Chen, Shuo Yang, Shiyi Cao, Justin Wong, Michael Luo, Xiaolong Wang, Hongxu Yin, Joseph E. Gonzalez, et al. Worldmodelbench: Judging video generation models as world models. *arXiv preprint arXiv:2502.20694*, 2025.

- Xiao Liu, Tianjie Zhang, Yu Gu, Iat Long Iong, Yifan Xu, Xixuan Song, Shudan Zhang, Hanyu Lai, Xinyi Liu, Hanlin Zhao, et al. Visualagentbench: Towards large multimodal models as visual foundation agents. *arXiv preprint arXiv:2408.06327*, 2024.
- Saman Motamed, Laura Culp, Kevin Swersky, Priyank Jaini, and Robert Geirhos. Do generative video models understand physical principles? *arXiv preprint arXiv:2501.09038*, 2025.
- Toby Perrett, Ahmad Darkhalil, Saptarshi Sinha, Omar Emara, Sam Pollard, Kranti Kumar Parida, Kaiting Liu, Prajwal Gatti, Siddhant Bansal, Kevin Flanagan, et al. Hd-epic: A highly-detailed egocentric video dataset. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 23901–23913, 2025.
- Li Puyin, Tiange Xiang, Ella Mao, Shirley Wei, Xinye Chen, Adnan Masood, Li Fei-Fei, and Ehsan Adeli. Quantiphy: A quantitative benchmark evaluating physical reasoning abilities of vision-language models. *arXiv preprint arXiv:2512.19526*, 2025.
- Christopher Rawles, Sarah Clinckemaille, Yifan Chang, Jonathan Waltz, Gabrielle Lau, Marybeth Fair, Alice Li, William Bishop, Wei Li, Folawiyi Campbell-Ajala, et al. Androidworld: A dynamic benchmarking environment for autonomous agents. *arXiv preprint arXiv:2405.14573*, 2024.
- Luke Rivard, Sun Sun, Hongyu Guo, Wenhui Chen, and Yuntian Deng. Neuralos: Towards simulating operating systems via neural generative models. *arXiv preprint arXiv:2507.08800*, 2025.
- Yale Song, Eugene Byrne, Tushar Nagarajan, Huiyu Wang, Miguel Martin, and Lorenzo Torresani. Ego4d goal-step: Toward hierarchical understanding of procedural activities. *Advances in Neural Information Processing Systems*, 36:38863–38886, 2023.
- Keith Stouffer, Michael Pease, CheeYee Tang, Timothy Zimmerman, Victoria Pillitteri, Suzanne Lightman, Adam Hahn, Stephanie Saravia, Aslam Sherule, and Michael Thompson. Guide to Operational Technology (OT) Security. Technical Report NIST Special Publication (SP) 800-82, Rev. 3, National Institute of Standards and Technology, 2023.
- Weihao Tan, Wentao Zhang, Xinrun Xu, Haochong Xia, Gang Ding, Boyu Li, Bohan Zhou, Junpeng Yue, Jiechuan Jiang, Yewen Li, Ruyi An, Molei Qin, Chuqiao Zong, Longtao Zheng, Yujie Wu, Xiaoqiang Chai, Yifei Bi, Tianbao Xie, Pengjie Gu, Xiyun Li, Ceyao Zhang, Long Tian, Chaojie Wang, Xinrun Wang, Börje F. Karlsson, Bo An, Shuicheng Yan, and Zongqing Lu. Cradle: Empowering foundation agents towards general computer control. In *Forty-second International Conference on Machine Learning, ICML 2025*, 2025.
- Yansong Tang, Dajun Ding, Yongming Rao, Yu Zheng, Danyang Zhang, Lili Zhao, Jiwen Lu, and Jie Zhou. Coin: A large-scale dataset for comprehensive instructional video analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1207–1216, 2019.
- Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018.
- Qineng Wang, Wenlong Huang, Yu Zhou, Hang Yin, Tianwei Bao, Jianwen Lyu, Weiyu Liu, Ruohan Zhang, Jiajun Wu, Li Fei-Fei, et al. Enact: Evaluating embodied cognition with world modeling of egocentric interaction. *arXiv preprint arXiv:2511.20937*, 2025.
- Thaddäus Wiedemer, Yuxuan Li, Paul Vicol, Shixiang Shane Gu, Nick Matarese, Kevin Swersky, Been Kim, Priyank Jaini, and Robert Geirhos. Video models are zero-shot learners and reasoners. *arXiv preprint arXiv:2509.20328*, 2025.
- Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh J Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, et al. Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments. *Advances in Neural Information Processing Systems*, 37:52040–52094, 2024.
- Xinrun Xu, Pi Bu, Ye Wang, Börje F. Karlsson, Ziming Wang, Tengtao Song, Qi Zhu, Jun Song, Zhiming Ding, and Bo Zheng. Deepphy: Benchmarking agentic vlms on physical reasoning. *arXiv preprint arXiv:2508.05405*, 2025. URL <https://arxiv.org/abs/2508.05405>.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujiu Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025a.

Rui Yang, Hanyang Chen, Junyu Zhang, Mark Zhao, Cheng Qian, Kangrui Wang, Qineng Wang, Teja Venkat Koripella, Marziyeh Movahedi, Manling Li, et al. Embodiedbench: Comprehensive benchmarking multi-modal large language models for vision-driven embodied agents. *arXiv preprint arXiv:2502.09560*, 2025b.

Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. Webshop: Towards scalable real-world web interaction with grounded language agents. *Advances in Neural Information Processing Systems*, 35:20744–20757, 2022.

Jiahao Zhang, Muqing Jiang, Nanru Dai, Taiming Lu, Arda Uzunoglu, Shunchi Zhang, Yana Wei, Jiahao Wang, Vishal M Patel, Paul Pu Liang, et al. World-in-world: World models in a closed-loop world. *arXiv preprint arXiv:2510.18135*, 2025a.

Zixin Zhang, Kanghao Chen, Xingwang Lin, Lutao Jiang, Xu Zheng, Yuanhuiyi Lyu, Litao Guo, Yinchuan Li, and Ying-Cong Chen. Phystoolbench: Benchmarking physical tool understanding for mllms. *arXiv preprint arXiv:2510.09507*, 2025b.

## A EXTENDED RELATED WORK

From the perspective of lifelong agents, evaluating large multimodal models and world models is not merely about measuring predictive accuracy or short-horizon task success. Lifelong agents are expected to operate continuously in dynamic environments. However, human environments are fundamentally shaped by physical tangible interfaces beyond basic tool affordances. These systems—ranging from switches and buttons to displays and smart appliances—encode human intent and regulate the spaces we live and work in. While a wide range of benchmarks has emerged—spanning perception, reasoning, action planning, tool use, and interface control Du et al. (2024); Yang et al. (2025b); Liu et al. (2024); Xu et al. (2025); Zhang et al. (2025b)—many lack the combined focus on semantic user interfaces, contextual control, and post-action verification that real-world interaction demands. We divide existing efforts into two main categories — *Multimodal & World-Model Evaluation* and *User Interface Control Interaction* — and discuss how our benchmark uniquely addresses the gaps.

### A.1 MULTIMODAL & WORLD-MODEL EVALUATION

**1/ Multimodal Evaluation** Recent benchmarks have progressively evolved from assessing static perception to evaluating dynamic reasoning and planning. For instance, the work on Phys-Bench Chow et al. (2025) introduces a foundation by testing physical-world understanding across object properties, relationships, and scene dynamics. Transitioning from static knowledge to temporal dynamics, subsequent works have focused on agent-centric world modeling. WorldPrediction Chen et al. (2025) targets more high-level abstraction, evaluating capabilities in both action choosing and long procedural planning. EnACT Wang et al. (2025) further formalizes this through forward (predicting future states) and inverse (inferring past actions) world modeling. This highlights that robust understanding requires sequencing actions and states, not just pattern matching. Furthermore, QuantiPhy Puyin et al. (2025) identifies a critical deficiency in current models regarding quantitative reasoning (*e.g.*, velocity, dimensions), advocating for metrics that go beyond qualitative descriptions. In parallel, large-scale embodied benchmarks like BEHAVIOR-1K Li et al. (2023) have expanded the scope to human-centered activities across diverse scenes, covering 1,000 everyday activities across 50 scenes with rich object annotations. Although BEHAVIOR-1K enhances scene diversity and long-horizon manipulation, it remains object-manipulation focused and does not systematically target semantic interface control, TCI-to-environment causality, or verification of control-action outcomes. These works provide strong foundations but leave interface-control reasoning and outcome verification under-explored.

**2/ World-Model Evaluation** The evaluation of world models has traditionally relied on visual fidelity metrics (*e.g.*, FVD Unterthiner et al. (2018)). WorldModelBench Li et al. (2025) concentrates on video generation models, with a special focus on nuanced world modeling violations and alignment with human preferences. Meanwhile, WorldScore Duan et al. (2025) proposes a unified benchmark for world generation and scene prediction—next-scene tasks driven by camera trajectories and layout dynamics. It centers on scene generation and layout controllability but overlooks interactive device/appliance control via their interfaces. A paradigm shift is currently occurring towards evaluating physical and logical compliance. Benchmarks such as Physics-IQ Motamed et al. (2025) and MMGR Cai et al. (2025) now evaluate generated content based on adherence to physical laws (*e.g.*, fluid dynamics) and multi-modal logical reasoning. Additionally, world models also exhibit capabilities beyond mere pixel generation, even simulating tool use within zero-shot settings Wiedemer et al. (2025). Meanwhile, World-in-World Zhang et al. (2025a) argues that "Visual Quality does not equal Task Success," using the world model as a simulator and focusing more on task success rates.

### A.2 USER INTERFACE CONTROL INTERACTION

A second line of research investigates agents controlling digital interfaces or tools via screen-based or scripted simulation environments. Benchmarks such as WebShop Yao et al. (2022), Android-World Rawles et al. (2024), Cradle Tan et al. (2025), and OSWorld Xie et al. (2024), require agents to click buttons, navigate GUIs, or perform scripted operations from textual instructions. These tasks advance procedural reasoning and interface grounding, but remain within virtual or digital regimes and rarely involve real-world UI elements or their downstream environmental effects.

In summary, despite significant progress in world modeling and interactive agents, existing benchmarks typically focus on either purely digital interactions or physical interactions with objects/tools (e.g., PhysToolBench Zhang et al. (2025b)). They largely overlook the critical domain of physical Tangible Control Interface (TCI) interactions. For AI to transition from passive perception toward truly embodied assistance, it must gain the ability to interpret and operate these widespread designed affordances. Mastering “physical UI” interactions represents a crucial milestone in realizing AI agents that can meaningfully collaborate with humans. Our benchmark—SWITCH—fills this gap by focusing explicitly on TCIs that combine UI semantics, physical device states, and user intent. As shown in Figure 5, SWITCH addresses the challenges of semantic understanding of capabilities, grounding of the interface, action planning, effect causality (including safety issues), and result verification. By evaluating a model’s ability to perceive, manipulate, and verify real-world UI-based interactions, we target a key capability for interactive and embodied multimodal intelligence.

## B ADDITIONAL QUALITATIVE EXAMPLES

Figure 6 illustrates a complex, multi-step task (printing a document) to demonstrate complex interaction reasoning. Unlike the single-step task shown in the main text, this example requires sequential Semantic UI Comprehension to interpret changing interface layouts, and the execution of corresponding Action Generation steps. Throughout the process, the model observes state transitions as the UI updates after each interaction, requiring adaptive planning based on visual context. Finally, through Result Verification Planning, the agent repositions to check the printed output, completing the full perception–action–verification loop in a dynamic, real-world setting.

## C SUCCESSFUL RESPONSE RATES OF BASELINES WITH VIDEOS

Table 4: Percentage of questions for which a valid response was successfully returned. VT: video in question, text choices as answers, IV: image in question, video choices as answers, VV: video in question, video choices as answers, and VI: video in question, image choices as answers.

Models	1/ Task-Aware VQA		3/ Action			4/ State Transition Prediction		5/ Verification				
	(b) Goal-Oriented Reasoning		VT	IV	VV	VT	VI	(a) Planning			(b) Expected State Prediction	
	VT							VT	IV	VV	VT	VI
Gemini 2.5 Flash Comanici et al. (2025)	100.00		100.00	99.02	98.67	100.00	100.00	100.00	100.00	100.00	100.00	100.00
Gemini 2.5 Pro Comanici et al. (2025)	100.00		100.00	98.69	99.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
Qwen3-VL-8B-Instruct Yang et al. (2025a)	100.00		100.00	97.71	97.67	100.00	100.00	100.00	100.00	100.00	100.00	100.00
Qwen3-VL-235B-Instruct Yang et al. (2025a)	100.00		100.00	97.71	97.67	100.00	100.00	100.00	100.00	100.00	100.00	100.00

Table 4 presents the successful response rates (%) of various baseline models when evaluated on video-based questions. Each value indicates the percentage of questions for which the model was able to return a valid response. This table provides additional details on model robustness and response consistency in the video setting.

## D CASE STUDY EXAMPLES

In this section, we provide detailed visualizations of the failure cases and organize these examples according to the three primary cognitive deficiencies discussed in the Section 4.2.

**1/ Weakness in Fine-grained State Estimation:** Models struggle to identify precise physical states despite correctly locating objects.

- **General State Misclassification:** Figure 7 and Figure 8 illustrate the difficulty in distinguishing dynamic fluid states (flowing vs. static), leading to incorrect action choices.
- **Task-Induced Hallucination:** Figure 9 demonstrates how instruction bias (e.g., "turn on") overrides visual evidence, causing the model to hallucinate an "on" state even when the switch is visually being turned "off". Similarly, Figure 10 shows multiple models failing to identify the functional state of a light switch.

**2/ Deficiencies in Causal Reasoning and Device Logic:** Models often confuse visual correlations with functional causality.

- **Elevator Logic:** Figure 11 shows Claude Sonnet 4 and Gemini 2.5 Flash misinterpreting a "lit display arrow" as a signal that the elevator is called, ignoring the "lit button" which is the true causal indicator.
- **State Transition Prediction:** Figure 14 (Printer) and Figure 15 (AC Fan) highlight failures in predicting the key immediate visual outcome of an action, such as paper appearing in the output tray or the fan speed indicator changing.

**3/ Failure in Adaptive Adjustment and Situated Grounding:** Models fail to adapt generic knowledge to specific environmental constraints.

- **Physical Affordance Mismatch:** Figure 12 reveals a critical failure where top-tier models (Gemini 2.5 Pro/Flash, Qwen3-VL) suggest "pressing the power button" on a wall-mounted AC unit that is physically out of reach, ignoring the need for a remote control.
- **Viewpoint Adaptation:** Figure 13 demonstrates a lack of meta-cognitive awareness, where models fail to recognize the need for an intermediate "adjust view" action to gain better visual information before proceeding.

## E BENCHMARK EVOLUTION

The current release of SWITCH represents only the first phase of our benchmark. In the upcoming stages, we plan to expand both the scale and diversity of the dataset. Specifically, we aim to increase the coverage of TCI-centric environments by incorporating a broader range of device interaction scenarios, as well as multiple variants of similar devices with distinct layouts, control mechanisms, and capabilities.

Crucially, we will also include evaluation beyond the current discriminative QA format, enhance the evaluation dimensions of existing tasks, and propose metrics to more precisely evaluate models' world modeling and ground abilities. In Task-Oriented Visual Question Answering, for example, future extensions will assess models' ability to perform context-aware inference — evaluating whether models can infer (TCI and world) states from indirect environmental cues when the interface itself is partially occluded or not directly visible.

As the collected dataset grows, we intend to partition it into training, validation, and test splits, enabling the community to use SWITCH not only for evaluation but also for training or adaptation of novel models.

Moreover, since many real-world interactions require multi-step operations before their outcome can be verified, the current benchmark data includes fewer Result Verification samples compared to other tasks. As we expand it, the number and diversity of verification examples should increase accordingly, allowing more comprehensive assessment of models' ability to reason about important delayed or multi-stage effects.

A focus point for next iterations is also to cover cases of adapting to incorrect actions and adapting to retry. SWITCH-Basic currently only covers 5 instances of situations where an interaction needs to be canceled or restarted.

Finally, we will introduce multi-level difficulty settings across tasks, designed to more comprehensively evaluate models' reasoning depth, generalization, and adaptability. We envision SWITCH evolving into a collaborative progressively richer and more diagnostic-oriented benchmark for studying the foundations of interactive and embodied intelligence in multimodal models.

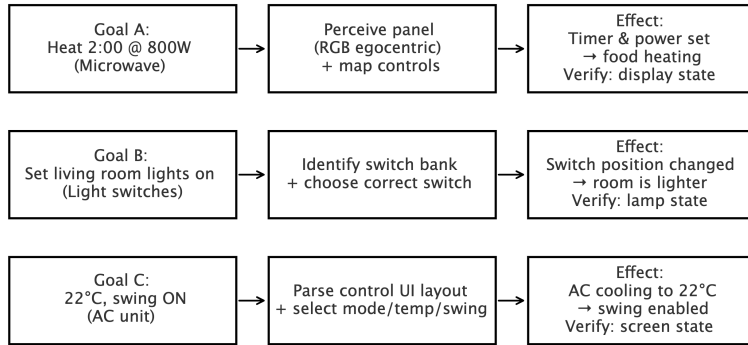


Figure 5: Example tasks modeled in SWITCH. Device Instruction (Goal) - Interface Understanding & Action - Effect / Verification.

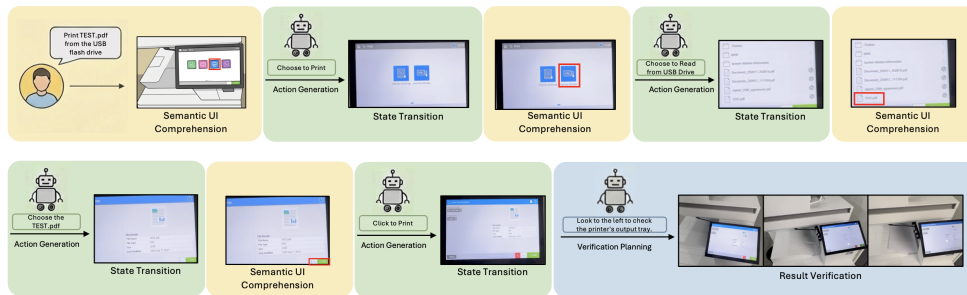



Figure 6: Example of a multi-step document printing task demonstrating complex interaction reasoning. Given the a user instruction, the agent must perform sequential Semantic UI Comprehension to interpret changing interface layouts, and execute corresponding Action Generation steps (A1–A4). Throughout the process, the model observes state transitions as the UI updates after each interaction, requiring adaptive planning based on visual context. Finally, through Result Verification Planning, the agent repositions to check the printed output, completing the full perception–action–verification loop in a dynamic, real-world setting.

**Query:**  
 Following the content of the image input, which of the four provided options correctly represents the next action in the flow of "Get water".  
 A. Check if the cup is filled.  
 B. Stop the dispenser.  
 C. Find the button.  
 D. Move to the dispenser.

**Input Image:**



**Ground Truth: D**

**qwen3\_vl**

Option A ("Check if the cup is filled") is the most immediate and logical next step while water is still being dispensed.

Option B ("Stop the dispenser") would come after confirming the cup is filled.

Option C ("Find the button") is irrelevant here – the dispenser is already operating, so the button has already been found and pressed.

Option D ("Move to the dispenser") is also irrelevant – the cup is already in place under the dispenser.


Thus, the correct next action is to check if the cup is filled.

### Answer  
 A

Figure 7: An example of Qwen3-VL-235B-Instruct on Task 3/ Action Generation in question format: Image in question, text as option. The model failed to recognize that "The water is not flowing down", thus giving the incorrect answer.

**Query:**  
 Following the content of the image input, which of the four provided options correctly represents the next action in the flow of "Get purified water".  
 A. Move to the dispenser.  
 B. Stop the water dispenser.  
 C. Look down to see if the water flow down ,and if the cup is full.  
 D. Turn on the dispenser.

**Input Image:**



**Ground Truth: B**

**Model Outputs (All Wrong)**

**claude\_sonnet\_4**


## Thinking Process  
 Looking at the image, I can see a water dispenser with a black touchscreen panel displaying "净化水" (Purified Water). There are two dispensing spouts visible at the bottom. The dispenser appears to be in a ready state with the purified water option selected on the screen, but no water is currently flowing and there's no visible cup positioned under the spouts yet.

In the typical flow of getting purified water:  
 1. Approach/move to the dispenser  
 2. Select the water type (purified water is already shown selected)  
 3. Position a cup (not visible yet)  
 4. Activate/turn on the dispenser  
 5. Monitor the water flow and cup level  
 6. Stop the dispenser when done

Figure 8: An example of Claude Sonnet 4 on Task 3/ Action Generation in question format: Image in question, text as option. The model failed to recognize that "The water is flowing down", thus chose the wrong answer "D. Turn on the dispenser."

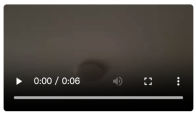
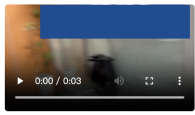

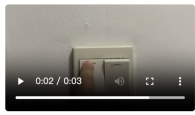
**Query:**  
Answer the question based on the input image (current state) and the four video response options, marked as A/B/C/D. Following the content of the first image input, which of the four provided video options correctly represents the next action in the flow of "Turn on the light in the kitchen".

**Input:**



**Options (A/B/C/D):**

**A** **B** **C** **D**



**Ground Truth: B**

**qwen3\_8b**

Looking at the video options:

- Video A shows a double light switch being pressed, which would turn on lights.
- Video B shows a single light switch being pressed, which would also turn on lights.
- Video C shows a conference room with lights already on, which doesn't represent turning on lights.
- Video D shows a ceiling light fixture, but no action is taken to turn it on.

Since both Video A and Video B show someone pressing a light switch (which is the typical way to turn on lights), they both represent valid next actions. Video C shows lights already on, so it's not the next action. Video D doesn't show any action being taken to turn on the light.

Therefore, the correct answer should be either A or B, as both represent turning on lights by pressing a switch.

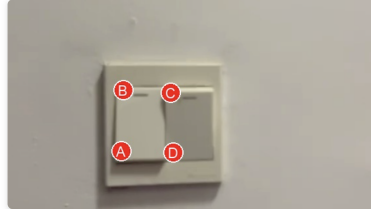
**## Answer**  
A or B

Figure 9: An example of Qwen3-VL-8B-Instruct on Task 3/ Action Generation in question format: Image in question, videos as option. The model mistakenly identified the actions in the first two options as turning on the light, leading to the wrong answer.

**Query:**

Which of the marked points in the image can be interacted with to "Turn off the light"? Choose only one option.

Input Image:



Ground Truth: C

## ✗ Model Outputs (All Wrong)

claude\_sonnet\_4

## ### Thinking Process

This appears to be a dual light switch plate. In a typical rocker-style light switch:

- The top portion (B and C) represents the "ON" position
- The bottom portion (A and D) represents the "OFF" position

To turn off a light, you would press the bottom part of the rocker switch. Since the question asks about turning off "the light" (singular), and points A and D are both in the bottom/OFF positions of their respective switches, either could work. However, point A is on the left switch, which is typically the primary switch position in a dual-switch configuration.

## ### Answer

A

gemini\_25\_flash

## ### Thinking Process

- **Right Switch (with points C and D):** The top part (C) is pressed in, and the bottom part (D) is sticking out. This is the typical 'off' position. Pressing the top part (C) will do nothing as it is already depressed. Pressing the bottom part (D) would turn this switch 'on'.

4. The question asks which point can be interacted with to "Turn off the light".

5. Based on the analysis:

- Interacting with **A** will do nothing, as the left switch is already in this position.
- Interacting with **B** will press the top of the left switch, toggling it from 'on' to 'off'. This matches the goal.
- Interacting with **C** will do nothing, as the right switch is already in this position.
- Interacting with **D** will press the bottom of the right switch, toggling it from 'off' to 'on'. This is the opposite of the goal.

6. Therefore, the only point that can be interacted with to turn a light off is B.

## ### Answer

B

gemma

## ### Thinking Process

The image shows a double switch. To turn off the light, you need to flip one of the switches to the "off" position. Looking at the image, points A and D are the switches. Since the question asks for only one option, and both A and D are switches, I'll choose the one that appears to be more clearly a switch.

## ### Answer


D

Figure 10: An example of Claude Sonnet 4, Gemini 2.5 Flash and Gemma3-4B on Task 2/ UI Comprehension in question format: Image in question, text as option. The model fails to identify that only the right light switch is on and only the upper part of the light switch is functional, thus choosing an incorrect answer.

**Query:**  
 Following the content of the image input, which of the four provided options correctly represents the next action in the flow of "Take the elevator".

A. Look up to see the elevator screen.  
 B. Wait until the elevator arrive 5 floor.  
 C. Wait until the elevator arrive.  
 D. Select to go up.

Input Image:



Ground Truth: D

✗ Model Outputs (All Wrong)

**claude\_sonnet\_4**

### Thinking Process

Looking at the image, I can see an elevator control panel with:

- A digital display showing an upward arrow (↕) and the number "1"
- Two call buttons below the display (up and down arrows)
- The display is already active and showing information

The illuminated display with the up arrow indicates that the elevator has already been called (the up button has already been pressed). The current floor appears to be 1, and the system is active.

Given this state, the person has already selected their direction (up), so the logical next step in the flow of "Take the elevator" would be to wait for the elevator car to arrive at the current floor before they can enter it.

**gemini\_25\_flash**


### Thinking Process

1. **Analyze the image:** The image shows an elevator call panel next to an elevator door. The panel has a digital display and two call buttons (up and down).
2. **Interpret the display:** The digital display shows a large arrow pointing upwards and the number "1". This signifies that the elevator is currently at the 1st floor and is moving in the upward direction.
3. **Evaluate the context:** The overall scenario is "Take the elevator". The image captures a moment \*after\* the elevator has been called or while it is in transit. The user is observing the elevator's status.
4. **Assess the options:**
  - \* **A. Look up to see the elevator screen:** The image itself is a view of the elevator screen. This action is what is currently happening, not the *next* action.
  - \* **B. Wait until the elevator arrive 5 floor:** The display shows the elevator is at floor 1. There is no information to suggest the user is on or going to the 5th floor. This is an unsupported and overly specific assumption.
  - \* **C. Wait until the elevator arrive:** Since the elevator is currently at floor 1 and moving up, the logical next step for someone waiting to

Figure 11: An example of Claude Sonnet 4 and Gemini 2.5 Flash on Task 3/ Action Generation in question format: Image in question, text as option. The model failed to recognize that the true indicator of a called elevator is the "lit button", not the upward arrow on the elevator screen.

**Query:**  
 Following the content of the video input, which of the four provided options correctly represents the next action in the flow of "Turn on the AC".  
 A. Wait to feel the cool air.  
 B. Look down to find the switch.  
 C. Swipe the screen to lower down the temperature.  
 D. Press the power button.

**Input Video:**



**Ground Truth: B**

**gemin\_25\_flash**  
 ... look down to find the switch... while one might need to find a switch or a remote control, this is a preparatory step. The core action of turning the device on is pressing the button, not just looking for it.  
 \* \*\*C. Swipe the screen to lower down the temperature:\*\* This implies using a smartphone app. While possible, it's not the most universal method. More importantly, adjusting the temperature is typically done \*after\* the unit is powered on. The primary action is turning it on first.  
 \* \*\*D. Press the power button:\*\* This is the direct action required to turn on most electronic appliances, including air conditioners. Whether this button is on a remote control, a wall panel, or the unit itself, pressing it is the fundamental step to initiate power.  
 5. **Conclusion:** Out of the given choices, pressing the power button is the most direct and accurate description of the action required to turn the air conditioner on. The other options describe actions that happen either before (looking for the switch) or after (feeling the cool air, adjusting the temperature) the core action of powering the device on.  
 ### Answer  
 D

**qwen3\_8b**  
 ### Thinking Process  
 The video shows an air conditioner mounted on the wall, plugged in, and presumably ready to be turned on. Since the video only displays the AC unit itself without any human interaction or interface shown, the next logical step to "turn on the AC" would involve identifying how to activate it - either by pressing a button (if physical controls exist) or swiping a screen (if it's a smart model). However, since no screen is visible and no person is shown interacting with it, option D ("Press the power button") is the most direct and universally applicable action for turning on an AC unit. Option A ("Wait to feel the cool air") implies the AC is already running, which contradicts the goal of turning it on. Options B and C require additional context not provided in the video. Therefore, D is the best answer.  
 ### Answer  
 D


**qwen3\_vl**  
 ### Thinking Process  
 The video shows a stationary shot of an air conditioner mounted on a wall. Since the task is to "Turn on the AC," the logical next step after locating the AC (which is already shown) would be to activate it. Among the options, pressing the power button (D) is the most direct and necessary action to turn on the AC. The other options either assume the AC is already on (A), involve unnecessary steps (B), or relate to adjusting settings after turning it on (C). Therefore, pressing the power button is the correct next action.  
 ### Answer  
 D

**real\_gemin\_25\_pro**  
 ... look down to find the switch... while one might need to find a switch or a remote control, this is the most direct or common primary action. The primary action involves using the remote.  
 \* \*\*C. Swipe the screen to lower down the temperature:\*\* This action implies control via a smartphone or a smart home display. While possible, the most fundamental action is simply turning the power on first. Adjusting the temperature is a subsequent step after the unit is powered on.  
 \* \*\*D. Press the power button:\*\* This is the universal first step to turn on almost any electronic appliance. For an air conditioner, this button is typically found on its dedicated remote control. Pressing the power button sends the signal to the unit to start its operation. This is the most direct and necessary action to achieve the goal of "Turn on the AC."  
 5. **Conclusion:** Comparing the options, pressing the power button is the most logical and immediate next action to turn on the air conditioner. The other options are either results of the action (A) or secondary adjustments made after the unit is on (C), or a less common method of operation (B).  
 ### Answer  
 D

Figure 12: An example of Gemini 2.5 Flash, Gemini 2.5 Pro, Qwen3-VL-8B-Instruct and Qwen3-VL-235B-Instruct on Task 3/ Action Generation in question format: Video in question, text as option. The model failed to realize the "power button" is out of reach for this wall-mounted AC and should switch to "find a remote control".

**Query:**  
 Following the content of the video input, which of the four provided options correctly represents the next action in the flow of "Go upstairs".  
 A. Swipe the card.  
 B. Move away to see which elevator come.  
 C. Step into the elevators.  
 D. Wait for the elevator to arrive.

Input Video:



**Ground Truth: B**

**Model Outputs (All Wrong)**

**gemini\_25\_flash**

Thinking Process

- \* \*\*B. Move away to see which elevator come:\*\* While a person might physically step back to get a better view of both elevators, the fundamental action is "waiting". Moving is a secondary action performed "while" waiting. The core next step is the waiting itself.
- \* \*\*C. Step into the elevators:\*\* This is impossible. The video ends with the button being pressed. The elevator has not yet arrived at the floor, and the doors are still closed. You cannot step into an elevator that is not there.
- \* \*\*D. Wait for the elevator to arrive:\*\* This is the most logical and necessary next step. After calling the elevator, a period of waiting is required before it arrives and the doors open, allowing the person to proceed.

5. **Conclusion:** Based on the logical sequence of using an elevator, after pressing the call button, the immediate next step is to wait.

Therefore, option D is the correct answer.

### Answer  
D

**qwen3\_8b**

Thinking Process

The video shows a person approaching an elevator panel and pressing the button to go upstairs (the upward arrow button is pressed). After pressing the button, the next logical step in the flow of "Go upstairs" would be to wait for the elevator to arrive. The other options (swiping a card, stepping into the elevator immediately, or moving away) do not logically follow the action of pressing the button to go upstairs.

### Answer  
D

**qwen3\_vl**

Thinking Process

The video shows a person interacting with an elevator call panel. The person presses the "up" button on the touchless keypad, which is typically the first step in requesting an elevator to go upstairs. After pressing the button, the logical next step is to wait for the elevator to arrive, as the elevator needs time to reach the current floor. There is no indication of swiping a card or stepping into the elevator yet, and moving away to see which elevator comes is not a necessary or typical action after pressing the button.

### Answer  
D

Figure 13: An example of Gemini 2.5 Flash, Qwen3-VL-8B-Instruct and Qwen3-VL-235B-Instruct on Task 3/ Action Generation in question format: Video in question, text as option. The model fails to recognize that a procedural action is needed between pressing the button and waiting, to adjust the view for better visual information.

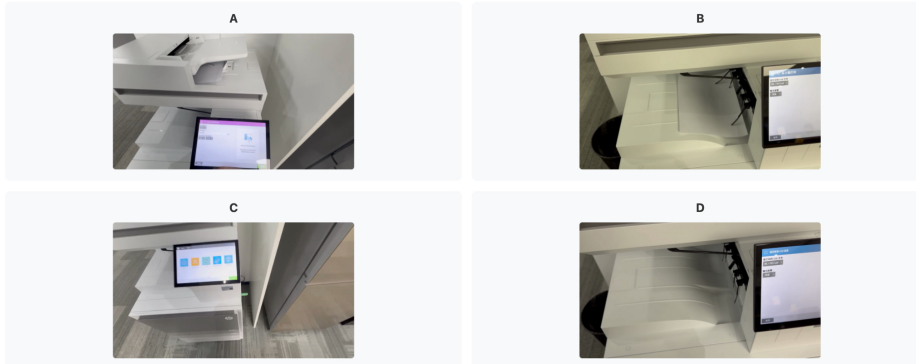
**Query:**

Answer the question based on the input image (current state) and the four image response options, marked as A/B/C/D. Following the content of the first image input, which of the four provided image options correctly represents the expected state upon outcome verification in the flow of "Print the file from the USB flash disk"?

**Input Image:**



**Option Images:**



**Ground Truth: B**

**claude\_sonnet\_4**

1. Insert USB and select file (current state)
2. Configure print settings
3. Execute print job
4. Outcome verification

After completing the print job and verifying the outcome, the printer would typically return to its main menu/home screen, ready for the next operation.

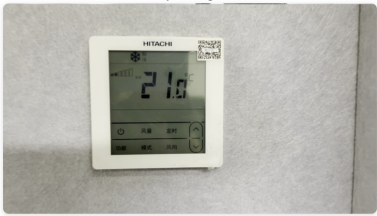
Looking at the options:

- Option A: Shows a scan interface - not related to print completion
- Option B: Shows the USB print selection screen (similar to starting state)
- Option C: Shows the main menu with multiple function icons - represents the printer returning to its default state after job completion
- Option D: Shows the same USB device warning screen as the input

Figure 14: An example of Claude Sonnet 4 on Task 4/ State Transition Prediction in question format: Image in question, images as option. The model failed to recognize that the completion of printing is indicated by paper coming out of the printer's output tray, which is exactly what option B depicts.

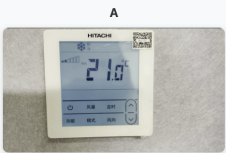
**Query:**  
 Answer the question based on the input image (current state) and the four image response options, marked as A/B/C/D. Following the content of the first image input and the action "Increase the fan speed", which one of the four provided image options will happen first?

**Input Image:**




**Option Images:**

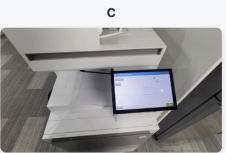
**A**



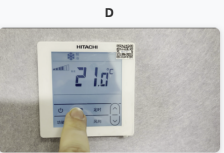
**B**



**C**



**D**



**Ground Truth: D**

✗ Model Outputs (All Wrong)

**claude\_sonnet\_4**

Let me analyze the four options:

- **Option A**: Shows the same thermostat display at 21.0°C in cooling mode, potentially with an updated fan speed indicator
- **Option B**: Shows a computer screen with a PDF printing interface - this appears unrelated to the thermostat
- **Option C**: Shows a printer from above - also unrelated to adjusting thermostat fan speed
- **Option D**: Shows a finger pressing a button on the thermostat (the action being performed)

The question asks what will happen **FIRST** after the action "Increase the fan speed."

Options B and C show printing-related screens that are completely unrelated to adjusting the air conditioner's fan speed.

Option D shows someone in the process of pressing the button (the action itself).

Option A shows the thermostat display after the adjustment, which would be the first observable result after increasing the fan speed - the display

Figure 15: An example of Claude Sonnet 4 on Task 4/ State Transition Prediction in question format: Image in question, images as option. The model fails to predict that the immediate outcome of increasing the fan speed is the change in the visible fan speed indicator displayed in the top-left part of the screen, thus chose the wrong option "A".