NovaPlan: Zero-Shot Long-Horizon Manipulation via Closed-Loop Video Language Planning

Jiahui Fu[1] **Junyu Nan**[1,2] **Lingfeng Sun**[1] **Hongyu Li**[1,3] **Jianing Qian**[4]
Yilun Du[5] **Jennifer L. Barry**[1] **Kris Kitani**[2] **George Konidaris**[3]
[1]Robotics and AI Institute [2]Carnegie Mellon University [3]Brown University
[4]University of Pennsylvania [5]Harvard University

Abstract: Solving complex long-horizon robotic tasks requires joint reasoning over abstract task structure and low-level physical interaction. While combining Vision-Language Models (VLMs) and video generation models offers a promising path for zero-shot planning, their individual tendencies to hallucinate physics or violate geometric consistency often compound over time, preventing reliable real-world execution. We introduce NovaPlan, a hierarchical framework that enables robust, zero-shot long-horizon manipulation by systematically proposing, verifying, and repairing visual plans. At the high level, a VLM planner decomposes tasks and filters out dynamically inconsistent futures by verifying multiple candidate video rollouts. To translate these imagined futures into reliable physical actions, NovaPlan utilizes a hybrid geometric representation that adaptively switches between object-centric flow and human hand flow. Finally, NovaPlan closes the loop by continuously monitoring execution to verify outcomes and synthesize local, non-prehensile corrective behaviors, such as fingertip poking, when failures occur. Across diverse multi-stage tasks, NovaPlan substantially outperforms prior zero-shot systems, achieving complex assembly and dexterous error recovery entirely without task-specific training or demonstrations. Please visit our project website for additional results: https://nova-plan.github.io/.

Keywords: Long-horizon planning, Zero-shot manipulation, Foundation models

1 Introduction

A central challenge in robotics is enabling agents to solve complex long-horizon tasks. Solving such tasks requires an agent to jointly reason about *abstract task structure*, deciding what sequence of actions should achieve a goal as well as *physical interaction*, deciding how objects, contacts, and constraints evolve under execution. These two forms of reasoning are tightly coupled: an action that is semantically correct may still make later interactions infeasible if it blocks access, causes misalignment, or leaves the scene in an unexpected state. Classical robotic systems typically separate these problems: symbolic planners reason over high-level task abstractions, while low-level controllers operate on geometric or dynamical models [1, 2, 3]. However, such modular pipelines struggle to scale to open-world environments due to extensive task- and object-specific modeling.

Recent foundation models suggest an alternative modular decomposition. Vision-language models (VLMs) have demonstrated strong semantic understanding and high-level reasoning capabilities [4, 5], while video generation models increasingly exhibit implicit knowledge of object dynamics, contact, and temporal evolution [6, 7, 8]. By jointly planning using both VLMs and video models through visual imagination [9, 10], we can combine both high- and low-level knowledge and construct a long-horizon robot planning system that can reason about both *abstract task structure* and *physical interaction*. In principle, this allows a robot to plan over imagined object interactions without task-specific demonstrations or manually specified dynamics.

10th Conference on Robot Learning (CoRL 2026), Austin, Texas, USA.

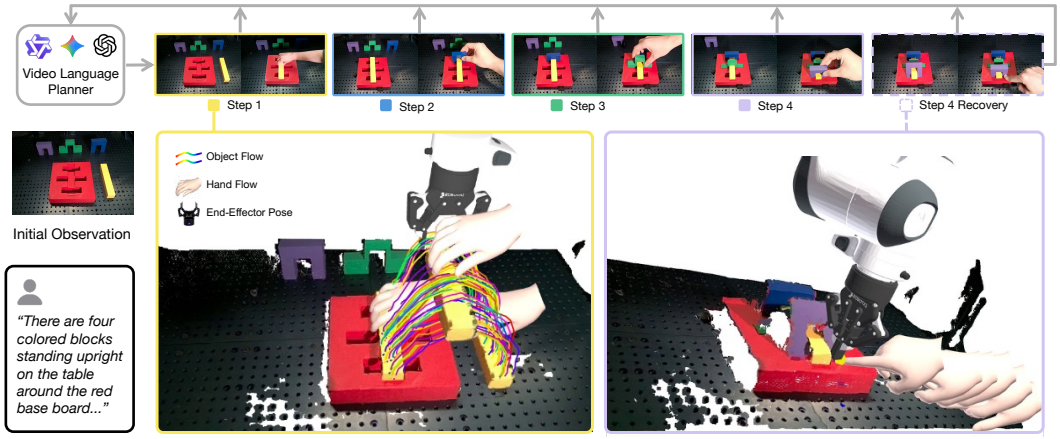


Figure 1: **Zero-shot long-horizon manipulation with NovaPlan.** Given a task command and an initial observation, a video language planner is used to select and generate the best next-step video. Executable robot actions are extracted from the video through object or hand flow tracking. The updated observation is sent back to the planner to enable closed-loop reasoning and recovery from failure states.

However, generated visual futures are not reliable transition models. VLMs may propose semantically plausible but physically infeasible actions, while video models may hallucinate objects and violate geometry. These errors are especially damaging in long-horizon manipulation, where deviations in imagined states can compound into full-task failure. Robust long-horizon operation therefore requires a closed-loop procedure that can effectively recover from these failures.

We present **NovaPlan**, a closed-loop framework for robustly solving zero-shot long-horizon manipulation tasks, where video planning is treated as a structured proposal mechanism within an online verification and repair system. At a high level, NovaPlan combines semantic task reasoning from VLMs with visual future prediction from video generation models, while augmenting these components with explicit verification, adaptive geometric grounding, and closed-loop recovery, enabling reliable execution of long-horizon tasks in a zero-shot manner.

NovaPlan robustly solves long-horizon tasks through three mechanisms. First, to construct *robust high-level planning*, the system generates multiple candidate video plans, verifies each of them using structured semantic and physical consistency checks, and filters out implausible futures from both VLM and video models. Second, to enable *geometrically grounded execution* from plans, NovaPlan converts generated videos into robot trajectories through a hybrid representation that adaptively switches between object-centric flow and hand-centric flow, enabling us to robustly infer executable actions across videos. Third, to enable *closed-loop execution and recovery from inaccurate plans*, NovaPlan continuously re-grounds and regenerates plans from current observations, verifies successful execution, and synthesizes local corrective behaviors, including non-prehensile recovery actions, when failures occur.

We evaluate NovaPlan on diverse long-horizon manipulation tasks and the challenging Functional Manipulation Benchmark (FMB). Across multi-step tasks including block stacking, object sorting, and exploratory hidden-object search, NovaPlan shows substantially improved robustness over various prior zero-shot manipulation approaches [10, 11, 12, 13, 14]. NovaPlan also demonstrates zero-shot local recovery behaviors, such as fingertip poking for object realignment.

Overall, our contributions are as follows: **(1)** We present NovaPlan, a closed-loop video language planning framework that enables robust long-horizon task execution by systematically proposing, verifying, and re-grounding visual plans. **(2)** We develop a hybrid visual grounding module that converts accepted visual transitions into executable robot trajectories by adaptively switching between object-centric flow and hand-centric flow. **(3)** We validate NovaPlan on diverse multi-step manipulation tasks, where it outperforms various leading zero-shot baselines.

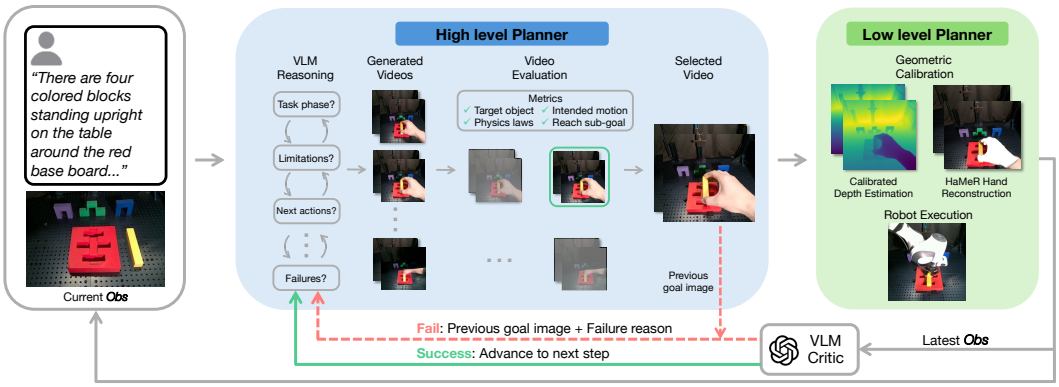


Figure 2: **NovaPlan system overview.** A high-level planner takes in the task instruction and current observation, and proposes multiple task-solving videos after VLM reasoning about the task progress. Videos are selected based on flow and semantic consistency. The low-level planner calculates the robot action from the extracted hand or object flow. Updated observations are sent to a VLM critic to determine whether the robot should proceed to the next step or recover from failure.

2 Related Work

Zero-shot manipulation. Inspired by recent advances in Large Language Models (LLMs) and Vision-Language Models (VLMs) [15, 16, 17], there is growing interest in solving robotic manipulation tasks through zero-shot generalization. One line of work trains Vision-Language-Action (VLA) models [18, 19, 20] on large-scale real-world robot trajectories to learn generalizable control policies. Several recent VLAs further inject 3D structure into the policy backbone to improve geometric precision [21, 22, 23, 24], although that structure is learned inside a policy trained on robot demonstrations. VLA scalability also remains limited by the cost of collecting physical interaction data relative to web-scale text and images. An alternative line of work uses VLMs to generate symbolic or spatial plans rather than direct motor commands [12, 25, 10, 26, 27], including agents that emit robot code against a designer-provided API [28]. While effective at capturing task structure, these methods often lack the fine-grained 3D spatial awareness required for robust manipulation.

Video generation for robotics. Video generation models provide a promising representation for dynamic reasoning and future-state prediction, which are leveraged by robotic applications in several distinct ways. One family of work predicts future frames from an initial observation and a language instruction (image-to-video, I2V), which requires no action labels at inference time. Other work instead conditions video generation on robot actions, yielding controllable world models [29], simulators for policy evaluation [30], and world action models that act as zero-shot policies [13, 31]. While this enables tighter controllability, it requires action-annotated data and ties the model to its training embodiments. Prior work converts generated videos into robot actions using inverse dynamics models [32, 9, 33] or imitation policies [34, 35, 36], both of which require task-specific demonstrations. Other work extracts control signals directly using 6D pose estimation [37, 38] or flow tracking [11, 34]. However, these approaches remain sensitive to perception error, tracking drift, and open-loop execution, which limits robustness over long horizons. Closest to our setting, video language planning [10] and compositional hierarchical planning [9] search jointly over language actions and imagined video futures, and our beam search follows their search method. VLP prompts a VLM to measure progress to final goals, and notes in its limitations that this often leads to inconsistent video plans. NovaPlan instead applies target, physical, motion, and result checks, converts videos into metric robot motion through object- or hand-flows with no robot training, and regrounds each plan on the latest visual observation with local recovery rather than receding-horizon replanning, which can leave later actions conditioned on stale predicted states (Figure 5).

3 Method

Our goal is to develop a closed-loop hierarchical video language planner for zero-shot long-horizon manipulation. As shown in Figure 2, NovaPlan treats generated videos as visual transition proposals

inside an online planning, verification, and repair loop. The planner proposes candidate macro-actions, samples video rollouts, verifies their semantic and physical consistency, and re-grounds selected transitions on the latest real observation before execution (Section 3.1). Accepted visual transitions are then converted into executable robot trajectories through a hybrid object-flow and hand-flow grounding module (Section 3.2).

3.1 Closed-Loop Video Language Planning

NovaPlan's high-level planner combines VLM-based text action generation with video-based physical verification to reason about possible future actions. The planner uses a proposal–rollout–verification procedure: the VLM proposes one or more candidate text actions from a source image and task context, the video model instantiates each action as multiple visual rollouts, and a VLM critic filters physically or semantically inconsistent futures to choose the most plausible rollout for execution. This procedure is used inside strategic beam search to build full text plans for coupled tasks, reactively to select the next action for exploratory or non-coupled tasks, and locally to produce recovery actions after execution failures. As summarized in Algorithm 1, NovaPlan closes the loop by verifying real execution outcomes and invoking local recovery when the observed transition deviates from the intended visual target. Please see Appendix C for full design details.

High-level action proposal. Given a source image I_{src}, task goal g, action history $\mathcal{H}$, and constraint context ξ, the VLM proposes candidate macro-actions $\Pi = \{(a^i, o^i, \xi^i)\}_{i=1}^{N_c}$, where a^i is a language-level manipulation action, o^i is the object or region to track, and ξ^i updates reasoning over task progress. I_{src} is the last frame of the previously selected video during strategic planning and the latest real observation during reactive planning, execution, and recovery. NovaPlan proposes multiple actions each step in strategic planning to reason over future action orders, and only the next action in reactive planning and recovery for efficiency.

Algorithm 1 Closed-loop video language planning

Input: goal g, initial RGB I_0^{real}, horizon cap N, video model G

$I \leftarrow I_0^{\text{real}}$, $\mathcal{H} \leftarrow \emptyset$, $\xi \leftarrow \emptyset$, $\mathcal{Q} \leftarrow \emptyset$
$m \leftarrow \textsc{AssessTaskStructure}(I, g)$ ▷ strategic or reactive
if $m = \text{strategic}$ **then**
  $\mathcal{Q}, \xi \leftarrow \textsc{StrategicBeamSearch}(G, I, g, \mathcal{H}, \xi, N)$
end if
while task not complete **do**
  if $m = \text{strategic}$ **then**
    if $\mathcal{Q} = \emptyset$ **then**
      break ▷ planned sequence completed
    end if
    $(a, o) \leftarrow \textsc{Pop}(\mathcal{Q})$
  else
    $(a, o, \xi) \leftarrow \textsc{ProposeAction}(I, g, \mathcal{H}, \xi)$
  end if
  $V \leftarrow \textsc{GenerateAndSelectRollout}(G, I, a, o)$
  $I^{\text{target}} \leftarrow \text{LastFrame}(V)$
  $\tau \leftarrow \textsc{GroundVideo}(V, o)$; execute τ and observe I^+
  if $\textsc{VerifyTransition}(I, I^+, I^{\text{target}}, a)$ **then**
    $\mathcal{H} \leftarrow \mathcal{H} \cup \{a\}$; $I \leftarrow I^+$
  else
    $r, I^+ \leftarrow \textsc{AttemptRecovery}(G, I^+, I^{\text{target}}, a, g)$
    if $r = \text{success}$ **then**
      $\mathcal{H} \leftarrow \mathcal{H} \cup \{a\}$; $I \leftarrow I^+$
    else
      break ▷ recovery failed
    end if
  end if
end while

Video rollouts as physical tests. The proposed actions are next converted into imagined visual futures. For each action a^i, NovaPlan samples L rollouts from the video model, $\{V^{i,k}\}_{k=1}^{L} = G(I_{\text{src}}, a^i)$. Each rollout is treated as a hypothesis about what the action would physically do from the given source image. This same rollout-generation step is used in three places: in strategic beam search to imagine future states, in execution to instantiate the selected text command from the latest real scene, and in recovery to generate candidate corrective motions toward the previous target state.

Hierarchical rollout verification. As generated videos can look plausible while being physically wrong, NovaPlan verifies each rollout before using it for planning, execution, or recovery. As described in Appendix C.2, a VLM critic applies four ordered checks: *target consistency*, whether the intended object moves; *physical consistency*, whether the video avoids artifacts such as deformation or hallucinated objects; *motion consistency*, whether the tracked flow matches the action command; and *result consistency*, whether the final frame achieves the subgoal. This hierarchy reduces hallu-

cination in the VLM critic and filters out goal-like final frames reached via impossible or irrelevant motions. NovaPlan batch-scores rollouts by arranging each candidate's final frame and its flow image (object flow across the rollout) into a grid, enabling efficient comparison of multiple futures.

Verified lookahead. The planning horizon is selected based on the structure of the task. For exploratory tasks, such as opening drawers to search for a hidden object, the next useful action depends on information revealed by execution. In this case, NovaPlan uses a greedy horizon $h = 1$, executes one verified action, and observes again. For coupled tasks, such as assembly, the order of actions can determine whether later insertions remain feasible. In this case, NovaPlan uses a strategic horizon $h = N$ and performs beam search [10] in advance to generate the full task plan to be followed during execution, as detailed in Appendix C.1.2. Each beam stores a starting image, an action history, a constraint context, and a rollout score. The final frame of the last selected rollout becomes the starting image for the next search depth, allowing video generation to serve as a coarse object-dynamics model to reason over action orders for full-step planning before execution.

Execution-time re-grounding. During execution, NovaPlan follows the text action command proposed by the planner: for strategic tasks, the command is popped from the beam-search plan, while for reactive tasks it is proposed from the current observation. Before executing the selected action a_t, NovaPlan generates candidate videos from the latest real observation, $\{V_t^k\} = G(I_t^{\text{real}}, a_t)$. The same rollout verifier selects the execution video V_t^{exec}, and its last frame defines the local target state, $I_t^{\text{target}} = \text{LastFrame}(V_t^{\text{exec}})$. This re-grounding step ensures that the low-level planner always receives a visual transition aligned with the current scene rather than relying on an imagined starting condition produced earlier during planning.

State verification after execution. After the grounded trajectory is executed, NovaPlan checks whether the real scene undergoes the intended transition. As detailed in Appendix C.3.1, the VLM critic compares three images: the pre-execution observation I_t^{real}, the post-execution observation I_{t+1}^{real}, and the generated target frame I_t^{target}, and determines if $I_t^{\text{real}} \rightarrow I_{t+1}^{\text{real}}$ matches the planned visual transition $I_t^{\text{real}} \rightarrow I_t^{\text{target}}$. If verification succeeds, the planner advances to the next action. If verification fails, the failure explanation is passed to the recovery stage.

Local recovery. When execution fails, NovaPlan attempts a local repair operation. The previous target frame I_t^{target} remains the intended local goal, while the current observation I_{t+1}^{real} becomes the new start state. As described in Appendix C.3.2, the VLM compares these two images given the intended action, selects a recovery mode $S_t \in \{\text{grasp}, \text{non-prehensile}\}$, and proposes a recovery action following the selected mode. The same proposal–rollout–verification pipeline is then used to generate and select a recovery video for local repair. If the recovered state reaches the previous target, NovaPlan resumes the high-level plan; otherwise, the recovery attempt is marked as failed.

3.2 Geometric Grounding of Verified Video Plans

During execution, after the high-level planner selects a verified video transition, NovaPlan grounds the visual motion into an executable end-effector trajectory. As the generated videos may contain geometric artifacts and are not metric robot trajectories, NovaPlan extracts motion by leveraging two complementary 3D cues: object flow and hand flow. Object flow recovers target object motion when the object remains visible and rigidly trackable, while hand flow uses the generated human hand as a kinematic prior when object tracking fails due to occlusion, large rotations, or unstable depth. This hybrid design allows NovaPlan to robustly convert imperfect generated videos into useful robot motion. Please see Appendix D for full design details.

Object-centric grounding. For object-centric grounding, we recover a metric 3D trajectory of the target object from the generated video, following the general actionable flow pipeline of NovaFlow [11]. As described in Appendix D.1, given a generated video and the target object specified by the planner, we first estimate per-frame geometry using a monocular geometry model [39], optionally refine the depth sequence for temporal consistency using CVD [40], and align the initial generated depth to the real sensor depth through affine calibration [37, 11]. We then segment the target object, sample object points, and track them in 3D using a dense point tracker [41]. The

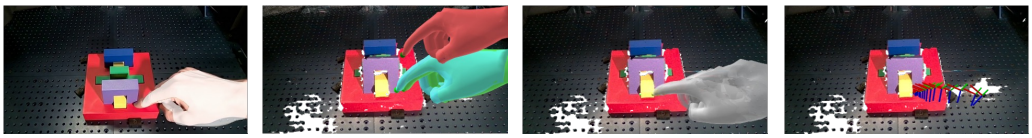


Figure 3: **Geometric grounding for non-prehensile recovery.** (a) HaMeR [46] predicts the MANO hand mesh on the generated recovery video. (b) Red: mesh from raw prediction. Green: mesh after scale calibration. Cyan: mesh after additional translation offset enforcing fingertip and object surface contact. (c) The hand trajectory after calibration. (d) Calibrated trajectory converted into per-frame SE(3) transforms.

resulting object flow is denoted as

$$\mathcal{F}_{\text{obj}} = \{\mathbf{f}_i^t \in \mathbb{R}^3 \mid i = 1, \ldots, K,\ t = 1, \ldots, T\},$$

where $\mathbf{f}_i^t$ is the 3D position of object point i at frame t. We fit a rigid transform between consecutive frames from the tracked points, yielding an object motion sequence $\{T_{\text{obj}}^t\}_{t=1}^T$.

To execute the object motion, we generate a grasp on the initial object geometry using a grasp proposal model [42]. The selected grasp defines a fixed object-to-end-effector transform $T_{\text{obj}\to\text{ee}}$. Assuming the object remains fixed relative to the gripper after grasping [37, 43, 38, 44], we convert the recovered object trajectory into an end-effector trajectory $T_{\text{ee}}^t = T_{\text{obj}}^t T_{\text{obj}\to\text{ee}}$. This representation is accurate when the target object is visible, rigid, and consistently tracked.

Hand-centric grounding. Object flow becomes unreliable when the target object is heavily occluded by the generated hand, when depth estimates flicker around contact, or when the fitted object motion contains implausibly large rotations. In such cases, NovaPlan uses the generated human hand as a kinematic prior. We reconstruct a MANO hand mesh [45] for each video frame using HaMeR [46], obtaining a hand trajectory $\mathcal{H} = \{\mathbf{H}^t \mid t = 1, \ldots, T\}$.

Since generated hands often exhibit scale errors and projective drift, the raw hand reconstruction is not directly metric. We thus apply a lightweight contact-based calibration (detailed in Appendix D.2) to the video hand motion: we detect the interaction interval from object-mask motion, align the hand to the object at contact onset, and compensate for residual drift near release. The calibrated hand motion is converted into a rigid $SE(3)$ trajectory by using the calibrated contact fingertip as the translation reference and a palm frame estimated from hand landmarks as the orientation reference. The resulting geometric effect is illustrated in Figure 3, where the raw hand reconstruction is brought into metric contact with the object and converted into per-frame $SE(3)$ transforms.

Flow switching and execution. NovaPlan first attempts object-centric grounding and switches to hand-centric grounding when the recovered object motion is unstable. We reject object flow if tracking fails or if the fitted inter-frame rotation exceeds a predefined threshold, i.e., $\exists t,\ \theta_t > \theta_{\max}$, where θ_t denotes the axis-angle rotation magnitude of the per-frame fitted object transform and $\theta_{\max}$ is the rotation threshold for triggering the flow switch. Appendix B presents relevant analyses on the design choice of the hand flow and the effect of varying $\theta_{\max}$ on execution performance. When object flow is rejected, the system switches to hand-centric grounding. For hand flow, the initial hand–object contact pose guides grasp selection (see Appendix D.3), and the robot follows the calibrated hand trajectory after grasping.

For non-prehensile recovery videos, the contact finger is specified by the recovery prompt rather than inferred from fingertip-mask overlap. We therefore constrain hand-flow calibration to the prompted contact finger, enforcing metric contact between the fingertip and the selected object surface. This contact-constrained grounding is illustrated in Figure 3 and described in detail in Appendix D.4.

4 Experiments

NovaPlan is designed to solve diverse long-horizon manipulation problems in a zero-shot manner. To evaluate its capabilities, we design a set of long-horizon tasks that assess reasoning, task decomposition, sub-task execution, error diagnosis, and recovery. We also include tasks from [11] to evaluate improvements in short-horizon zero-shot settings. Finally, we use an assembly task from

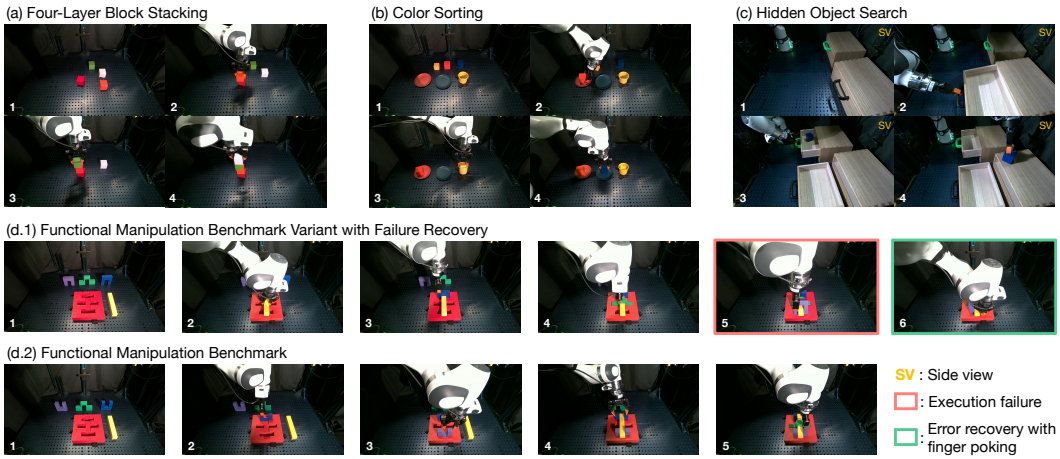


Figure 4: **Robot Experiments.** (a-c) The long-horizon tasks in Section 4.1, each with three steps. (d) Functional Manipulation Benchmark Multi-Object Multi-Stage Assembly 1 task and its variant in Section 4.3, including one recovery step.

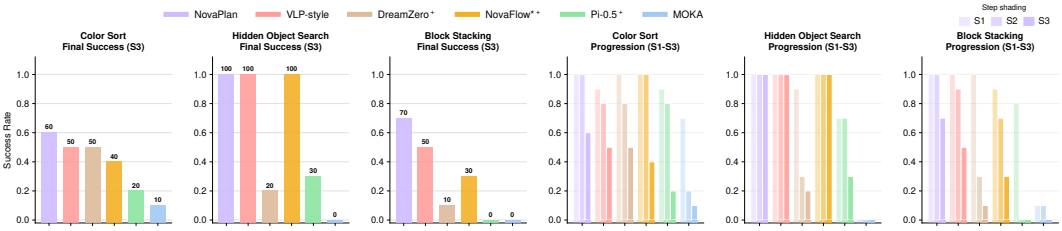


Figure 5: **Long-horizon task performance.** (Left) Final success rate of different methods on three long-horizon tasks. (Right) Success rate of each method after every step of the three long-horizon tasks, where shading denotes the step index. Methods marked with $^+$ (NovaFlow, DreamZero, and $\pi_{0.5}$) receive oracle high-level plans; MOKA and VLP-style generate their own plans.

the Functional Manipulation Benchmark (FMB) [47] to discuss the boundaries of NovaPlan under the same zero-shot setup. Appendix A presents implementation details and runtime analysis, where each macro-action computation takes roughly 40 s end to end.

For baselines, we compare NovaPlan against five representative zero-shot methods from distinct architectural paradigms: NovaFlow [11] (video-based planning), VLP-style planning [10] (video-language planning), DreamZero [13] (world action model), $\pi_{0.5}$ [14] (VLA model), and MOKA [12] (VLM-based planning). The VLP-style baseline keeps our video generation and training-free hybrid-flow execution and replaces only the planner, so that the comparison reflects the planning loop rather than choices of backbone models. NovaFlow, DreamZero, and $\pi_{0.5}$ receive oracle subtask decompositions, isolating grounding and execution for these baselines; MOKA and VLP-style use their own task planners (Appendix E).

4.1 Long-Horizon Task Solving

To test NovaPlan's long-horizon robustness, we evaluate the system's performance on three tasks:

- **Four-layer Block Stacking:** Stack four colored 2-inch blocks into a vertical tower, with each step placing a block onto the increasingly taller stack. This task evaluates spatial reasoning and extracted action accuracy under increasing placement height.
- **Color Sorting:** Sort three blocks into color-matched containers. Beyond semantic reasoning, the yellow block requires precise vertical alignment due to its tight fit to avoid jamming.
- **Hidden Object Search:** Find a target object hidden in one of the two closed drawers and place it on the other drawer's top surface. This exploratory task evaluates conditional planning, requiring a two- or three-step plan depending on which drawer contains the object.

We conduct ten trials for each task with varying initial setups. For each step in the long-horizon tasks, we mark the trial as a failure only when the method *cannot recover* from an error. For *Color*

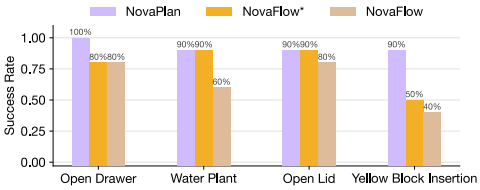


Figure 6: **Single-step task performance.** Comparison with NovaFlow* (implemented) and NovaFlow (from [11]).

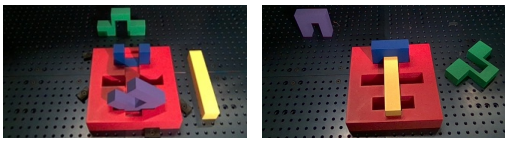


Figure 7: **Failure cases on FMB task** due to (a) video model generating physically implausible motions or (b) the need for object reorientation.

Sorting, since sub-tasks are order-independent, we use the step number to represent the number of correctly sorted objects after full execution. We report both overall long-horizon task success rates and task progression based on single-step success rates in Figure 5. For more results on the generated plans and experiment rollout visualizations, please see Appendices F and G.

In *Block Stacking*, performance decreases as the stack grows taller because later placements require more precise spatial reasoning and control: $\pi_{0.5}$ stacks up to two layers but fails beyond that, while NovaFlow reaches the third block with 70% success but drops to 30% on the fourth due to object-flow instability. By switching to hand flow when object flow becomes unreliable, NovaPlan improves execution stability and succeeds in seven out of ten trials. In *Color Sorting*, the low-tolerance "yellow block and cup" step causes all methods to drop in performance; for NovaPlan, the remaining failures are mainly caused by inaccurate pose extraction from depth estimation errors, which could be mitigated by stronger depth models or additional camera views. In *Hidden Object Search*, success depends on both drawer opening and object retrieval: $\pi_{0.5}$ performs better on drawer opening than retrieval, likely reflecting the task distribution in the DROID dataset [48], while MOKA fails to find a valid grasp for the horizontal drawer handle. Both NovaPlan and NovaFlow succeed in all trials, as their sub-skill execution remains reliable across the task sequence. VLP-style planning matches NovaPlan on reactive search, but degrades on the coupled tasks, as a mismatch between predicted and real states changes later placement geometry, and incurs 2–3 times as many re-planning calls due to mismatch-caused execution failures. DreamZero degrades sharply after the first action. Overall, these results show that NovaPlan combines stable short-horizon execution with robust switching and recovery, enabling stronger long-horizon manipulation performance.

4.2 Short-Horizon Task Solving

To evaluate the efficacy of our hybrid flow mechanism, we compare NovaPlan with NovaFlow on its original single-step task suite: *Open Drawer*, *Water Plant*, *Open Lid*, and *Yellow Block Insertion*, using ten trials per task under the same setup. We exclude *Hang Mug*, which requires a goal image for plan video generation, and *Cup on Saucer*, whose pick-and-place behavior is well covered by our long-horizon tasks. This comparison tests whether hand flow can mitigate self-occlusion failures in object-centric tracking. As shown in Figure 6, NovaPlan achieves higher success rates across all tasks, indicating that using the hand as a reference improves stability when object flow is noisy.

4.3 Zero-Shot FMB Solving

We use FMB as a case study to probe the capability boundary of NovaPlan. FMB requires the system to compose individual manipulation skills into functionally meaningful long-horizon behaviors, making it a challenging stress test for zero-shot manipulation. The task demands long-horizon planning, millimeter-level assembly precision, and irregular object shapes that are often unseen by pre-trained foundation models, with failure modes that cannot be resolved by simple re-grasping.

Since here *none* of the chosen VLA/VLM baselines can complete even a single step, we focus on the qualitative success and failure mode analysis of NovaPlan. We evaluate on standard Initial Layout 3 of the FMB Multi-Object Multi-Stage Task Assembly 1, shown in Figure 4(d.2), where all parts lie flat on the table and require no reorientation. We also test a variant with the U-shaped parts placed upside down, shown in Figure 4(d.1), to evaluate execution on a different assembly target.

We identify two main bottlenecks in this setting. First, video generation becomes unreliable for irregular assembly objects. Only Veo3.1 can generate physically feasible videos for all FMB stages,

Table 1: **(a) Component ablations and (b) failure attribution** (%; 10 trials per task). In (a), each row swaps one module, all else fixed; ‡ scored by long-horizon task success, others by module-level correctness. In (b), * covers video generation, object/hand segmentation, and recovery reasoning.

(a) Controlled component ablations

Component	Variant	Ours	Ablated
Strategic beam search	Reactive planning	**85.0**	15.0
Execution re-grounding‡	Reuse planned rollout	**76.7**	43.3
Rollout verification	Without flow image	**98.8**	42.4
Hybrid grounding‡	Object flow only	**76.7**	46.7
Hand calibration	Raw HaMeR	**93.3**	13.3
State verification	Without target frame	**98.8**	92.9
Local recovery‡	Terminate on failure	**76.7**	50.0

(b) Failure attribution by module

Module	Failure rate
3D object tracking	20.0
Depth estimation	12.9
Hand calibration	6.7
Grasp selection and execution	5.9
Rollout selection	1.2
State verification	1.2
Video generation, segmentation, recovery*	0.0

and even then with a much lower success rate, especially during failure recovery where contact dynamics of the human finger are more complex. In this case, NovaPlan keeps regenerating plans until a feasible recovery video is obtained. Second, insertion recovery remains difficult because both object displacement and hand motion are small after partial failure. As shown in Figure 7, this causes object keypoint flow and hand pose estimates to be dominated by noise, making both object-flow and grasp-mode hand-flow execution unreliable. In these cases, NovaPlan can still ground execution on fingertip motion during non-prehensile pushing, providing a cleaner signal that enables recovery.

4.4 Component Ablations and Failure Attribution

Table 1(a) presents the module ablation results, where we replace one module in the pipeline at a time for a variant and observe two patterns. First, the pre-execution modules each prevent a potential failure. Scoring rollouts from frames alone collapses valid selection to 42.4%, since a goal-like final frame can be reached by motion that never occurs. Calibration and planning are also meaningful: raw HaMeR output is not metric, and an ordering error cannot be undone once a step blocks access. Second, the execution-time mechanisms are not incremental refinements. Removing execution re-grounding, local recovery, or hand flow each costs less than half of the long-horizon success rate. The switch to hand flow takes up only a modest fraction of steps, but those are the steps that are hard and would otherwise fail a trial (Appendix B). Table 1(b) then shows the breakdown of module failure rates. The reasoning stages are most reliable, and infeasible rollouts are filtered before the robot moves, so NovaPlan's bottleneck lies in metric geometry rather than in the planner. Please refer to Appendices H.1 and H.2 for full details.

5 Limitations and Conclusion

Limitations. NovaPlan relies on foundation and perception models for vision-language reasoning, video generation, depth estimation, keypoint tracking, and hand pose estimation, and its performance is bounded by their capabilities, with metric geometry the weakest part (Table 1(b)). (1) *Scope*: We evaluate tabletop scenes of rigid object motion under a single fixed camera, and do not address clutter, deformables, or transparent objects. (2) *Evaluation*: Ten trials per task suffice for real-robot comparison, but not for claims of broad object or scene diversity, and while our prompts are robust to paraphrase, they degrade once the reasoning rubric is removed (Appendix H.3). (3) *Recovery and safety:* Recovery only chooses between grasping and non-prehensile repair, and video models still fail to generate corrections for many failure modes, leaving reorientation and Figure 7(a) out of reach. Since trajectories come from generated videos, a verified rollout can still be physically wrong, so we give no formal safety guarantee and would not deploy NovaPlan unsupervised. Yet, given its principled interface structure, NovaPlan will naturally improve with advances in foundation models and geometric reasoning.

Conclusion. We introduce NovaPlan, a zero-shot, long-horizon manipulation framework coupling closed-loop video-language planning with geometrically grounded execution. By extracting object keypoints and hand poses from generated videos, NovaPlan connects high-level task reasoning with low-level control. Experiments across diverse long-horizon tasks and the FMB boundary case study show decent manipulation and error recovery without task-specific demonstrations or training.

Acknowledgments

The authors sincerely thank Benjamin Yang for their assistance with the evaluation experiments. This work was fully supported by the Robotics and AI Institute.

References

[1] S. Cambon, R. Alami, and F. Gravot. A hybrid approach to intricate motion, manipulation and task planning. *The International Journal of Robotics Research*, 28(1):104–126, 2009.

[2] J. Wolfe, B. Marthi, and S. Russell. Combined task and motion planning for mobile manipulation. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 20, pages 254–257, 2010.

[3] L. P. Kaelbling and T. Lozano-Pérez. Hierarchical task and motion planning in the now. In *2011 IEEE International Conference on Robotics and Automation*, pages 1470–1477. IEEE, 2011.

[4] A. Singh, A. Fry, A. Perelman, A. Tart, A. Ganesh, A. El-Kishky, A. McLaughlin, A. Low, A. Ostrow, A. Ananthram, A. Nathan, A. Luo, A. Helyar, A. Madry, A. Efremov, A. Spyra, A. Baker-Whitcomb, A. Beutel, A. Karpenko, A. Makelov, A. Neitz, A. Wei, A. Barr, A. Kirchmeyer, A. Ivanov, A. Christakis, A. Gillespie, A. Tam, A. Bennett, A. Wan, A. Huang, A. M. Sandjideh, A. Yang, A. Kumar, A. Saraiva, A. Vallone, A. Gheorghe, A. G. Garcia, A. Braunstein, A. Liu, A. Schmidt, A. Mereskin, A. Mishchenko, A. Applebaum, A. Rogerson, A. Rajan, A. Wei, A. Kotha, A. Srivastava, A. Agrawal, A. Vijayvergiya, A. Tyra, A. Nair, A. Nayak, B. Eggers, B. Ji, B. Hoover, B. Chen, B. Chen, B. Barak, B. Minaiev, B. Hao, B. Baker, B. Lightcap, B. McKinzie, B. Wang, B. Quinn, B. Fioca, B. Hsu, B. Yang, B. Yu, B. Zhang, B. Brenner, C. R. Zetino, C. Raymond, C. Lugaresi, C. Paz, C. Hudson, C. Whitney, C. Li, C. Chen, C. Cole, C. Voss, C. Ding, C. Shen, C. Huang, C. Colby, C. Hallacy, C. Koch, C. Lu, C. Kaplan, C. Kim, C. Minott-Henriques, C. Frey, C. Yu, C. Czarnecki, C. Reid, C. Wei, C. Decareaux, C. Scheau, C. Zhang, C. Forbes, D. Tang, D. Goldberg, D. Roberts, D. Palmie, D. Kappler, D. Levine, D. Wright, D. Leo, D. Lin, D. Robinson, D. Grabb, D. Chen, D. Lim, D. Salama, D. Bhattacharjee, D. Tsipras, D. Li, D. Yu, D. Strouse, D. Williams, D. Hunn, E. Bayes, E. Arbus, E. Akyurek, E. Y. Le, E. Widmann, E. Yani, E. Proehl, E. Sert, E. Cheung, E. Schwartz, E. Han, E. Jiang, E. Mitchell, E. Sigler, E. Wallace, E. Ritter, E. Kavanaugh, E. Mays, E. Nikishin, F. Li, F. P. Such, F. de Avila Belbute Peres, F. Raso, F. Bekerman, F. Tsimpourlas, F. Chantzis, F. Song, F. Zhang, G. Raila, G. McGrath, G. Briggs, G. Yang, G. Parascandolo, G. Chabot, G. Kim, G. Zhao, G. Valiant, G. Leclerc, H. Salman, H. Wang, H. Sheng, H. Jiang, H. Wang, H. Jin, H. Sikchi, H. Schmidt, H. Aspegren, H. Chen, H. Qiu, H. Lightman, I. Covert, I. Kivlichan, I. Silber, I. Sohl, I. Hammoud, I. Clavera, I. Lan, I. Akkaya, I. Kostrikov, I. Kofman, I. Etinger, I. Singal, J. Hehir, J. Huh, J. Pan, J. Wilczynski, J. Pachocki, J. Lee, J. Quinn, J. Kiros, J. Kalra, J. Samaroo, J. Wang, J. Wolfe, J. Chen, J. Wang, J. Harb, J. Han, J. Wang, J. Zhao, J. Chen, J. Yang, J. Tworek, J. Chand, J. Landon, J. Liang, J. Lin, J. Liu, J. Wang, J. Tang, J. Yin, J. Jang, J. Morris, J. Flynn, J. Ferstad, J. Heidecke, J. Fishbein, J. Hallman, J. Grant, J. Chien, J. Gordon, J. Park, J. Liss, J. Kraaijeveld, J. Guay, J. Mo, J. Lawson, J. McGrath, J. Vendrow, J. Jiao, J. Lee, J. Steele, J. Wang, J. Mao, K. Chen, K. Hayashi, K. Xiao, K. Salahi, K. Wu, K. Sekhri, K. Sharma, K. Singhal, K. Li, K. Nguyen, K. Gu-Lemberg, K. King, K. Liu, K. Stone, K. Yu, K. Ying, K. Georgiev, K. Lim, K. Tirumala, K. Miller, L. Ahmad, L. Lv, L. Clare, L. Fauconnet, L. Itow, L. Yang, L. Romaniuk, L. Anise, L. Byron, L. Pathak, L. Maksin, L. Lo, L. Ho, L. Jing, L. Wu, L. Xiong, L. Mamitsuka, L. Yang, L. McCallum, L. Held, L. Bourgeois, L. Engstrom, L. Kuhn, L. Feuvrier, L. Zhang, L. Switzer, L. Kondraciuk, L. Kaiser, M. Joglekar, M. Singh, M. Shah, M. Stratta, M. Williams, M. Chen, M. Sun, M. Cayton, M. Li, M. Zhang, M. Aljubeh, M. Nichols, M. Haines, M. Schwarzer, M. Gupta, M. Shah, M. Y. Guan, M. Huang, M. Dong, M. Wang, M. Glaese, M. Carroll, M. Lampe, M. Malek, M. Sharman, M. Zhang, M. Wang, M. Pokrass, M. Florian, M. Pavlov, M. Wang, M. Chen, M. Wang, M. Feng, M. Bavarian, M. Lin, M. Abdool, M. Rohaninejad,

N. Soto, N. Staudacher, N. LaFontaine, N. Marwell, N. Liu, N. Preston, N. Turley, N. Ansman, N. Blades, N. Pancha, N. Mikhaylin, N. Felix, N. Handa, N. Rai, N. Keskar, N. Brown, O. Nachum, O. Boiko, O. Murk, O. Watkins, O. Gleeson, P. Mishkin, P. Lesiewicz, P. Baltescu, P. Belov, P. Zhokhov, P. Pronin, P. Guo, P. Thacker, Q. Liu, Q. Yuan, Q. Liu, R. Dias, R. Puckett, R. Arora, R. T. Mullapudi, R. Gaon, R. Miyara, R. Song, R. Aggarwal, R. Marsan, R. Yemiru, R. Xiong, R. Kshirsagar, R. Nuttall, R. Tsiupa, R. Eldan, R. Wang, R. James, R. Ziv, R. Shu, R. Nigmatullin, S. Jain, S. Talaie, S. Altman, S. Arnesen, S. Toizer, S. Toyer, S. Miserendino, S. Agarwal, S. Yoo, S. Heon, S. Ethersmith, S. Grove, S. Taylor, S. Bubeck, S. Banesiu, S. Amdo, S. Zhao, S. Wu, S. Santurkar, S. Zhao, S. R. Chaudhuri, S. Krishnaswamy, Shuaiqi, Xia, S. Cheng, S. Anadkat, S. P. Fishman, S. Tobin, S. Fu, S. Jain, S. Mei, S. Egoian, S. Kim, S. Golden, S. Mah, S. Lin, S. Imm, S. Sharpe, S. Yadlowsky, S. Choudhry, S. Eum, S. Sanjeev, T. Khan, T. Stramer, T. Wang, T. Xin, T. Gogineni, T. Christianson, T. Sanders, T. Patwardhan, T. Degry, T. Shadwell, T. Fu, T. Gao, T. Garipov, T. Sriskandarajah, T. Sherbakov, T. Korbak, T. Kaftan, T. Hiratsuka, T. Wang, T. Song, T. Zhao, T. Peterson, V. Kharitonov, V. Chernova, V. Kosaraju, V. Kuo, V. Pong, V. Verma, V. Petrov, W. Jiang, W. Zhang, W. Zhou, W. Xie, W. Zhan, W. McCabe, W. DePue, W. Ellsworth, W. Bain, W. Thompson, X. Chen, X. Qi, X. Xiang, X. Shi, Y. Dubois, Y. Yu, Y. Khakbaz, Y. Wu, Y. Qian, Y. T. Lee, Y. Chen, Y. Zhang, Y. Xiong, Y. Tian, Y. Cha, Y. Bai, Y. Yang, Y. Yuan, Y. Li, Y. Zhang, Y. Yang, Y. Jin, Y. Jiang, Y. Wang, Y. Wang, Y. Liu, Z. Stubenvoll, Z. Dou, Z. Wu, and Z. Wang. Openai gpt-5 system card, 2026. URL https://arxiv.org/abs/2601.03267.

[5] G. Comanici, E. Bieber, M. Schaekermann, I. Pasupat, N. Sachdeva, I. Dhillon, M. Blistein, O. Ram, D. Zhang, E. 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.

[6] Team Wan, A. Wang, B. Ai, B. Wen, C. Mao, C.-W. Xie, D. Chen, F. Yu, H. Zhao, J. Yang, J. Zeng, J. Wang, J. Zhang, J. Zhou, J. Wang, J. Chen, K. Zhu, K. Zhao, K. Yan, L. Huang, M. Feng, N. Zhang, P. Li, P. Wu, R. Chu, R. Feng, S. Zhang, S. Sun, T. Fang, T. Wang, T. Gui, T. Weng, T. Shen, W. Lin, W. Wang, W. Wang, W. Zhou, W. Wang, W. Shen, W. Yu, X. Shi, X. Huang, X. Xu, Y. Kou, Y. Lv, Y. Li, Y. Liu, Y. Wang, Y. Zhang, Y. Huang, Y. Li, Y. Wu, Y. Liu, Y. Pan, Y. Zheng, Y. Hong, Y. Shi, Y. Feng, Z. Jiang, Z. Han, Z.-F. Wu, and Z. Liu. Wan: Open and Advanced Large-Scale Video Generative Models. arXiv preprint arXiv:2503.20314, Apr. 2025.

[7] T. Wiedemer, Y. Li, P. Vicol, S. S. Gu, N. Matarese, K. Swersky, B. Kim, P. Jaini, and R. Geirhos. Video models are zero-shot learners and reasoners. arXiv preprint arXiv:2509.20328, Sept. 2025.

[8] Gemini Robotics Team, K. Choromanski, C. Devin, Y. Du, D. Dwibedi, R. Gao, A. Jindal, T. Kipf, S. Kirmani, I. Leal, F. Liu, A. Majumdar, A. Marmon, C. Parada, Y. Rubanova, D. Shah, V. Sindhwani, J. Tan, F. Xia, T. Xiao, S. Yang, W. Yu, and A. Zhou. Evaluating Gemini Robotics Policies in a Veo World Simulator. arXiv preprint arXiv:2512.10675, Jan. 2026.

[9] A. Ajay, S. Han, Y. Du, S. Li, A. Gupta, T. Jaakkola, J. Tenenbaum, L. Kaelbling, A. Srivastava, and P. Agrawal. Compositional foundation models for hierarchical planning. *Advances in Neural Information Processing Systems*, 36:22304–22325, 2023.

[10] Y. Du, S. Yang, P. Florence, F. Xia, A. Wahid, B. Ichter, P. Sermanet, T. Yu, P. Abbeel, J. B. Tenenbaum, L. P. Kaelbling, A. Zeng, and J. Tompson. Video Language Planning. In *The Twelfth International Conference on Learning Representations*, 2024.

[11] H. Li, L. Sun, Y. Hu, D. Ta, J. Barry, G. Konidaris, and J. Fu. NovaFlow: Zero-shot manipulation via actionable flow from generated videos. In *2026 IEEE International Conference on Robotics and Automation (ICRA)*, June 2026.

[12] K. Fang, F. Liu, P. Abbeel, and S. Levine. MOKA: Open-World Robotic Manipulation through Mark-Based Visual Prompting. In *Robotics: Science and Systems XX*, July 2024. doi:10.15607/RSS.2024.XX.062.

[13] S. Ye, Y. Ge, K. Zheng, S. Gao, S. Yu, G. Kurian, S. Indupuru, Y. L. Tan, C. Zhu, J. Xiang, et al. World Action Models are Zero-Shot Policies. *arXiv preprint arXiv:2602.15922*, 2026.

[14] K. Black, N. Brown, J. Darpinian, K. Dhabalia, D. Driess, A. Esmail, M. R. Equi, C. Finn, N. Fusai, M. Y. Galliker, D. Ghosh, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, D. LeBlanc, S. Levine, A. Li-Bell, M. Mothukuri, S. Nair, K. Pertsch, A. Z. Ren, L. X. Shi, L. Smith, J. T. Springenberg, K. Stachowicz, J. Tanner, Q. Vuong, H. Walke, A. Walling, H. Wang, L. Yu, and U. Zhilinsky. $\pi_{0.5}$: a vision-language-action model with open-world generalization. In *Proceedings of the 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pages 17–40. PMLR, Sept. 2025.

[15] Qwen, A. Yang, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Li, D. Liu, F. Huang, H. Wei, H. Lin, J. Yang, J. Tu, J. Zhang, J. Yang, J. Yang, J. Zhou, J. Lin, K. Dang, K. Lu, K. Bao, K. Yang, L. Yu, M. Li, M. Xue, P. Zhang, Q. Zhu, R. Men, R. Lin, T. Li, T. Tang, T. Xia, X. Ren, X. Ren, Y. Fan, Y. Su, Y. Zhang, Y. Wan, Y. Liu, Z. Cui, Z. Zhang, and Z. Qiu. Qwen2.5 Technical Report. arXiv preprint arXiv:2412.15115, Jan. 2025.

[16] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, and G. Lample. LLaMA: Open and Efficient Foundation Language Models. arXiv preprint arXiv:2302.13971, Feb. 2023.

[17] H. Liu, C. Li, Q. Wu, and Y. J. Lee. Visual Instruction Tuning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, Dec. 2023.

[18] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. P. Foster, P. R. Sanketi, Q. Vuong, T. Kollar, B. Burchfiel, R. Tedrake, D. Sadigh, S. Levine, P. Liang, and C. Finn. OpenVLA: An Open-Source Vision-Language-Action Model. In *Conference on Robot Learning (CoRL)*, Jan. 2025.

[19] D. Ghosh, H. R. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, T. Kreiman, C. Xu, J. Luo, Y. L. Tan, L. Y. Chen, Q. Vuong, T. Xiao, P. R. Sanketi, D. Sadigh, C. Finn, and S. Levine. Octo: An Open-Source Generalist Robot Policy. In *Robotics: Science and Systems (RSS)*, volume 20, July 2024. ISBN 979-8-9902848-0-7.

[20] J. Barreiros, A. Beaulieu, A. Bhat, R. Cory, E. Cousineau, H. Dai, C.-H. Fang, K. Hashimoto, M. Z. Irshad, M. Itkina, N. Kuppuswamy, K.-H. Lee, K. Liu, D. McConachie, I. McMahon, H. Nishimura, C. Phillips-Grafflin, C. Richter, P. Shah, K. Srinivasan, B. Wulfe, C. Xu, M. Zhang, A. Alspach, M. Angeles, K. Arora, V. C. Guizilini, A. Castro, D. Chen, T.-S. Chu, S. Creasey, S. Curtis, R. Denitto, E. Dixon, E. Dusel, M. Ferreira, A. Goncalves, G. Gould, D. Guoy, S. Gupta, X. Han, K. Hatch, B. Hathaway, A. Henry, H. Hochsztein, P. Horgan, S. Iwase, D. Jackson, S. Karamcheti, S. Keh, J. Masterjohn, M. Masuda, J. Mercat, P. Miller, P. Mitiguy, T. Nguyen, J. Nimmer, Y. Noguchi, R. Ong, A. Onol, O. Pfannenstiehl, R. Poyner, L. P. M. Rocha, G. Richardson, C. Rodriguez, D. Seale, M. Sherman, M. Smith-Jones, D. Tago, P. Tokmakov, M. Tran, B. Van Hoorick, I. Vasiljevic, S. Zakharov, M. Zolotas, R. Ambrus, K. Fetzer-Borelli, B. Burchfiel, H. Kress-Gazit, S. Feng, S. Ford, and R. Tedrake. A careful examination of large behavior models for multitask dexterous manipulation. *Science Robotics*, 11(113):eaea6201, Apr. 2026. doi:10.1126/scirobotics.aea6201.

[21] L. Sun, B. Xie, Y. Liu, H. Shi, T. Wang, and J. Cao. GeoVLA: Empowering 3D Representations in Vision-Language-Action Models. arXiv preprint arXiv:2508.09071, Aug. 2025.

[22] D. Qu, H. Song, Q. Chen, Y. Yao, X. Ye, J. Gu, Z. Wang, Y. Ding, B. Zhao, D. Wang, and X. Li. SpatialVLA: Exploring spatial representations for visual-language-action models. In *Proceedings of Robotics: Science and Systems*, June 2025. doi:10.15607/RSS.2025.XXI.011.

[23] X. Li, L. Heng, J. Liu, Y. Shen, C. Gu, Z. Liu, H. Chen, N. Han, R. Zhang, H. Tang, S. Zhang, and H. Dong. 3DS-VLA: A 3D Spatial-Aware Vision Language Action Model for Robust Multi-Task Manipulation. In *Proceedings of the 9th Conference on Robot Learning (CoRL)*, volume 305 of *Proceedings of Machine Learning Research*, pages 2344–2359, 2025.

[24] H. Huang, M. Cen, K. Tan, X. Quan, G. Huang, and H. Zhang. GraphCoT-VLA: A 3D Spatial-Aware Reasoning Vision-Language-Action Model for Robotic Manipulation with Ambiguous Instructions. *arXiv preprint arXiv:2508.07650*, 2025.

[25] W. Huang, C. Wang, Y. Li, R. Zhang, and L. Fei-Fei. ReKep: Spatio-Temporal Reasoning of Relational Keypoint Constraints for Robotic Manipulation. In *Conference on Robot Learning (CoRL)*, Jan. 2025.

[26] G. Yin, Y. Li, Y. Wang, D. McConachie, P. Shah, K. Hashimoto, H. Zhang, K. Liu, and Y. Li. CodeDiffuser: Attention-Enhanced Diffusion Policy via VLM-Generated Code for Instruction Ambiguity. In *Robotics: Science and Systems (RSS) XXI*, June 2025. doi:10.15607/RSS.2025.XXI.072.

[27] H. Liu, S. Yao, H. Chen, J. Gao, J. Mao, J.-B. Huang, and Y. Du. SIMPACT: Simulation-enabled action planning using vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20790–20801, June 2026.

[28] L. Fu, J. Yu, K. El-Refai, E. Kou, H. Xue, H. Huang, W. Xiao, F.-F. Li, G. Shi, J. Wu, S. S. Sastry, Y. Zhu, K. Goldberg, and L. Fan. CaP-X: A Framework for Benchmarking and Improving Coding Agents for Robot Manipulation. In *International Conference on Machine Learning (ICML)*, 2026.

[29] Y. Guo, L. X. Shi, J. Chen, and C. Finn. Ctrl-World: A controllable generative world model for robot manipulation. In *The Fourteenth International Conference on Learning Representations*, 2026.

[30] J. Quevedo, A. K. Sharma, Y. Sun, V. Suryavanshi, P. Liang, and S. Yang. WorldGym: World Model as an Environment for Policy Evaluation. In *International Conference on Learning Representations (ICLR)*, 2026.

[31] T. Yuan, Z. Dong, Y. Liu, and H. Zhao. Fast-WAM: Do World Action Models need Test-Time Future Imagination? *arXiv preprint arXiv:2603.16666*, 2026.

[32] Y. Du, S. Yang, B. Dai, H. Dai, O. Nachum, J. Tenenbaum, D. Schuurmans, and P. Abbeel. Learning Universal Policies via Text-Guided Video Generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, Dec. 2023.

[33] C. Luo, Z. Zeng, Y. Du, and C. Sun. Solving New Tasks by Adapting Internet Video Knowledge. In *The Thirteenth International Conference on Learning Representations*, 2025.

[34] H. Bharadhwaj, D. Dwibedi, A. Gupta, S. Tulsiani, C. Doersch, T. Xiao, D. Shah, F. Xia, D. Sadigh, and S. Kirmani. Gen2Act: Human video generation in novel scenarios enables generalizable robot manipulation. In *Proceedings of the 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pages 3936–3951. PMLR, Sept. 2025.

[35] S. Li, Y. Gao, D. Sadigh, and S. Song. Unified video action model. In *Proceedings of Robotics: Science and Systems*, June 2025. doi:10.15607/RSS.2025.XXI.074.

[36] J. Liang, P. Tokmakov, R. Liu, S. Sudhakar, P. Shah, R. Ambrus, and C. Vondrick. Video Generators are Robot Policies. arXiv preprint arXiv:2508.00795, Aug. 2025.

[37] S. Patel, S. Mohan, H. Mai, U. Jain, S. Lazebnik, and Y. Li. Robotic manipulation by imitating generated videos without physical demonstrations. In *The Fourteenth International Conference on Learning Representations*, 2026.

[38] Z.-H. Yin, S. Yang, and P. Abbeel. Object-centric 3D motion field for robot learning from human videos. In *Advances in Neural Information Processing Systems*, volume 38, pages 62307–62327. Curran Associates, Inc., 2025. doi:10.52202/085713-1872.

[39] R. Wang, S. Xu, Y. Dong, Y. Deng, J. Xiang, Z. Lv, G. Sun, X. Tong, and J. Yang. MoGe-2: Accurate monocular geometry with metric scale and sharp details. In *Advances in Neural Information Processing Systems*, volume 38, pages 40377–40408. Curran Associates, Inc., 2025. doi:10.52202/085713-1207.

[40] X. Luo, J.-B. Huang, R. Szeliski, K. Matzen, and J. Kopf. Consistent video depth estimation. *ACM Trans. Graph.*, 39(4):71:1–71:13, Aug. 2020. ISSN 0730-0301. doi:10.1145/3386569.3392377.

[41] B. Zhang, L. Ke, A. Harley, and K. Fragkiadaki. TAPIP3D: Tracking any point in persistent 3D geometry. In *Advances in Neural Information Processing Systems*, volume 38, pages 149980–149999. Curran Associates, Inc., 2025. doi:10.52202/085713-4514.

[42] A. Murali, B. Sundaralingam, Y.-W. Chao, W. Yuan, J. Yamada, M. Carlson, F. Ramos, S. Birchfield, D. Fox, and C. Eppner. GraspGen: A diffusion-based framework for 6-DOF grasping with on-generator training. In *2026 IEEE International Conference on Robotics and Automation (ICRA)*, June 2026.

[43] H. Bharadhwaj, R. Mottaghi, A. Gupta, and S. Tulsiani. Track2Act: Predicting Point Tracks from Internet Videos Enables Generalizable Robot Manipulation. In *European Conference on Computer Vision (ECCV)*, pages 306–324, 2024. ISBN 978-3-031-73116-7. doi:10.1007/978-3-031-73116-7_18.

[44] H. Huang, K. Schmeckpeper, D. Wang, O. Biza, Y. Qian, H. Liu, M. Jia, R. Platt, and R. Walters. IMAGINATION POLICY: Using Generative Point Cloud Models for Learning Manipulation Policies. In *Conference on Robot Learning (CoRL)*, Sept. 2024.

[45] J. Romero, D. Tzionas, and M. J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics*, 36(6):245, 2017. doi:10.1145/3130800.3130883.

[46] G. Pavlakos, D. Shan, I. Radosavovic, A. Kanazawa, D. Fouhey, and J. Malik. Reconstructing Hands in 3D with Transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9826–9836, 2024.

[47] J. Luo, C. Xu, F. Liu, L. Tan, Z. Lin, J. Wu, P. Abbeel, and S. Levine. FMB: A functional manipulation benchmark for generalizable robotic learning. *The International Journal of Robotics Research*, 44(4):592–606, Apr. 2025. doi:10.1177/02783649241276017.

[48] A. Khazatsky, K. Pertsch, S. Nair, A. Balakrishna, S. Dasari, S. Karamcheti, S. Nasiriany, M. K. Srirama, L. Y. Chen, K. Ellis, et al. DROID: A large-scale in-the-wild robot manipulation dataset. In *Proceedings of Robotics: Science and Systems*, July 2024. doi:10.15607/RSS.2024.XX.120.

[49] Z. Li, R. Tucker, F. Cole, Q. Wang, L. Jin, V. Ye, A. Kanazawa, A. Holynski, and N. Snavely. MegaSaM: Accurate, Fast and Robust Structure and Motion from Casual Dynamic Videos. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.

[50] T. Ren, S. Liu, A. Zeng, J. Lin, K. Li, H. Cao, J. Chen, X. Huang, Y. Chen, F. Yan, Z. Zeng, H. Zhang, F. Li, J. Yang, H. Li, Q. Jiang, and L. Zhang. Grounded SAM: Assembling Open-World Models for Diverse Visual Tasks. arXiv preprint arXiv:2401.14159, Jan. 2024.

[51] N. Carion, L. Gustafson, Y.-T. Hu, S. Debnath, R. Hu, D. Suris Coll-Vinent, C. Ryali, K. V. Alwala, H. Khedr, A. Huang, J. Lei, T. Ma, B. Guo, A. Kalla, M. Marks, J. Greer, M. Wang, P. Sun, R. Rädle, T. Afouras, E. Mavroudi, K. Xu, T.-H. Wu, Y. Zhou, L. Momeni, R. Hazra,

S. Ding, S. Vaze, F. Porcher, F. Li, S. Li, A. Kamath, H. K. Cheng, P. Dollár, N. Ravi, K. Saenko, P. Zhang, and C. Feichtenhofer. SAM 3: Segment anything with concepts. In *The Fourteenth International Conference on Learning Representations*, 2026.

[52] LightX2V Contributors. LightX2V: Light video generation inference framework. GitHub repository, 2025.

[53] N. Karaev, Y. Makarov, J. Wang, N. Neverova, A. Vedaldi, and C. Rupprecht. CoTracker3: Simpler and better point tracking by pseudo-labelling real videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6013–6022, Oct. 2025.

[54] W. Kabsch. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A: Crystal Physics, Diffraction, Theoretical and General Crystallography*, 32(5):922–923, Sept. 1976. ISSN 0567-7394. doi:10.1107/S0567739476001873.

Appendix overview. Appendix A presents implementation details and discusses runtime. Appendix B analyzes the hand-flow design and its sensitivity to the switching threshold. Appendix C provides full details of the planning, verification, and recovery procedures, along with the prompts used at each stage. Appendix D describes the geometric grounding procedure, including flow extraction, hand calibration, and grasp selection. Appendix E specifies the baseline setups and oracle information. Appendix F provides additional experiment rollout visualizations, while Appendix G presents example action plans and video-generation prompts. Appendix H reports failure attribution, component ablations, and prompt sensitivity analysis.

A Implementation Details and Runtime Analysis

Hardware setup. We utilize a Franka Research 3 arm equipped with a Robotiq 2F-85 gripper, with RGB and depth images captured by a front-facing RealSense D455 camera fixed on the table. All experiments use a joint impedance controller to track the calculated joint reference.

Implementation details. We upgrade several modules from [11]: MegaSaM [49] to MoGe2 [39] for metric depth recovery and Grounded-SAM2 [50] to SAM3 [51]. We use TAPIP3D [41] for 3D point tracking, HaMeR [46] for hand tracking, Wan2.2 [6, 52] and Veo3.1 [7] for video generation, and GPT-5.2 as the VLM for all reasoning tasks. We use Wan2.2 and Veo3.1 each to generate $L = 8$ rollout videos as candidates for selection. All perception backbones run with their released default inference settings, and the geometric thresholds are held fixed across every task. The maximum rotation magnitude between frames is set as $\theta_{\max} = 45^\circ$ for switching from object to hand flow. We provide an ablation analysis over $\theta_{\max}$ in Section B.

Runtime analysis. We deploy perception-heavy components, including video generation, monocular depth estimation, and 3D point tracking, on an NVIDIA H100 GPU, while hand-flow extraction is performed on an NVIDIA RTX A6000 workstation. Table 2 reports the runtime of the main NovaPlan modules for a 41-frame 720P video, including video generation with Wan2.2 or Veo3.1, monocular depth estimation with MoGe2 or with optional CVD-based depth refinement (see Section D.1), object segmentation with SAM3, object point tracking with TAPIP3D, and hand-flow extraction with HaMeR. With parallelized object- and hand-flow extraction, NovaPlan achieves an end-to-end runtime of approximately 40 s.

Table 2: Runtime breakdown for major modules with a 41-frame 720P video (seconds).

Model	**Wan / Veo**	**MoGe2 / MoGe2+CVD**	**SAM3**	**TAPIP3D**	**HaMeR**
Time	30	3 / 90	3.5	3.5	0.8/frame

End-to-end wall-clock time. Aggregating Table 2 at the task level, video generation dominates planning at roughly 30 s of the $\sim$40 s spent per macro-action, and including network latency and robot execution, each macro-action takes about 90 s. A reactive two- to three-step task therefore takes $\sim$3–4.5 min without recovery and $\sim$4.5–6 min with one recovery cycle. Most runs need no recovery, and when they do it adds one step cycle to the total time. Three-step coupled tasks additionally pay $\sim$2.8 min for strategic beam search before execution begins, giving $\sim$7–9 min in total. Since planning cost is incurred once per macro-action, end-to-end time grows roughly linearly with the task horizon.

B Analysis of Hand-Flow Grounding

How often hand flow is used. Hand flow acts as a fallback, so its value depends on how often it is actually invoked. Across the long-horizon experiments, hand flow is selected for $7/30$ block-stacking actions, $6/30$ color-sorting actions, and $1/25$ hidden-object-search actions. The overall utilization ratio is therefore modest. However, the switches concentrate on the steps that are the hardest and would otherwise terminate a trial, e.g., placing the fourth block or inserting the tightly

fitting yellow block. This is consistent with the hybrid-grounding ablation in Table 1(a), where disabling the switch lowers long-horizon success from 76.7% to 46.7%.

Why a human hand. We use a human hand, rather than a robot end-effector, as the source morphology for video generation for two reasons. First, current video generation models produce physically plausible human manipulation far more reliably than robot arm motion, since human-related data are more common than robot data in their training distribution. Second, the hand stays partially visible and informative even when the manipulated object is occluded, which is exactly the regime in which object flow becomes unreliable. We emphasize that NovaPlan does not imitate human joint angles, so only the calibrated task-space $SE(3)$ motion of the designated contact fingertip is retargeted to the gripper, as described in Section D.2.

Sensitivity to the switching threshold. We analyze the sensitivity of the hand-flow switching threshold $\theta_{\max}$, which controls when NovaPlan will switch from object-flow-based execution to hand-flow-based execution. In our main experiments, we use $\theta_{\max} = 45^\circ$ as a simple heuristic default corresponding to a visually large inter-frame rotation, rather than as a heavily tuned hyperparameter.

To evaluate whether performance depends sensitively on this choice, we conduct a sensitivity analysis over $\theta_{\max} \in \{20^\circ, 30^\circ, 45^\circ, 60^\circ\}$, along with an object-flow-only variant ($\theta_{\max} = \infty$) that disables hand-flow switching. We evaluate on the two most failure-prone steps in the long-horizon experiments and report single-step success rates over 10 trials per setting.

Across both tasks, enabling a switch to hand flow consistently outperforms object-flow-only execution. Thresholds in the 20°–45° range perform similarly well, suggesting that NovaPlan is not highly sensitive to the exact value of $\theta_{\max}$. Performance degrades slightly at 60°, likely because the system is too tolerant to switch away from unreliable object flow. These results support our key design choice of allowing execution to switch from object flow to hand flow when object tracking becomes unreliable, while indicating that $\theta_{\max}$ itself is not a fragile threshold.

Table 3: Ablation study of the hand-flow switching threshold $\theta_{\max}$ on single-step success rates.

$\theta_{\max}$	**20°**	**30°**	**45°**	**60°**	**∞ (Object-flow Only)**
4-layer Block Stacking: 4th Block	70%	60%	70%	60%	40%
Color Sorting: Yellow Block	70%	60%	60%	40%	20%

C Vision-Language Planning Implementation Details

To ensure robust long-horizon reasoning, our high-level planner operates within a physics-compliant reasoning framework established via structured prompt conditioning. Rather than relying on unconstrained generation, this approach restricts the model's output space to adhere to environmental affordances and causal dependencies. The following sections detail the specific reasoning strategies governing our VLM prompting.

C.1 Constraint-Aware Task Decomposition and Action Proposal

The high-level planner takes as input the task goal, the latest scene observation, and the action history. It decomposes the goal into a sequence of macro-actions. To reduce physically invalid proposals and redundant moves, we use a structured prompt that enforces basic physical constraints and ordering rules.

C.1.1 Dynamic Context and Prompt Structure

The VLM action proposal is conditioned on the current observation, the task goal, and the full action history. It also receives a small amount of step-dependent context: the remaining planning horizon T_{remain} and an optional *constraint context* carried over from the previous step (when available). The prompt encodes three types of guidance.

- **Horizon-dependent urgency.** We add a short note based on T_{remain}. If $T_{remain} \leq 2$, we append *"Time is critical. Actions must be aggressive."* Otherwise, we append *"You have time. Prioritize precision and safety."*
- **State awareness to avoid redundant moves.** The model is instructed to treat objects that are already in their final configuration as *completed* and to avoid moving them, while treating misplaced objects that overlap target zones as potential obstructions.
- **Ordering constraints.** When two objects interact, the model is instructed to test both orders. If executing A then B is feasible but B then A blocks access, then A is enforced as the prerequisite, and all proposed actions must start with A.

Prompt for Task Decomposition and Action Proposal

Analyze the scene and history. Propose EXACTLY {num_actions} distinct narrative macro-actions.
Goal: {goal}
Action History: {prev_act_str}
{urgency_note}
Constraint Context: {constraint_context}
Critical visual analysis (intersections & hierarchy)

1. **Trace target zones:** Visually trace the target locations (slots, containers, stacking areas) for *all* objects.
2. **Detect overlaps:** Do any target zones visually cross or intersect?
3. **Determine hierarchy (access & occlusion):**
   - **Occlusion principle:** If placing object B visually ''covers'' or ''bridges'' the access path to object A's target, then object A must be placed first.
   - **Constrained access:** Placing objects in *deep* or *central* zones first is safer than placing peripheral objects that may narrow the workspace.
   - **Generic rule:** Place objects that sit *lowest* or *deepest* in the assembly structure first.

State verification (goal state vs. obstruction)

- **Analyze objects in the workspace:** Distinguish objects that are *correctly placed* from those that are *clutter/obstacles*.
- **Correctly placed:** If an object is fully aligned with its target (e.g., inside its slot, neatly stacked, or in its final position), treat it as a completed step, not an obstruction. It usually does not need to be moved.
- **Misplaced/loose:** If an object lies haphazardly across target zones or rests on top of other destinations without being properly placed, it is an obstruction that might need to be moved.
- **Correction:** Do not propose removing objects that are already in their correct final state unless they clearly block a necessary path.

Mental simulation (''order of operations'' test)
Compare potential sequences for interacting objects (A and B):

1. **Sequence 1 (A then B):** ''If I place A first, can I still place B?''
2. **Sequence 2 (B then A):** ''If I place B first, can I still place A?''

Critical conclusion

- If (A then B) works but (B then A) fails because B blocks A's path/slot, then A is the prerequisite. You must propose A first.
- **Generic indicator:** The prerequisite object often lies ''under,'' ''behind,'' or ''inside'' the other object's zone, or its insertion path is crossed by the other object.
- **Risk assessment:** If paths cross or zones overlap, assume strict ordering is required until proven otherwise.

Selection filter

- Only propose actions that pass the ''order of operations'' test.
- If multiple valid actions exist (independent sub-tasks), propose diverse options.
- **Mandatory filter:** If you identify a prerequisite object (e.g., A must be first), then every single one of the {num_actions} proposals must start with A.
- Do not offer options that violate the prerequisite. If A is the bottleneck, diversity should come from *how* you handle A (e.g., different grasp or approach), not from proposing B or C.

Step 1: phase identification
Where are we? (approaching → interacting → transporting → finishing).
Step 2: apply constraints (critical)

```
   1. No telekinesis (active voice): Objects never move on their own.
        • Bad: ''The pot moves to the drawer.''
        • Good: ''Grasp the pot and move it.''
   2. Abstraction level:
        • Bad: ''Move forward 10 cm.''
        • Good: ''Reach down and grasp the handle.''
   3. Single hand only: The robot has one right hand. Do not use bimanual actions.
   4. Pacing & termination:
        • Do not complete the task instantly if the object is far; approach first.
        • If done: output FINISH: Hold position.
   5. Action description format:
        • Do not mention ''the robot'' or ''the robot right hand.''
        • Start directly with the verb.
        • Be definitive; do not use ''/'' or ''or'' inside an action.
        • Bad: ''Grasp the block from the top/side.''
        • Good: ''Grasp the block from the top.''
Step 3: generate {num_actions} strategies
For each action, identify the primary object being manipulated or interacted with. This
object will be tracked visually during video generation.
Output format (JSON only). You must provide EXACTLY {num_actions} actions:
{
  "phase": "Current Phase Analysis...",
  "dependency_analysis": "Explain the 'Future Block' test results.",
  "valid_objects": ["List of objects that are safe to manipulate NOW"],
  "proposals": [
    {"action": "Action description 1", "track_object": "object being manipulated (e.g.,
drawer handle, cup, lid)", "reasoning": "Why this is valid now?"},
    {"action": "Action description 2", "track_object": "object being manipulated",
"reasoning": "Why this is valid now?"},
    ...
  ]
}

Important: the proposals array must contain EXACTLY {num_actions} items. Each item must
include both action and track_object.
```

We post-process the returned actions to remove boilerplate prefixes (e.g., "the robot ..."), enforce single-hand phrasing, and ensure that exactly the desired number of actions is produced. In greedy mode (see Section 3.1), we set $N_c = 1$ and propose the next viable action based on the latest observation. In strategic mode, we perform tree search and use the video generation model as an object dynamics verifier to select a full action plan in advance; details are given in Section C.1.2. Section G shows example multi-step plans for all long-horizon tasks in our experiments.

C.1.2 Full Action Plan Generation in Strategic Mode

As described in Section 3.1, for coupled tasks we use a strategic beam search to generate a single feasible plan over the full horizon before execution. At each step in strategic mode, we ask the VLM to propose $N_c = 2$ action candidates per step and sample $L = 4$ video rollouts per action candidate. We then select the top $A = 2$ beams whose rollouts receive the highest rankings and use them to initialize the next step. This process continues until the final step, where the best beam is selected to form the full-horizon plan. The full procedure is summarized in Algorithm 2.

C.2 Physics-Based Video Rollout Selection

To select a rollout for execution, we use a VLM as an automatic evaluator and apply a hierarchical scoring rubric over the generated candidates. In our experiments, for each proposed action at each step during execution, we use Wan2.2 and Veo3.1 each to generate $L = 8$ rollout videos as candidates for selection.

To reduce evaluation cost, we batch-score candidates by stitching them into a single grid image. Each tile corresponds to one candidate rollout and is labeled with its original candidate index (ID).

Algorithm 2 Strategic Beam Search for Full Task Plan Generation

Input: start frame I_0, goal g, horizon H
Input: candidate set size A, actions per candidate N_c, rollouts per action L
Input: video generation model $\mathcal{G}$, VLM $\mathcal{V}$
Output: best candidate $b^\star$

candidate state. each candidate (beam) b stores:
current frame I_b; action history $\mathcal{H}_b$; metadata list $\mathcal{M}_b$; score s_b.
metadata. each entry $m \in \mathcal{M}_b$ is a per-step dictionary used for state carryover, e.g.,
{action: a, track_object: o, candidate_id: id, VLM_context: ξ}.
context carryover. GETCONTEXT($\mathcal{M}_b$) returns the most recent VLM_context (or $\emptyset$) to enforce constraints consistently across steps.

$\mathcal{B} \leftarrow \{ (I_0, \emptyset, \emptyset, 0) \}$
for $t \leftarrow 1$ **to** H **do**
  $\mathcal{C} \leftarrow \emptyset$ ▷ all rollout candidates at step t
  $id \leftarrow 0$
  for all $b = (I_b, \mathcal{H}_b, \mathcal{M}_b, s_b) \in \mathcal{B}$ **do**
    $h \leftarrow$ FORMATHISTORY($\mathcal{H}_b$)
    $c \leftarrow$ GETCONTEXT($\mathcal{M}_b$)
    $\Pi \leftarrow \mathcal{V}$.PROPOSEACTIONS($I_b, g, N_c, h, H - t + 1, c$) ▷ Π is a list of proposals $\{(a_i, o_i, \xi_i)\}_{i=1}^{N_c}$
    for $i \leftarrow 1$ **to** $|\Pi|$ **do**
      $a \leftarrow \Pi[i]$.action; $o \leftarrow \Pi[i]$.track_object; $\xi \leftarrow \Pi[i]$.VLM_context
      $\{V^{(1)}, \ldots, V^{(L)}\} \leftarrow \mathcal{G}$.GENERATEROLLOUTS($I_b, a, L$) ▷ $V^{(k)}$ is the k^{th} generated video
      for $l \leftarrow 1$ **to** L **do**
        $I' \leftarrow$ LASTFRAME($V^{(l)}$)
        $F \leftarrow$ FLOW($V^{(l)}, o$)
        $\mathcal{C} \leftarrow \mathcal{C} \cup \{(id, b, a, o, \xi, V^{(l)}, F, I')\}$
        $id \leftarrow id + 1$
      end for
    end for
  end for
  if $\mathcal{C} = \emptyset$ **then**
    break
  end if
  $\mathcal{R} \leftarrow \mathcal{V}$.RANKROLLOUTSBATCH($g, \mathcal{C}$, top_n $= A$) ▷ VLM scores candidates using a stitched grid of (flow, final-frame) tiles
  initialize default score map $s(\cdot) \leftarrow s_{min}$ for all candidate ids
  for all $r \in \mathcal{R}$ **do**
    $s(r.\text{candidate_id}) \leftarrow r.\text{score}$
  end for
  $\mathcal{B}_{next} \leftarrow \emptyset$
  for all $(id, b, a, o, \xi, V, F, I') \in \mathcal{C}$ **do**
    $\mathcal{H}' \leftarrow \mathcal{H}_b \cup \{a\}$
    $m' \leftarrow \{ \text{action} : a,\ \text{track_object} : o,\ \text{candidate_id} : id,\ \text{VLM_context} : \xi \}$
    $\mathcal{M}' \leftarrow \mathcal{M}_b \cup \{m'\}$
    $b' \leftarrow (I', \mathcal{H}', \mathcal{M}', s(id))$
    $\mathcal{B}_{next} \leftarrow \mathcal{B}_{next} \cup \{b'\}$
  end for
  $\mathcal{B} \leftarrow$ TOPB($\mathcal{B}_{next}, A$) ▷ sort by candidate score and keep best A
end for
$b^\star \leftarrow \arg\max_{b \in \mathcal{B}} s_b$
return $b^\star$

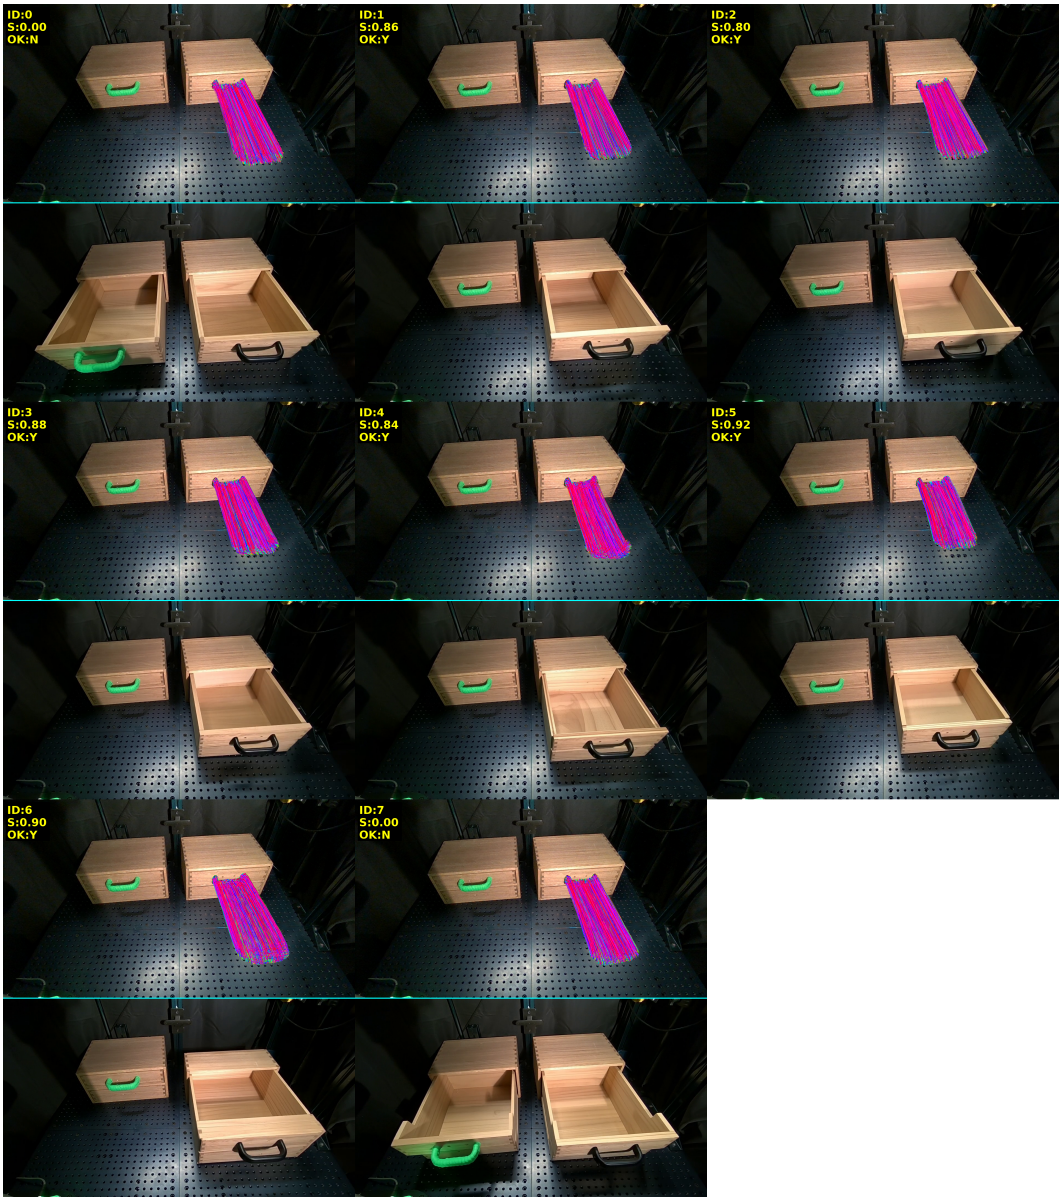


Figure 8: **Example grid image for video rollout selection.** The grid contains eight candidate video rollouts for the hidden object search task at the first step ("open the right drawer").

Concretely, each tile contains (i) a flow image (2D object motion flow overlaid on the start frame from CoTracker3 [53]) and (ii) the last frame of the generated video, stacked vertically and separated by a cyan divider. Given the grid and the associated action text for each ID, the VLM assigns a score to every candidate and follows a strict hierarchy of checks, stopping at the first violated constraint:

- **Target check.** Did the correct object move? If the wrong object moves or there is no motion, the score is 0.0.
- **Physics check.** Is the motion physically plausible (no melting, deformation, teleportation, or objects appearing from nowhere)? If violated, the score is at most 0.1.
- **Motion check.** Does the motion direction and trajectory in the flow visualization match the action language instruction? If violated, the score is at most 0.2.
- **Result check.** Does the final frame match the intended outcome and make progress toward the goal? If satisfied, the score lies in $[0.3, 1.0]$.

In Figure 8 we show an example of the grid image, where each tile is labeled with the returned score and success flag.

Prompt for Video Rollout Selection

```
Rank {N} robot rollouts.
Goal: ``{goal}''
IMAGE: Grid with {N} tiles. Each tile has TOP=motion flow overlay on initial frame, CYAN
LINE=divider, BOTTOM=final frame. Yellow ``ID: k'' label on each tile.
Actions:
{ID 0: ``...''}
{ID 1: ``...''}
...
SCORING RULES (Strict Hierarchical Check - apply in order, stop at first failure):
1. Target Check: Did the CORRECT object move? If WRONG object moved or NO motion at all
→ score=0.0
2. Physics Check: Is physics realistic? If object MELTS/DEFORMS/TELEPORTS or NEW objects
appear from nowhere → score≤0.1
3. Motion Check: Look at TOP (flow overlay on initial frame). Does the motion
direction/trajectory match what the ACTION describes?
      • NO (wrong direction/drifted/stuck/incomplete path) → score≤0.2, ok=false
      • YES (motion trajectory matches action description) → continue
4. Result Check: Look at BOTTOM (final frame, below cyan line). Does it show the
expected outcome of the ACTION?
      • NO (object deformed/new objects appeared/wrong final position/action not completed)
        → score≤0.2, ok=false
      • YES (final state matches what action intended, progresses toward goal) → score
        0.3--1.0, ok=true
OUTPUT: JSON only. Score ALL candidates.
{
  ``rankings'': [
   {``candidate_id'': 0, ``success'': true, ``score'': 0.9, ``reason'': ``perfect
insertion''},
   {``candidate_id'': 1, ``success'': false, ``score'': 0.1, ``reason'': ``wrong object
moved''}
  ]
}
Key: candidate_id=ID (0-based index matching the grid labels), success=boolean,
score=float(0--1), reason=string rationale
```

C.3 Closed-Loop State Verification

After each executed action, NovaPlan verifies the outcome with a VLM-based critic. Unlike open-loop pipelines that assume success, we explicitly evaluate the observed state transition to detect action-level failures, such as grasp slippage, partial insertions, or acting on the wrong object.

C.3.1 State Verification

Given an intended action and goal, the critic compares three images:

- **Start state (I_t):** the observation before execution.
- **Current state (I_{t+1}):** the observation after the robot stops.
- **Target state (I_{target}):** the last frame of the generated video plan.

The planned video provides a concrete visual reference for the expected outcome. The critic checks whether the change observed in the real world ($I_t \rightarrow I_{t+1}$) is consistent with the intended change in the video plan ($I_t \rightarrow I_{target}$). This comparison distinguishes *no-change* outcomes (the action did not take effect) from *incorrect-change* outcomes (a change occurred but the result is wrong), such as moving the wrong object or stopping before the action is completed.

Prompt for State Verification

```
You are a Robotic Systems Verifier. Your task is to determine whether the CURRENT STATE
(Image 2) satisfies the physical requirements of the intended ACTION, given the START STATE
(Image 1).
```

INTENDED ACTION: {action}
GOAL SPECIFICATION: {goal}
You are provided with three images:
Image 1 (START STATE): state before the action.
Image 2 (CURRENT STATE): state after attempting the action.
Image 3 (TARGET STATE): expected outcome from the video plan.
Verification logic:

1. **Analyze the action.** List the physical constraints implied by ``{action}'' (e.g., *insert* ⇒ inside a cavity and flush; *place* ⇒ resting on a surface; *grasp* ⇒ firm contact).
2. **Compare Image 1 vs. Image 2.** What changed? Did the intended object move as expected?
3. **Inspect Image 2.** Are the action-specific physical constraints satisfied in the current state?
4. **Compare Image 2 vs. Image 3.** Using Image 3 as a visual reference, is the object state close enough to the target outcome?

OUTPUT (JSON only):
{ ``success'': boolean, ``reason'': string }
The reason should identify which physical constraint is (or is not) satisfied, and whether the change from Image 1 to Image 2 matches the intended change from Image 1 to Image 3.

C.3.2 Recovery Action Proposal and Video Generation

A core capability of NovaPlan is recovery from runtime errors. When the verifier returns `"success": false`, the system enters a corrective planning loop. We use hierarchical prompting that prioritizes local corrections over full resets when the remaining discrepancy is small and the goal remains achievable. This design reduces unnecessary motion and limits error accumulation that can arise from repeatedly replanning full pick-and-place sequences.

Recovery mode selection. The VLM first serves as a decision-maker. It compares the current observation with the target state I_{target} and selects a recovery strategy $S \in \{\text{grasp}, \text{non-prehensile}\}$. If grasp recovery is selected, the system proceeds with the standard planning pipeline (i.e., a normal non-recovery step). Otherwise, it invokes the non-prehensile recovery procedure described below.

Spatial grounding and annotation. When non-prehensile recovery is selected, the system grounds the abstract corrective plan in the scene geometry. We query the VLM to identify a single contact region where a small external interaction (e.g., a poke or nudge) can move the manipulated object toward the target configuration. To bridge the semantic planner and the video generation model, the VLM produces two outputs:

- **Visual annotation.** The VLM edits the latest observation image by overlaying a visual marker (e.g., a red star) at the selected contact point, yielding $I_{t+1,anno}$. See Figure 9 for an example in which the VLM marks the contact point between the purple block and a human index finger with a red star for a corrective poke.
- **Textual guidance.** The VLM generates a detailed prompt P that specifies the desired correction video, including the scene setup, the finger contact point, the hand motion trajectory, and the intended final state.

Prompt for Recovery Action Proposal and Video Generation

You are a Recovery-Policy Controller for a manipulation task. You must (i) select a recovery mode and (ii) if non-prehensile recovery is chosen, spatially ground a single contact point, return an annotated image, and synthesize a structured prompt for video generation.
INPUTS (provided):
• Current observation and image to annotate: **Image 1**
• Goal / target image: **Image 2**
• Goal text: {goal_text}
TASK 1: Recovery Mode Selection.
Compare **Image 1** with **Image 2** and choose

$$S \in \{\text{grasp}, \text{non_prehensile}\}.$$

Choose non_prehensile if a small perturbation (e.g., a gentle poke/nudge) can plausibly correct the remaining discrepancy; otherwise choose grasp. Provide a brief justification grounded in visible evidence.

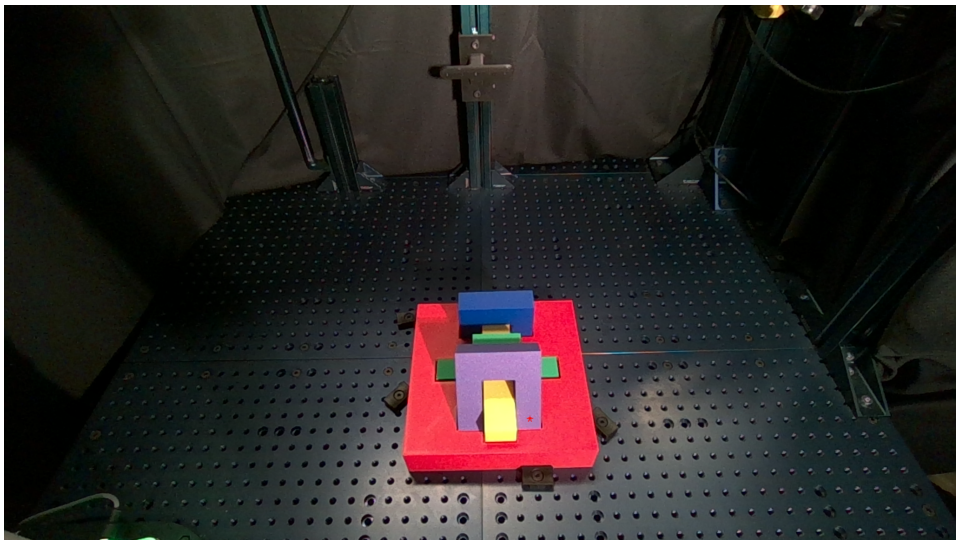


Figure 9: **Example annotated start image for recovery video generation**. The VLM marks the contact point between the purple block and the human index finger with a red star on the lower right part of the purple block, indicating where the corrective poke should occur.

TASK 2: Spatial Grounding + Annotation (only if $S =$ non_prehensile).
Identify *one* manipulated object and *one* contact point that will nudge the object toward **Image 2**.

- Specify: object_name, contact_surface, contact_location.
- Define a **single anchor** contact point (no alternatives).
- Overlay a **small solid red star** centered exactly at the contact point on **Image 1**, producing the **annotated Image 1**.
- Return the **annotated Image 1** in the response as a base64-encoded PNG (no surrounding text).

TASK 3: Textual Guidance Prompt P (only if S is non_prehensile).
Write a detailed prompt P for a first-last-frame conditioned video generator, using:

$$\text{START} = \textbf{annotated Image 1}, \quad \text{GOAL} = \textbf{Image 2}.$$

P must be highly structured with labeled sections and must include:

- Global constraints (single continuous shot, fixed camera, no cuts; last frame matches goal exactly).
- Scene description (what is present; what must remain still / never be touched).
- Framing constraints (camera view, margins, full visibility requirements).
- Contact target (single anchor; explicitly tie to the red star center).
- Hard spatial constraints (where the hand may move; what must remain in frame).
- Hand appearance + stability constraints (realistic proportions; no warping; stable pose).
- Hand approach constraints (straight-line approach; height constraints; no arcs).
- Continuous-contact rule (no sliding; no switching surfaces; no second contact point).
- Action description (direction and components; intended rotation/settling toward goal).
- End behavior (pause, retract, exit; no second push; realistic contact physics; silent).

Use unambiguous language (*must/only/never*), include **exactly one** contact point and **exactly one** push, do not introduce new objects, and keep all untouched objects unchanged.
OUTPUT (strict JSON only). Return *only* valid JSON with exactly this schema:

```
{
  "recovery_mode": "grasp" | "non_prehensile",
  "mode_justification": {
    "discrepancy_summary": "string (1{3 sentences)",
    "why_this_mode": "string (1{3 sentences)"
  },
  "annotation": {
    "enabled": boolean,
    "object_name": "string" | null,
```

```
    "contact_point_definition": "string" | null,
    "edit_spec": {
      "marker": "solid_red_star",
      "anchor": "center_of_star_is_contact_point",
      "placement": "string"
    },
    "annotated_image_png_base64": "string" | null
  },
  "prompt_P": {
    "enabled": boolean,
    "title": "string",
    "text": "string (full structured prompt P)"
  }
}
If recovery_mode=grasp, set annotation.enabled=false and prompt_P.enabled=false, and set
image/prompt fields to null.
```

Conditioned video generation. The final recovery video is generated by a video generation model. We use Veo3.1 to generate the recovery video in the First-Last-Frame setting. Specifically, the model generates a recovery video V conditioned on the annotated start image $I_{t+1,anno}$, the proposed correction prompt P, and the target image I_{target} as the required last frame:

$$V = \text{VideoModel}(I_{t+1,anno}, P, I_{target}) \tag{1}$$

Using $I_{t+1,anno}$ rather than the raw observation constrains the video model to the specific contact dynamics selected by the VLM. This yields a targeted non-prehensile correction (e.g., a single poke) that steers the scene toward the target configuration without requiring a full reset. Here, we show the example video prompt for generating the recovery video of a human index finger poking the purple block for the fourth step of the FMB variant:

Example Prompt for Non-prehensile Recovery Video Generation

```
Use the provided START image as the first frame and the provided GOAL image as the last
frame.  The last frame must match the provided goal image:  same camera view, same object
poses.  Single continuous shot, fixed camera, no cuts.
Scene:  slightly top-down view of a perforated metal workbench.  A red puzzle base is
centered.  A purple U-shaped block sits on the base.  A yellow rectangular bar stays
perfectly still.  The green and blue pieces behind remain perfectly still and are never
touched.
Framing:  medium-wide shot with extra margin around the red base; the entire red base is
visible with space around it.
Contact target (single anchor, no alternatives):  A small solid red star is painted on the
outer camera-right-facing vertical side wall of the purple block's right leg, very close to
the bottom edge near the red base surface.  The center of this red star is the only contact
point of the human index finger on the purple block.
Hard spatial constraint:  All visible hand motion stays in the front half of the red base
(closest to the bottom of the frame) for the entire shot and never moves into the back
half.  Make sure the full hand, including the palm and the thumb, is always visible in the
view.
Hand appearance + stability (reduce DoF): It should be a normal adult human hand (realistic
proportions), not oversized, not close to the camera.  The hand should enter the scene from
the right side of the screen.  The index finger extended and the full hand always visible
in the view.  The fingertip shape and perspective stay consistent (no warping).  The finger
does not twist or rotate; the pose is stable.
Hand approach (straight line, no vertical motion):  The index finger is low, nearly
parallel to the tabletop, and approaches horizontally along a single straight line at
the same height as the red star on the purple block.  The finger never descends from above,
never arcs, and never rises above the purple block's top edge.
Continuous-contact rule (no sliding, no switching surfaces):  The fingertip pad presses
onto the center of the red star on the bottom edge of the purple block's right leg and
fully covers/occludes the star during contact.  From first contact until the purple block
is fully seated, the fingertip stays glued to that exact star location (no sliding upward,
no rolling over an edge, no moving to the top face, no second contact point).  The top face
of the purple block is never touched at any time.
Action (correct tilt direction, one push only):  With the fingertip still glued to the red
star on the bottom edge of the purple block's right leg, perform one smooth inward push
from camera-right to camera-left with a slight downward component.  This makes the purple
block rotate forward toward the camera (toward the bottom of the frame) like a hinge:  the
front edge (bottom of frame) goes down into the socket first while the back edge (top of
frame) rises slightly.  Immediately after this forward tilt, the block drops/slides into
the socket underneath and becomes fully seated/flush exactly like the goal image.
End:  As soon as it seats, the finger stops pushing, pauses for a fraction of a second
still at the same low height, then retracts straight back to camera-right along the same
straight line and exits.  No second push, no re-contact elsewhere.  Realistic contact
physics.  Silent.
```

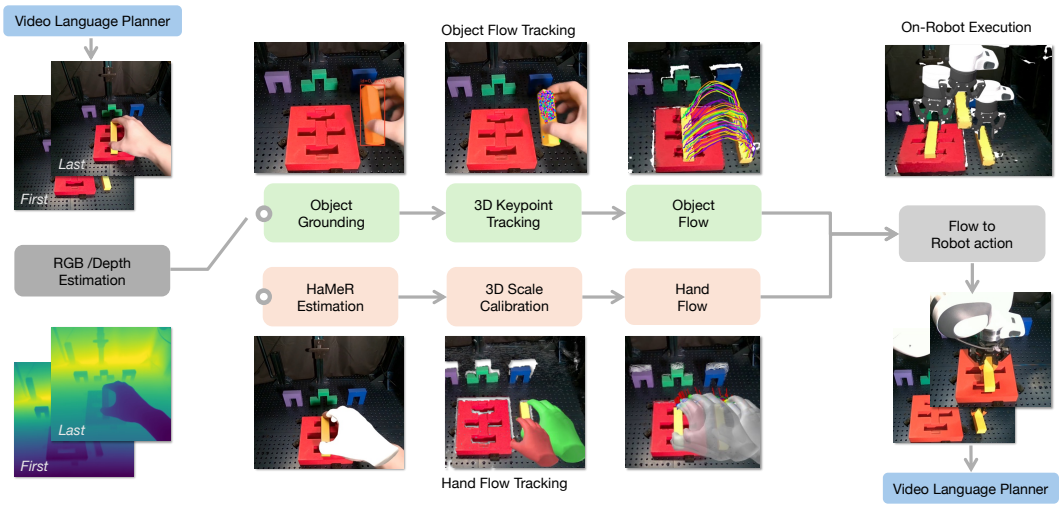


Figure 10: **NovaPlan geometric grounding overview.** Given the generated video plan (RGB+Depth), we switch between (top) grounding the target object, tracking 3D keypoints, and recovering object flow, and (bottom) estimating the hand pose with HaMeR, calibrating 3D scale, and computing the hand flow. The resulting object/hand flows are converted into robot actions for on-robot execution, with a video language planner enabling closed-loop re-planning and recovery.

D Geometric Grounding Implementation Details

This section provides implementation details for the geometric grounding module described in Section 3.2. Following the components outlined in Figure 10, we describe metric object-flow extraction, hand-flow calibration, contact-constrained grounding for non-prehensile repair, and grasp selection.

D.1 Metric Depth and Object Flow Extraction

Given a generated RGB video and the target object specified by the planner, we extract a metric object flow by combining monocular depth prediction, temporal depth refinement, object segmentation, and 3D point tracking.

Metric depth recovery. We estimate per-frame depth from the generated RGB video using MoGe2 [39]. Given known camera intrinsics, we pass the intrinsics to the model during inference so that the predicted depth is tied to the physical camera geometry. To reduce frame-to-frame depth flicker, we optionally refine the MoGe2 depth sequence with Consistent Video Depth (CVD) [40]. CVD minimizes optical-flow-induced geometric reprojection errors and regularizes the solution with gradient and normal-consistency terms. In our implementation, we set the gradient weight $w_{\text{grad}} = 2.0$ and the normal-consistency weight $w_{\text{normal}} = 5.0$, and we run the optimization at a fixed processing budget of 196608 pixels. Because MoGe2 already provides metric depth, we freeze the global shift term during CVD to preserve the metric offset while improving temporal coherence.

RANSAC-based affine calibration. To align the generated depth D_{gen} with the reference sensor depth D_{sensor} for the first frame, we estimate a global affine transformation:

$$D_{\text{metric}}(\mathbf{u}) = s_{\text{depth}} D_{\text{gen}}(\mathbf{u}) + t_{\text{depth}}, \tag{2}$$

where $\mathbf{u}$ denotes an image pixel. We estimate $(s_{\text{depth}}, t_{\text{depth}})$ with RANSAC to reduce sensitivity to invalid sensor measurements, specular regions, and occasional depth prediction failures.

We first build a calibration mask $\mathcal{M}$ that keeps pixels with finite, positive depth in both D_{gen} and D_{sensor}. To suppress isolated sensor artifacts, we remove small connected components in the valid-depth mask with area smaller than 50 pixels. We then run RANSAC for 1000 iterations to find the

affine parameters that maximize the number of inliers:

$$(s_{\text{depth}}, t_{\text{depth}}) = \underset{s,t}{\text{argmax}} \sum_{\mathbf{u}\in\mathcal{M}} \mathbb{I}(|sD_{\text{gen}}(\mathbf{u}) + t - D_{\text{sensor}}(\mathbf{u})| < \tau)\,, \tag{3}$$

where $\tau = 0.15$ by default. Finally, we re-fit $(s_{\text{depth}}, t_{\text{depth}})$ using all RANSAC inliers and apply the calibrated parameters to the full generated depth sequence.

3D point tracking. After depth recovery, we obtain per-frame target-object masks and sample K keypoints on the object in the initial frame. We track these points through the generated video using a 3D dense point tracker [41], producing the object flow:

$$\mathcal{F}_{\text{obj}} = \{\mathbf{f}_i^t \in \mathbb{R}^3 \mid i = 1, \ldots, K,\ t = 1, \ldots, T\}. \tag{4}$$

Rigid object-motion fitting. Given tracked object points at consecutive frames, we estimate the rigid object motion $(\mathbf{R}^t, \mathbf{t}^t)$ by aligning $\{\mathbf{f}_i^{t-1}\}_{i=1}^K$ and $\{\mathbf{f}_i^t\}_{i=1}^K$. Let $\mathbf{c}^{t-1}$ and $\mathbf{c}^t$ be the centroids of the two point sets. We solve for rotation using the Kabsch algorithm [54]:

$$\mathbf{R}^t = \underset{\mathbf{R}\in SO(3)}{\text{argmin}} \sum_{i=1}^{K} \left\|\mathbf{R}(\mathbf{f}_i^{t-1} - \mathbf{c}^{t-1}) - (\mathbf{f}_i^t - \mathbf{c}^t)\right\|^2. \tag{5}$$

The translation is then computed as

$$\mathbf{t}^t = \mathbf{c}^t - \mathbf{R}^t\mathbf{c}^{t-1}. \tag{6}$$

The frame-to-frame transforms are composed to obtain the object trajectory $\{T_{\text{obj}}^t\}_{t=1}^T$.

D.2 Hand Flow Calibration

Generated human hands provide useful kinematic priors, but raw hand reconstructions are often not directly metric due to scale hallucination, floating contacts, and projective drift. We therefore calibrate the reconstructed hand trajectory using object-contact constraints.

Interaction interval. We first identify the interval during which the generated hand interacts with the target object. Assuming the object is static before contact, we define the contact onset t_{start} as the first frame where the object mask changes relative to the initial mask by more than a threshold ϵ:

$$t_{\text{start}} = \underset{t}{\text{argmin}} \left\{ \frac{\left|\mathbf{M}_{\text{obj}}^t \setminus \mathbf{M}_{\text{obj}}^{t_0}\right|}{\left|\mathbf{M}_{\text{obj}}^{t_0}\right|} \geq \epsilon \right\}, \tag{7}$$

where $\mathbf{M}_{\text{obj}}^t$ is the object mask at frame t and t_0 is the first frame. The contact end t_{end} is defined analogously as the last frame satisfying the same criterion. For grasp-mode hand flow, we use $\epsilon = 0.9$.

Scale recovery at contact onset. At t_{start}, we recover a metric scale that brings the reconstructed hand into contact with the object. We project the reconstructed hand skeleton into the image plane and select fingertips whose projections overlap with the object mask $\mathbf{M}_{\text{obj}}^{t_{\text{start}}}$. To allow small projection error, each fingertip projection is represented as a circle with radius 15 pixels. This produces a candidate contact set $\mathcal{F}_{\text{contact}}$. For each fingertip $f \in \mathcal{F}_{\text{contact}}$, we estimate an isotropic scale s_f by snapping the fingertip to the center of the corresponding metric object points. We use

$$s_{\text{start}} = \max_{f\in\mathcal{F}_{\text{contact}}} s_f \tag{8}$$

as the global hand scale to avoid under-scaling from partially occluded fingertips. The fingertip attaining this maximum is used as the designated contact finger for subsequent drift compensation.

Drift compensation at release. Even after global scale correction, generated videos can exhibit projective drift as the hand moves relative to the camera. We therefore re-estimate the contact-based scale at release, denoted as s_{end}, and compute a drift offset for the designated contact fingertip:

$$\mathbf{d}_{\text{err}} = \mathbf{p}_{\text{tip}}^{t_{\text{end}}}(s_{\text{end}}) - \mathbf{p}_{\text{tip}}^{t_{\text{end}}}(s_{\text{start}}). \tag{9}$$

To avoid perturbing the approach and transport phases, this correction is only applied near release. We define t_{corr} as the first frame whose globally scaled fingertip comes within distance δ of its release configuration:

$$t_{\text{corr}} = \underset{t}{\text{argmin}} \left\{ \left\| \mathbf{p}_{\text{tip}}^{t}(s_{\text{start}}) - \mathbf{p}_{\text{tip}}^{t_{\text{end}}}(s_{\text{start}}) \right\|_2 < \delta \right\}. \tag{10}$$

We then linearly ramp the offset over $[t_{\text{corr}}, t_{\text{end}}]$:

$$\begin{aligned} \alpha(t) &= \mathbb{I}(t \geq t_{\text{corr}}) \frac{t - t_{\text{corr}}}{t_{\text{end}} - t_{\text{corr}}}, \\ \mathbf{P}_{\text{cal}}^{t} &= \mathbf{P}_{\text{raw}}^{t}(s_{\text{start}}) + \alpha(t)\mathbf{d}_{\text{err}}. \end{aligned} \tag{11}$$

For grasp-mode hand flow, we use $\delta = 5\,\text{cm}$.

SE(3) hand trajectory. After calibration, we convert the hand motion into a rigid $SE(3)$ trajectory. The translation is set to the calibrated 3D position of the designated contact fingertip. For orientation, we construct a palm frame from 3D hand landmarks: the palm normal is estimated by fitting a plane to the wrist and MCP joints, the in-plane axis is the wrist-to-middle-MCP direction projected onto this plane, and the third axis is their cross product. This yields a per-frame hand transform $\{T_{\text{hand}}^{t}\}_{t=1}^{T}$ used for hand-centric grounding.

Rejection criteria. Reliable hand trajectory estimation requires the hand to remain largely within the camera view. We project the estimated hand mesh into the image plane and reject any trajectory where more than 30% of the hand mask lies outside the image in any frame during the critical interaction phase. We also reject the trajectory if the primary calibration step fails, e.g., if no fingertip satisfies the projection constraints or the candidate scale set is empty. In either case, the system discards the current hand flow and prompts the video model for regeneration.

D.3 Hand-Guided Grasp Selection

NovaPlan also uses the generated video's human hand as a kinematic prior to select grasps. We implement a three-stage filtering and scoring pipeline to identify the optimal grasp pose.

We first generate $N = 5000$ candidate grasps using GraspGen, then filter the candidate grasps using the pipeline below:

- **Contact filtering.** Given contact fingers $\mathcal{F}_{contact}$ from the hand mesh calibration procedure, we calculate the distance d_{prox} between the grasp's "palm line" (segment connecting gripper finger bases) and these contact points. Grasps with $d_{prox} > 5\,\text{cm}$ are discarded to ensure semantic alignment with the hand-object interaction.
- **Collision avoidance ($S_{collision}$).** We verify collision-free poses against the scene point cloud (excluding the target object) with a threshold of 1 mm. Grasps that potentially collide with the scene point cloud are filtered.
- **Palm line support ($S_{support}$).** To verify that the gripper makes contact with the object, we compute the fraction of object points that are within $0.5\,\text{cm}$ of the gripper's opening. This metric is used as a soft filter for the final scoring of the remaining grasps.

We rank the unfiltered grasps using a composite score S_{total}:

$$S_{total} = S_{conf} \cdot S_{support} \tag{12}$$

where S_{conf} is the raw confidence from GraspGen and $S_{support}$ is the metric representing the fraction of the object surface that could be grasped with the predicted grasp direction.

Table 4: Example oracle subtask sequences provided to NovaFlow, DreamZero, and $\pi_{0.5}$.

Task	Oracle Sub-step Sequence
Four-layer Block Stacking	1.1 Grasp the blue block. 1.2 Put the blue block onto the red block. 2.1 Grasp the green block. 2.2 Put the green block onto the blue block. 3.1 Grasp the yellow block. 3.2 Put the yellow block onto the green block.
Color Sorting	1.1 Grasp the yellow block. 1.2 Put the yellow block into the yellow container. 2.1 Grasp the blue block. 2.2 Put the blue block into the blue container. 3.1 Grasp the red block. 3.2 Put the red block into the red saucer.
Hidden Object Search	1.1 Grasp the left drawer's black handle. 1.2 Pull the black drawer handle straight outward. 2.1 Grasp the right drawer's green handle. 2.2 Pull the green drawer handle straight outward. 3.1 Grasp the block in the right drawer. 3.2 Place the block onto the top surface of the left drawer.

D.4 Contact-Constrained Calibration for Non-Prehensile Recovery

For non-prehensile recovery, the recovery prompt specifies both the interaction mode and the contact finger, e.g., "poke the object with the index finger." However, generated recovery videos often distort the contact finger or shift the apparent contact point, making the standard fingertip-overlap calibration unreliable. We therefore modify the hand-flow calibration routine by explicitly enforcing contact between the prompted fingertip and the selected object surface.

Let f_{poke} denote the contact finger specified by the recovery prompt. Unlike grasp-mode calibration, we do not infer the contact finger from fingertip-mask overlap. Instead, we use the prompted contact finger directly. At contact onset t_{start}, we estimate a scale and translation correction $(s_{\mathrm{start}}, \Delta\mathbf{t}_{\mathrm{start}})$ such that the prompted fingertip aligns with the metric object surface:

$$\mathbf{p}_{\mathrm{obj}}^{t_{\mathrm{start}}} = s_{\mathrm{start}}\mathbf{p}_{\mathrm{tip}}^{t_{\mathrm{start}}} + \Delta\mathbf{t}_{\mathrm{start}}, \tag{13}$$

where $\mathbf{p}_{\mathrm{tip}}^{t_{\mathrm{start}}}$ is the reconstructed 3D location of the prompted fingertip and $\mathbf{p}_{\mathrm{obj}}^{t_{\mathrm{start}}}$ is the metric object surface point around the annotated contact region.

We apply this start-anchor correction to the full hand sequence:

$$\mathbf{P}_{\mathrm{hand,cal}}^{t} = s_{\mathrm{start}}\mathbf{P}_{\mathrm{hand,raw}}^{t} + \Delta\mathbf{t}_{\mathrm{start}}. \tag{14}$$

As in grasp-mode hand-flow calibration, we compensate residual projective drift near release using a localized ramp correction. For non-prehensile repair, the drift correction is computed with respect to the prompted contact finger f_{poke}, and we use a smaller threshold $\delta = 2\,\mathrm{cm}$ because the object displacement is typically small. This keeps the fingertip in metric contact with the selected object surface and prevents the recovery trajectory from sliding to a different contact mode.

The remaining steps, including palm-frame orientation estimation and conversion to an end-effector trajectory, are identical to the hand-flow grounding procedure in Section D.2.

E Baseline Setups and Oracle Definitions

For evaluation, both $\pi_{0.5}$ and DreamZero are evaluated zero-shot from their released checkpoints under the required DROID camera configuration.

We provide $\pi_{0.5}$, NovaFlow, and DreamZero with an "oracle planner": each receives a fixed sequence of ground-truth sub-goals in text form, so that their evaluation isolates visual grounding and execution from high-level planning. We present an example in Table 4. MOKA receives the task instruction and uses its own VLM-based planner to decompose it into subtasks, without oracle subtask sequences. Its evaluation therefore includes task planning as well as grounding and execution. VLP-style also produces its own plan and receives no oracle. It keeps our video generation and our training-free hybrid-flow execution, and replaces only the planner with VLP's beam search over language and video candidates scored by VLM-judged progress toward the final goal. Because this baseline replans over the remaining horizon whenever execution deviates from the predicted state, rather than repairing a single transition, it costs 2–3 times as many planning calls as NovaPlan, with the largest gap on the reactive search task.

F Visualizations of Long-horizon Experiment Rollouts

We provide a qualitative analysis of our video-based planning approach across all long-horizon manipulation tasks. The rollouts from Figure 11 to Figure 15 illustrate that NovaPlan can synthesize coherent multi-step plans and execute them with robustness and precision.

In the Four-layer Block Stacking task (Figure 11), the planner generates and follows a multi-stage vertical assembly plan, requiring consistent geometric reasoning over several sequential placements.

In the Color Sorting task (Figure 12), the planner uses semantic cues to match each object to its target container based on visual appearance. It also places the yellow block upright, allowing it to enter the yellow container without jamming at the opening.

In the Hidden Object Search task (Figure 13), the planner operates under partial observability and must first explore the scene (e.g., open multiple drawers if necessary) to reveal the target object before completing the placement.

In the assembly task from the Functional Manipulation Benchmark (Figure 14), NovaPlan is able to complete a long-horizon assembly sequence that requires millimeter-level alignment.

Finally, the FMB variant assembly task (Figure 15) showcases a last non-prehensile poking recovery step, highlighting the effectiveness of our hand-flow extraction strategy for proposing and executing corrective interactions.

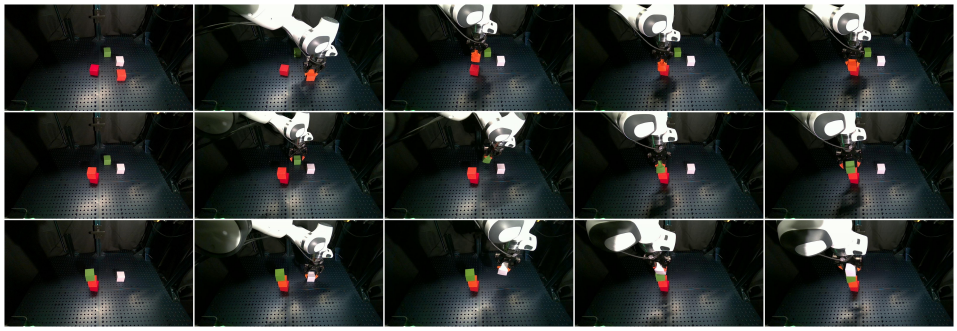


Figure 11: **Four-layer block stacking experiment.** Each row in this figure illustrates a distinct, sequential step in the block stacking task. The frames in each row display the execution of that specific step, ordered chronologically from left to right. This rollout demonstrates the model's ability to handle precise geometric constraints as the robot identifies, grasps, and aligns multiple blocks vertically.

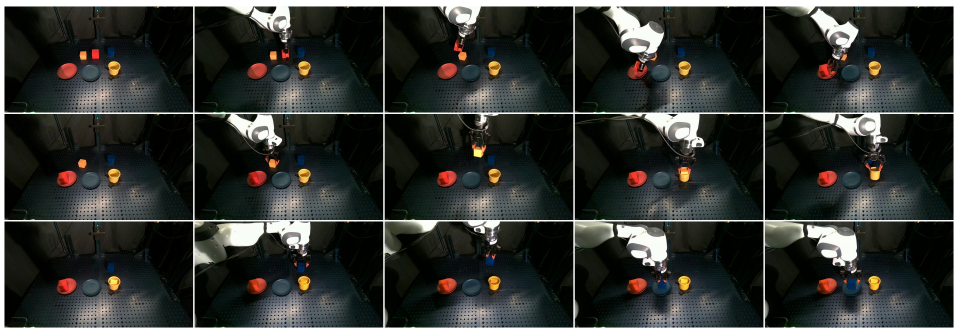


Figure 12: **Color sorting experiment.** Each row in this figure illustrates a distinct, sequential step in the color sorting task. The frames in each row display the execution of that specific step, ordered chronologically from left to right. This sequence highlights the semantic understanding of the planner as the robot sorts objects into containers with matching colors and the execution precision of preventing the yellow block from getting stuck at the yellow container.

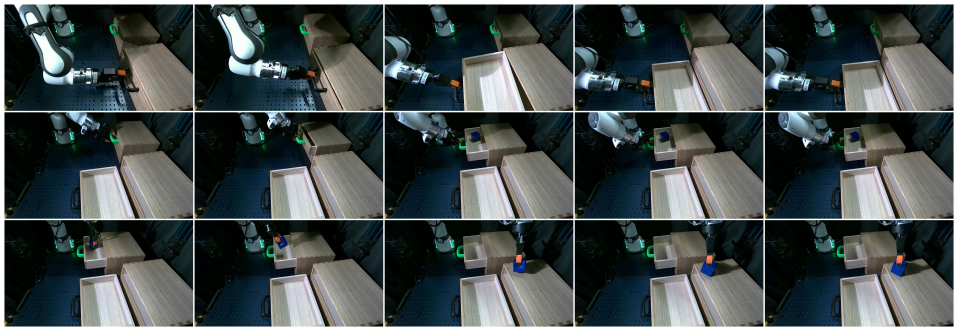


Figure 13: **Hidden object search experiment.** Each row in this figure illustrates a distinct, sequential step in the hidden object search task. The frames in each row display the execution of that specific step, ordered chronologically from left to right. This rollout shows an example of planning under partial observability where the robot must open drawers to reveal the target object.

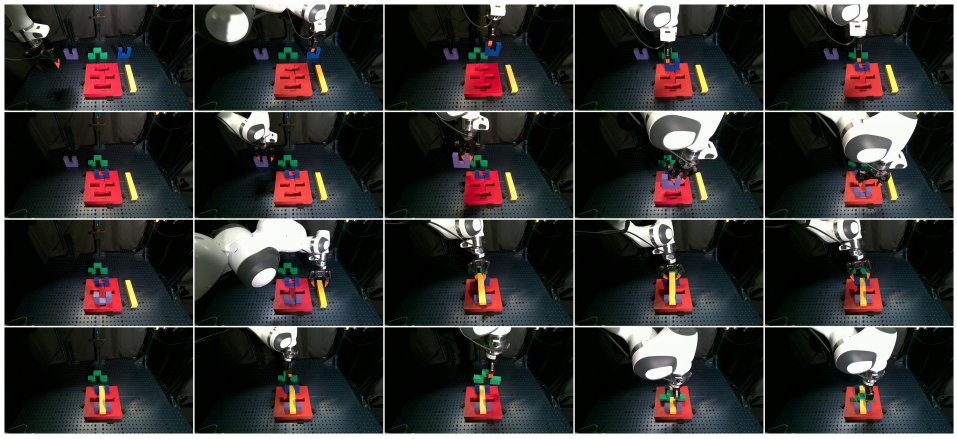


Figure 14: **Functional Manipulation Benchmark experiment.** Each row in this figure illustrates a distinct, sequential step in an assembly task from the Functional Manipulation Benchmark. The frames in each row display the execution of that specific step, ordered chronologically from left to right. This sequence shows that NovaPlan has the capability to solve long-horizon assembly tasks with millimeter precision.

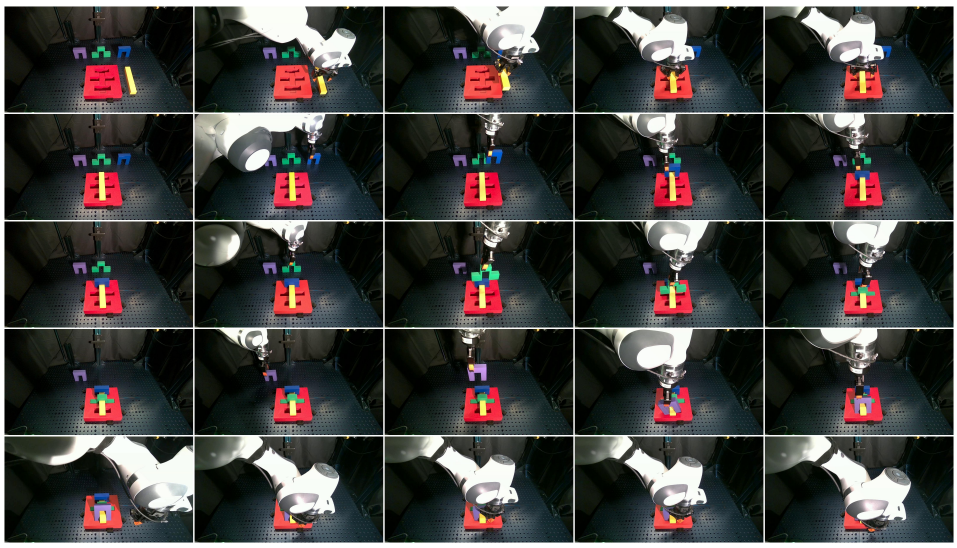


Figure 15: **Variant of Functional Manipulation Benchmark experiment.** Each row in this figure illustrates a distinct, sequential step in the variant of an assembly task from the Functional Manipulation Benchmark. The frames in each row display the execution of that specific step, ordered chronologically from left to right. The last row of this rollout demonstrates a non-prehensile poking recovery and shows the effectiveness of our hand flow extraction strategy.

G Example Plans and Action Prompts for Video Generation of Long-horizon Tasks

Here, in Table 5, we present representative multi-step plans produced by the planner. For each long-horizon task, we also provide the corresponding example prompts used to condition the video generation model, together with the object of interest tracked at each step.

Table 5: Examples of generated multi-step action plans, action prompts for step video generation, and objects to track.

Step	Proposed Action	Action Prompt for Video Generation	Object to Track
Task 1: 4-Layer Block Stacking			
Goal: Stack three blocks onto the other block.			
1	Reach to the orange block, grasp it from the top, lift it, move it to the red block, center it, and lower until seated.	A human hand reaches to the orange block sitting on the tabletop, grasps it firmly from the top, lifts it straight up, carries it to the red block on the tabletop on its right side, carefully centers it over the red block, and lowers it until fully seated and stable before releasing. The human hand should make every effort not to occlude the orange block during the action and keep the orange block fully visible. The hand completely leaves the frame after the action is done and the overall scene remains still. Do not introduce any new objects in the scene.	Orange block
2	Reach to the green block, grasp it from the top, lift it, move it above the orange block, align it, and lower to seat before releasing.	A human hand reaches to the green block sitting on the tabletop, grasps it firmly from the top, lifts it quickly, carries it directly above the orange block, puts it onto the orange block's top face, and releases it. The human hand should avoid occluding the green block during the action and keep the green block entirely in view. The hand fully exits the scene after the action is done and the scene stays motionless. Do not create any new objects in the scene.	Green block

Continued on next page

Table 5 – continued from previous page

Step	Proposed Action	Action Prompt for Video Generation	Object to Track
3	Reach to the white block, pinch-grasp from the side, lift, hover above the green block, touch the white block's lower edge to align, then release.	A human hand reaches to the white block sitting on the tabletop, grasps it from the side with a tight pinch, snap-lifts it, carries it quickly to hover just above the green block's top face, touches the white block's lower edge lightly to the green top to align, and releases it. The human hand should take care not to occlude the white block during the action and keep the white block fully in frame. The hand leaves the scene entirely after the action is done and the whole scene stays still. Do not create any new objects in the scene.	White block
Task 2: Color Sorting			
Goal: Put objects into containers of matching colors.			
1	Reach to the blue cube, grasp it from the top, lift it, move it above the center of the blue dish, lower it onto the dish, and release inside the boundary.	A human hand reaches to the blue cube on the tabletop, grasps it from the top, lifts it clear of the surface, moves it above the center of the blue dish, lowers it onto the dish surface, and releases after ensuring it is fully inside the dish boundary. The human hand should make every effort not to occlude the blue cube during the action and keep the blue cube clearly in view. The hand fully exits the scene after the action is done and the scene remains still.	Blue cube
2	Reach to the yellow cube, grasp it from the top, lift it, move it straight above the yellow cup opening, lower it into the cup, and release without contacting the rim.	A human hand reaches to the yellow cube on the left side of the tabletop in the scene, grasps it from the top, lifts it clear of the surface, moves it straight above the yellow cup opening, lowers it down into the cup until it is fully seated inside, and releases without contacting the rim. The human hand should avoid blocking the yellow cube during the action and keep the yellow cube fully in frame. The hand completely leaves the frame after the action is done and the whole scene stays still.	Yellow cube
3	Reach to the orange cube, grasp it from the top, lift it, move it over the center of the orange dish, lower until it contacts the bottom, and release fully inside.	A human hand reaches to the orange cube on the tabletop, grasps it firmly from the top, lifts it clear of the surface, moves it straight over the center of the orange dish, lowers it until it contacts the dish bottom, and releases with the cube fully inside the dish boundary. The human hand should take care not to occlude the orange cube during the action and keep the orange cube entirely visible. The hand fully exits the scene after the action is done and the whole scene stays still.	Orange cube
Task 3: Hidden Object Search			
Goal: Find the object hidden in one drawer unit and place it on top of the other unit.			
1	Reach to the right unit's black handle, grasp it firmly, and pull the drawer straight outward slowly until the interior is clearly visible.	A human hand reaches for the right unit's black handle, grips it firmly, and draws the drawer straight outward slowly until the inside is clearly visible. The hand should make every effort not to occlude the black handle and keep it fully visible throughout. After the drawer is open, the hand exits the frame completely and the scene remains static. Do not introduce any new objects in the scene.	Black handle
2	Reach to the left unit's green handle, grasp it securely, and pull the drawer straight out at a steady pace until the interior is clearly visible.	A human hand reaches to the left unit's green handle, grasps it firmly, holds it and pulls the drawer outward steadily until the interior is clearly visible. The hand should avoid blocking the green handle and keep it entirely in view during the motion. Once the drawer is opened, the hand fully exits the scene and everything stays still. Do not create any new objects in the scene.	Green handle

Continued on next page

Table 5 – continued from previous page

Step	Proposed Action	Action Prompt for Video Generation	Object to Track
3	Place the blue block in the left drawer onto the top surface of the right drawer and release it.	A human hand grasps the blue block from the top in the left drawer, lifts it up, and moves it onto the top surface of the right drawer and releases it. The human hand should take care not to occlude the blue block during the action and keep the blue block fully in frame. The hand fully exits the scene after the action is done and the whole scene stays still. Do not create any new objects in the scene.	Blue block
Task 4: Functional Manipulation Benchmark			
Goal: Insert all colored blocks on the table into their matching sockets in the red base.			
1	Reach to the blue U-shaped block, grasp it near the top edge, lift it, approach the right-side socket from behind the red base, align the inner notch, slide forward while lowering, then press down to seat flush.	A human hand picks up the blue U-shaped block from the back right side of the tabletop and places it into the furthest horizontal slot at the top of the red board. The open side of the U faces the center. The human hand then fully leaves the scene after the action is done and the whole scene stays still. The human hand should make every effort not to occlude the blue block during the action.	Blue block
2	Reach to the purple U-shaped block, grasp it near the top edge, lift it, align it with the lowest front socket, lower into the slot, and press down to seat flush.	A human hand picks up the purple U-shaped block from the back left side of the tabletop and inserts it into the lowest horizontal socket at the front of the red board. The open side of the U faces the center, aligning with the blue block. The human hand then fully exits the scene after the action is done and the scene remains still. The human hand should avoid occluding the purple block during the action.	Purple block
3	Reach to the yellow rectangular bar, grasp it at the midpoint, lift it, align it with the long central groove, lower it into the slot, and press until it sits fully seated.	A human hand picks up the long rectangular yellow block sitting on the right side of the red base on the tabletop and places it into the long central groove on the red base, seating it centered longitudinally within the vertical groove and aligning it with the U-shape openings of the blue and green blocks. The bottom edge of the yellow block extends past the purple block, terminating near the bottom edge of the recessed cross-shaped area of the red base. The human hand leaves the scene entirely after the action is done and the whole scene stays still. The human hand should take care not to occlude the yellow block during the action.	Yellow bar
4	Reach to the green block, grasp it near the top center, lift it, align it with the middle slot, lower and press down until it sits flush and locks over the yellow bar.	A human hand picks up the green block and inserts it into the middle horizontal slot on the red board. The green block sits in the middle horizontal slot and the arch of the green block straddles directly over the center of the yellow bar, locking it down into the red base. The human hand fully exits the scene after the action is done and the scene remains static. The human hand should make every effort not to occlude the green block during the action.	Green block

Continued on next page

Table 5 – continued from previous page

Step	Proposed Action	Action Prompt for Video Generation	Object to Track
Task 5: Functional Manipulation Benchmark-Variant			
Goal: Insert all colored blocks on the table into their matching sockets in the red base.			
1	Reach to the yellow rectangular bar, grasp it at the midpoint, lift it, align it with the long central groove, lower it into the slot, and press down until it sits fully flush.	A human hand holds the long rectangular yellow block and places it into the long central groove running vertically down the center of the red board. Ensure it is lying flat at the bottom of the pocket. The human hand exits the scene completely after the action is done and the whole scene stays still. The human hand should avoid occluding the yellow block during the action.	Yellow bar
2	Reach to the blue block, grasp it near the top edge, lift it, align it with the top horizontal slot, lower it to straddle over the yellow bar, and press down to seat fully.	A human hand picks up the blue block that is sitting on the right backside of the tabletop and inserts the blue block into the top horizontal slot at the top of the red base. It should straddle over the yellow bar and sit fully in the furthest horizontal slot. The human hand fully leaves the scene after the action is done and the whole scene stays still. The human hand should make every effort not to occlude the blue block during the action.	Blue block
3	Reach to the green block, grasp it near the top center, lift it, align it with the middle slot, lower to bridge over the yellow bar, and press down until fully seated.	A human hand picks up the green block that is sitting on the tabletop and inserts it into the middle horizontal slot. It should bridge over the yellow bar and sit fully in the middle horizontal slot. The human hand leaves the scene entirely after the action is done and the scene remains still. The human hand should take care not to occlude the green block during the action.	Green block
4	Reach to the purple U-shaped block, grasp it near the top edge, lift it, align it with the lowest remaining socket, lower it into place, and press down to seat flush.	A human hand picks up the purple U-shaped block that is sitting on the tabletop and inserts it into the last remaining empty horizontal socket at the lowest part of the red board. It should bridge over the yellow bar and sit upright in the lowest horizontal socket in the red base. The human hand fully exits the scene after the action is done and the whole scene stays still. The human hand should avoid occluding the purple block during the action.	Purple block

H Failure Attribution, Module Ablation, and Prompt Sensitivity

This appendix gives the full results of the analyses summarized in Section 4.4. All numbers use the three long-horizon tasks of Section 4.1 with ten trials per task.

H.1 Module-Level Failure Attribution

We record the intermediate outputs of each executed action. A failed action is assigned to the *first* stage whose output cannot be reliably consumed by the next stage. For example, if depth estimation returns a distorted scene but 3D tracking then produces an unusable trajectory, the failure is attributed to depth estimation rather than to tracking. Video generation is counted as a failure when the entire sampled set lacks a feasible rollout. The verifier filters individual infeasible rollouts before selection.

Table 1(b) reports the resulting distribution. The highest module failure rates are for 3D object tracking (20.0%) and monocular depth estimation (12.9%). Hand flow provides a fallback when the recovered object trajectory is unreliable. Video generation, object segmentation, hand segmentation,

and recovery reasoning have no attributed failures in these trials. Rollout selection and state verification each have a failure rate of 1.2%. Infeasible candidate videos are filtered and regenerated before execution.

Single-view generation is nonetheless a hard bottleneck for the irregular FMB assembly parts: only $17/64$ and $16/40$ individually generated videos are physically feasible in the variant and standard layouts respectively, which is why we use FMB to probe the boundary of the NovaPlan pipeline.

H.2 Controlled Component Ablation

Each ablation in Table 1(a) replaces exactly one module with a simpler alternative while holding every other module, prompt, and geometric threshold fixed. We evaluate pre-execution outputs by their correctness and execution-time mechanisms by full-task success.

Module-level output correctness applies to modules whose output can be judged directly against the archived trial data, independently of whether the surrounding trial eventually succeeds. *Rollout verification* measures the fraction of selections that pick a physically feasible rollout; the ablation removes the accumulated object flow image from each candidate tile and scores using frames alone. *State verification* measures the fraction of correct success/failure judgements; the ablation withholds the generated target frame I_t^{target}, leaving only the pre- and post-execution observations. *Hand calibration* measures the fraction of hand trajectories that reach metric contact with the object; the ablation uses raw HaMeR output without contact-based scale and drift correction. *Strategic beam search* measures the fraction of trials yielding a valid action ordering on the coupled tasks; the ablation substitutes purely reactive one-step proposal.

Long-horizon task success measures the effect of execution-time mechanisms on full-task completion. *Execution re-grounding* is ablated by reusing the rollout selected during planning instead of regenerating from the latest real observation. *Local recovery* is ablated by terminating the trial on the first verification failure. *Hybrid grounding* is ablated by disabling the switch to hand flow and executing object flow only.

H.3 Prompt Sensitivity

Because NovaPlan conditions its VLM stages on structured prompts, we test whether performance depends on exact phrasing. We test three prompt variants. The *paraphrase* preserves the original instructions with different wording. The *compressed* prompt shortens these instructions while retaining their structure. The *generic* prompt keeps the input and output schema but removes the physical and causal checks. Each variant is run three times over the archived examples from the same ten trials per long-horizon task, and we report mean correctness of each VLM stage in Table 6.

Table 6: Prompt sensitivity of each VLM stage (% correct, mean of three runs).

VLM role	Original	Paraphrase	Compressed	Generic
Action proposal	**100.0**	**100.0**	**100.0**	97.6
Rollout selection	98.8	**100.0**	97.6	78.8
State verification	**98.8**	97.6	95.3	89.4
Recovery reasoning	**100.0**	**100.0**	95.8	75.0

Paraphrasing and compression change correctness by at most 4.2% across the four VLM stages. Removing the physical and causal checks lowers rollout-selection correctness from 98.8% to 78.8% and recovery-reasoning correctness from 100.0% to 75.0%. These results indicate that the reasoning instructions contribute more to accuracy than their exact phrasing.