

# UC Merced

## Proceedings of the Annual Meeting of the Cognitive Science Society

### Title

Overhang Tower: Resource-Rational Adaptation in Sequential Physical Planning

### Permalink

<https://escholarship.org/uc/item/1rg9f26v>

### Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

### Authors

Shen, Ruihong

Li, Shiqian

Zhu, Yixin

### Publication Date

2026

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# Overhang Tower: Resource-Rational Adaptation in Sequential Physical Planning

Ruihong Shen<sup>1,2,3,4,5</sup>, Shiqian Li<sup>2,1,4,5</sup>, and Yixin Zhu<sup>1,2,4,5</sup>

<sup>1</sup> School of Psychological and Cognitive Sciences, Peking University <sup>2</sup> Institute for Artificial Intelligence, Peking University

<sup>3</sup> School of EECS, Peking University <sup>4</sup> State Key Laboratory of General Artificial Intelligence, Peking University

<sup>5</sup> Beijing Key Laboratory of Behavior and Mental Health, Peking University Project Website: <https://overhangtower.github.io>

## Abstract


Humans effortlessly navigate the physical world by predicting how objects behave under gravity and contact forces, yet how such judgments support sequential physical planning under resource constraints remains poorly understood. Research on intuitive physics debates whether prediction relies on the Intuitive Physics Engine (IPE) (Battaglia et al., 2013; Zhang et al., 2016) or fast, cue-based heuristics; separately, decision-making research debates deliberative lookahead vs. myopic strategies (Kahneman, 2011). These debates have proceeded in isolation, leaving the cognitive architecture of sequential physical planning underspecified. How physical prediction mechanisms and planning strategies jointly adapt under limited cognitive resources remains an open question. Here we show that humans exhibit a *dual transition* under resource pressure, simultaneously shifting both physical prediction mechanism and planning strategy to match cognitive budget. Using *Overhang Tower*, a construction task requiring participants to maximize horizontal overhang while maintaining stability, we find that IPE-based simulation dominates early stages while CNN-based visual heuristics prevail as complexity grows; concurrently, time pressure truncates deliberative lookahead, shifting planning toward shallower horizons—a dual transition unpredicted by prior single-mechanism accounts. These findings reveal a hierarchical, resource-rational architecture that flexibly trades computational cost against predictive fidelity. Our results unify two long-standing debates—simulation vs. heuristics and myopic vs. deliberative planning—as a dynamic repertoire reconfigured by cognitive budget.

**Keywords:** intuitive physics; sequential planning; mental simulation; heuristic inference; resource rationality

## Introduction

From infancy, humans develop a robust intuitive understanding of the physical world (McCloskey, 1983; Piloto et al., 2022; Spelke & Kinzler, 2007), enabling not only passive predictions—such as whether a stack of boxes will topple—but also active planning of complex arrangements, like building a cantilevered shelf. This capacity for *sequential physical planning*, the coordination of multiple interdependent actions to achieve a physical goal, is central to everyday cognition, yet its underlying mechanisms remain poorly understood.


Previous research on intuitive physics has primarily focused on *passive, single-step judgments*: observers view a scene and predict a future outcome without intervening (Battaglia et al., 2013; Bear et al., 2021; Hamrick et al., 2011; Zhang et al., 2016). While these accounts successfully explain how humans judge stability from static images (Calabro et al., 2025; Wang et al., 2024; Zhou et al., 2023), they leave a critical gap: how do physical prediction mechanisms guide *sequential physical planning*, where agents must coordinate multiple interdependent decisions over time? Some recent interactive benchmarks introduce agency into physical reasoning, covering single-



**Figure 1: Overview of *Overhang Tower* and computational models.** (a) In *Overhang Tower*, participants construct a tower by placing blocks from a given sequence to maximize horizontal overhang while maintaining continuous stability. Due to *resource-rational constraints* imposed by a large combinatorial search space, different planning strategies emerge. A myopic planner greedily seeks immediate overhang gains, often falling into local optima or causing collapse. In contrast, a deliberative planner looks ahead to build specific *physical structures*—such as counter-weighting and vertical anchoring (numbers denote placement order)—that may sacrifice early reward but unlock higher long-term overhang ( $R = 2.6$  vs.  $R = 1.7$ ). (b) The task inherently involves a *risk-reward trade-off*: extending blocks increases the overhang but shifts the Center of Mass (CoM), depleting the stability budget and making purely greedy approaches suboptimal. (c) To evaluate stability during planning, models rely on different *physical prediction* mechanisms. The IPE estimates stability via noisy probabilistic simulation of physical dynamics; alternatively, visual heuristics approximate stability through learned superficial visual patterns, bypassing explicit simulation.

shot interventions (Allen et al., 2020; Bakhtin et al., 2019) or multi-step interactions emphasizing action timing (Li et al., 2024, 2026). However, these tasks typically frame physical planning as a *binary satisfaction problem*—finding any successful causal chain to complete a goal—and none isolate the distinct computational challenges of *sequential planning under constraints*: (i) early actions restrict future possibilities, demanding deliberative lookahead; (ii) cumulative risk must be managed to balance safety and reward; (iii) computational resources are insufficient to exhaustively evaluate all candidate sequences (Binder et al., 2025; Huys et al., 2015; McCarthy et al., 2020).

To cope with such constraints, a natural hypothesis is that hu-



mans adaptively adjust the planning horizon, trading off speed against reward (Keramati et al., 2011; Kuperwajs et al., 2025; Snider et al., 2015): under time pressure, they rely on myopic, greedy action selection; when resources permit, they engage in deliberative lookahead, evaluating multi-step sequences before committing (Kahneman, 2011). While this dichotomy successfully explains behavior in abstract sequential decision tasks such as chess (Daw et al., 2011; Holding, 1989; Russek et al., 2025), applying it to the physical domain introduces a unique complication: the nature of the underlying physical prediction mechanism itself remains contested. Computational accounts have modeled physical prediction either as probabilistic mental simulation—formalized in the Intuitive Physics Engine (IPE) framework (Battaglia et al., 2013; Hamrick et al., 2011; Smith & Vul, 2013; Ullman et al., 2017)—or as fast, cue-based heuristics that bypass explicit physics (Calabro et al., 2025; Callaway et al., 2017; Sanborn et al., 2013; Tversky & Kahneman, 1974). Crucially, these two debates—myopic vs. deliberative planning strategies, and simulation vs. heuristics in physical prediction—have proceeded largely in isolation: planning research assumes a perfect oracle model of physics, while intuitive physics research focuses on passive judgments, neglecting how prediction mechanisms support sequential action.

This theoretical fragmentation leaves the cognitive architecture of sequential physical planning underspecified, raising two central questions: At the level of physical prediction, do

humans simulate physical dynamics or rely on learned visual heuristics to approximate stability? At the level of action selection, do humans employ fast, myopic strategies or engage in deliberative lookahead when constructing physical structures?

We address both questions through the lens of *resource-rationality* (Callaway et al., 2018, 2022; Griffiths et al., 2015). Recent work on passive judgment suggests simulation and heuristics constitute a dynamic repertoire, with costly simulation deployed only when cheaper heuristics prove insufficient (Li et al., 2025). Extending this logic to sequential planning, we hypothesize a *dual transition* under resource pressure: as construction progresses and scene complexity grows, participants should shift their physical prediction mechanism from costly IPE-style simulation toward visual heuristics; simultaneously, as time pressure increases, their planning strategy should transition from deliberative lookahead toward myopic responding. Together, these constitute a hierarchical resource-rational architecture in which both physical prediction mechanism and planning strategy adapt dynamically to match cognitive budget to task demands.

To test this hypothesis, we introduce *Overhang Tower*, an interactive construction task in which participants place blocks to maximize horizontal overhang while maintaining continuous stability, as shown in Fig. 1a. Unlike classic height-maximization stacking tasks (Groth et al., 2018), this objective enforces an explicit *risk-reward trade-off* (Fig. 1b): achieving high rewards requires placing blocks near stability limits,

rendering the planner acutely sensitive to physical prediction errors. Greedy strategies may over-commit to stable but suboptimal placements, whereas deeper lookahead can justify early scaffolding that unlocks subsequent high-reward opportunities—making *Overhang Tower* particularly diagnostic for dissociating physical prediction mechanism from planning strategy.

To systematically vary cognitive resources, we manipulated *time pressure*, testing participants under *time-constrained* and *unconstrained* conditions. To interpret the resulting behaviors, we developed a family of computational models that factorize physical prediction and planning independently. On the prediction dimension, we compare an IPE module running probabilistic forward simulations against a visual-heuristic predictor employing a learned classifier (Fig. 1c) (Calabro et al., 2025; Lindsay, 2021; Szegedy et al., 2017). On the planning dimension, we contrast myopic strategies with deliberative lookahead. By fitting resource parameters (*e.g.*, search depth, noise levels) and comparing predicted action sequences across conditions, we find that human physical reasoning reflects resource-rational adaptation at both dimensions: as task demands escalate, people deploy simpler physical prediction mechanisms and adopt shallower planning strategies, balancing cognitive cost against expected reward.

## Experimental Design

### Participants

A total of 82 participants (54% male, 46% female; mean age =  $21.5 \pm 2.6$ ) were recruited from an online campus forum. Participants received a base payment plus a performance-proportional bonus. One participant was excluded for failing to pass the practice phase.

### Stimuli

The experiment employs a web-based, interactive 2D construction task in which participants arrange a sequence of 6 blocks on a continuous grid of size  $8 \times 8$  (Fig. 2b). Blocks are randomly drawn from three distinct shapes defined by dimensions  $w \times h$  (width  $\times$  height), where  $w \in \{0.6, 1.2, 1.8\}$  and  $h = 0.6$ .

Physical dynamics are simulated via PyBullet (Coumans, 2015) at a temporal resolution of 1/240 seconds over a 5-second window per sequence. To approximate the visual input of real-world physical reasoning, scenes are rendered into photorealistic 2.5D representations using PyRender (Fig. 2a).

Each trial specifies a predetermined block sequence of length  $T = 6$ , denoted  $(b_1, \dots, b_T)$ , fully visible to participants. The environment is defined in the  $x$ - $z$  plane. After  $t$  placements, the tower geometry is represented as

$$G_t = ((b_1, \mathbf{p}_1), \dots, (b_t, \mathbf{p}_t)), \quad \mathbf{p}_i = (x_i, z_i), \quad (1)$$

where  $b_i$  indexes the block type and  $\mathbf{p}_i$  is its translational placement on the grid. The first block is pre-positioned at the ground center. The decision state at step  $t$  is defined as  $s_t := (G_t, b_{t+1:T})$ , augmenting the current geometry with the remaining blocks to be placed.

At each step  $t$ , the participant selects an action  $a_t = (x, z)$  from the legal action set  $\mathcal{A}(s_t)$ , where  $x$  specifies a continuous horizontal position and  $z$  a discrete height layer. The action set enforces two physical constraints: (i) *non-penetration*—no overlap with existing blocks or the ground; and (ii) *support*—the new block must rest on a supporting face of an existing block (Fig. 2c). Tower stability is given by a binary indicator  $\mathbb{I}_{\text{stable}}(G_t)$ , evaluated by the physics engine.

The objective is to maximize horizontal overhang—the maximum lateral distance of any block from the structural centerline (Fig. 2d):  $\Omega(G_T) = \max_{i \leq T} (|x_i| + \frac{1}{2}w_i)$ , where  $w_i$  is the width of the  $i$ -th placed block. The episode reward is contingent on continuous stability throughout construction: if the tower remains stable at every step, the participant receives the final overhang; otherwise, the reward is zero:  $R = \left( \prod_{k=1}^T \mathbb{I}_{\text{stable}}(G_k) \right) \cdot \Omega(G_T)$ .

### Task Characteristics

Unlike binary satisfaction tasks, the objective of *Overhang Tower* creates a continuous trade-off: extending blocks shifts the center of mass toward the edge, progressively depleting the “stability budget.” A greedy planning strategy that maximizes immediate overhang rapidly exhausts this budget, rendering myopic strategies suboptimal.

Optimal performance therefore requires non-greedy strategies that sacrifice immediate reward for long-term stability. For instance, *counter-weighting* shifts the center of mass inward, while *vertical anchoring* constructs sandwich-like interlocking structures to clamp extended blocks. Because these tactics demand precise causal ordering—blocks must be sequenced to maintain intermediate stability—they serve as a diagnostic signature of deliberative lookahead, making *Overhang Tower* well-suited for dissociating planning from physical prediction.

### Procedures

**Task interface** Participants used a 2D mouse interface to position a “ghost block” preview (Fig. 2) that provided geometric validity feedback (*e.g.*, collision or support validation) but no physical stability cues. We recorded continuous mouse trajectories and preview dwell times to reconstruct participants’ internal search processes.

**Familiarization** Prior to the main experiment, participants completed an interactive tutorial covering the web interface, physics dynamics, and scoring rules, followed by 3 practice trials excluded from the stimulus pool. To advance, participants were required to exceed a minimum performance threshold. Unlike the main task, the practice phase permitted trial-and-error repetition, and the optimal tower configuration was displayed after each trial to provide clear ground truth for the expected strategy.

**Experiment** Participants were randomly assigned to one of two between-subjects conditions: *time-constrained* or *unconstrained*. In the *time-constrained* condition, each placement was subject to a strict 5-second limit; failure to confirm a placement within this window terminated the trial and recorded a



reward of zero. In the *unconstrained* condition, participants were free to deliberate for an unlimited duration before each placement. Under both conditions, participants completed 20 experimental trials following familiarization, each presenting a distinct task sequence, for a total session duration of approximately 20 minutes. A comprehensive behavioral log—including action sequences, timestamps, and continuous mouse trajectories—was recorded for subsequent analysis.

## Models

To investigate the cognitive mechanisms underpinning sequential physical planning, we develop a computational framework that factorizes the planning process into two independent components: a physical prediction module that evaluates the stability of candidate configurations, and a planning module that expands the search space before committing to an action.

On the prediction dimension, we contrast an Intuitive Physics Engine that estimates stability via Monte Carlo simulation with a visual-heuristic model implemented as a Convolutional Neural Networks (CNN) (Calabro et al., 2025; Szegedy et al., 2017). On the planning dimension, we distinguish myopic action selection from deliberative lookahead via bounded forward search. We evaluate these paradigms based on both task performance and alignment with human behavioral data.

### Prediction Models

**IPE** This paradigm instantiates the physical prediction mechanism via Monte Carlo probabilistic simulation, representing physical stability reasoning as mental simulation under noise. For each candidate action, the model runs  $K$  forward simulations with stochastic perturbations to estimate stability, where  $\text{Stable}^{(i)}$  denotes whether the  $i$ -th run remains stable:

$$\hat{p}_{\text{IPE}}(\text{stable} \mid s_t, a) = \frac{1}{K} \sum_{i=1}^K \mathbb{I}[\text{Stable}^{(i)}(G_{t+1})]. \quad (2)$$

We apply Gaussian perturbations  $\mathcal{N}(0, \sigma^2)$  to the position, gravity, and friction of each block. Following Li et al. (2025) and Wang et al. (2024), we set  $K = 50$  and  $\sigma = 0.03$  to best align with human predictions. This approach is computationally expensive but captures uncertainty arising from perceptual and dynamic noise.

**Visual-heuristic model** This paradigm approximates physical prediction through perceptual heuristics rather than explicit simulation. Because CNN extracts statistical patterns without explicit physical laws—sharing visual feature preferences with human judgments (Calabro et al., 2025)—we train a network as a tractable proxy for cognitive models given the task’s complexity, mapping a rendered image of the post-action geometry to a stability probability:

$$\hat{p}_{\text{NN}}(\text{stable} \mid s_t, a) = f_{\theta}(I(G_{t+1})). \quad (3)$$

Training data consists of 200k diverse partial tower configurations sampled from the task space and excluded from the stimulus pool, with ground-truth stability labels provided by the physics engine. We employ Inception-V4 (Szegedy et al., 2017) as the model backbone, which has been shown to excel

at stability prediction by integrating multi-scale visual cues (Calabro et al., 2025). The model is trained for 15 epochs, reaching 97.5% accuracy on the validation set.

### Planning Algorithms

Both planning algorithms select actions by evaluating candidates according to their expected reward and stability, computed using the internal physical prediction module  $\mathcal{P}$ . The critical distinction lies in the temporal horizon of evaluation: myopic planning considers only immediate consequences, whereas deliberative lookahead evaluates multi-step futures before committing.

**Myopic planning** The myopic planner selects each action based solely on its immediate expected value, without anticipating how current choices constrain future possibilities. At decision state  $s_t$ , the planner evaluates each candidate action  $a \in \mathcal{A}(s_t)$  by:

$$a_t^* = \arg \max_{a \in \mathcal{A}(s_t)} \hat{p}_{\mathcal{P}}(s_t, a) \cdot r(s_t, a), \quad (4)$$

where  $r(s_t, a)$  is the immediate reward. This strategy is computationally cheap but systematically neglects downstream consequences.

**Deliberative lookahead** The lookahead planner imagines future trajectories before acting. At each decision state  $s_t$ , the planner expands candidate action sequences up to depth  $D$  and evaluates their utilities:

$$U(\pi \mid s_t) = \left( \prod_{k=t}^{t+D-1} \hat{p}_{\mathcal{P}}(s_k, a_k) \right) \cdot V(s_{t+D}), \quad (5)$$

where  $\pi = (a_t, \dots, a_{t+D-1})$  and  $V(s_{t+D})$  estimate the overhang of the resulting state (e.g., cumulative overhang at depth  $D$ ). Larger  $D$  enables deeper foresight at greater computational cost, sacrificing immediate reward for long-term gain—the behavioral signature of counter-weighting strategies.

## Results

The basic statistics between the two groups are shown in [Tab. 1](#). Results from 81 participants indicate that time constraints significantly reduced the overhang of successful solutions ( $p < .001$ ) without compromising the overall success rate ( $p = .484$ ). This pattern suggests that under time pressure, humans adopt a risk-averse planning strategy, sacrificing potential reward to ensure stability.

Table 1: **Overall performance.** Comparison between participants in the time-constrained ( $n = 43$ ) and unconstrained ( $n = 38$ ) conditions. Average overhang is computed over successful trials only; decision time indicates the mean decision time per placement. Values indicate mean  $\pm$  SEM unless otherwise stated. Statistical significance assessed using permutation tests.

| Metric                | Time Limit        |                   |            |
|-----------------------|-------------------|-------------------|------------|
|                       | 5s                | None              | $p$ -value |
| Total reward          | 18.25 $\pm$ 0.57  | 19.57 $\pm$ 0.56  | .107       |
| Stable proportion (%) | 70.2 $\pm$ 2.5    | 67.8 $\pm$ 2.1    | .482       |
| Average overhang      | 1.321 $\pm$ 0.023 | 1.460 $\pm$ 0.029 | < .001     |
| Decision time (s)     | 2.47 $\pm$ 0.07   | 7.27 $\pm$ 0.88   | < .001     |




Figure 3: **Planning trajectory distributions and  $\Gamma_{G_T}$ .** (a) Planning distributions for a block sequence with widths 0.6, 1.8, 1.2, 1.8, 0.6, 1.8. Flow color encodes stability (native hue = stable; crimson = unstable), opacity indicates normalized transition probability. The near-optimal solution yields overhang 2.4, whereas the best human result was 1.94. Time-constrained participants exhibited a stronger tendency toward vertical stacking; early differences remain modest because both groups first establish a conservative base, with divergence clearer later. (b) Left: distribution of  $\Gamma_{G_T}$  across individual participants under each condition. Right: distribution of  $\Gamma_{G_T}$  across tasks.

### Physical Prediction Mechanism Switching

To validate our hypothesis concerning physical prediction mechanisms, we compared human behavior against two candidate models—a visual-heuristic model and an IPE—with a noise-free veridical simulator serving as baseline.

Assuming stepwise stability judgments and rational action selection, we treat all executed actions as positive samples of perceived physical plausibility, and measure mean log-likelihood across all human actions (Fig. 4). Both approximate models significantly outperformed the veridical baseline ( $p < .001$ ), confirming that human intuitive physics systematically deviates from Newtonian mechanics—people rely on noise-corrupted representations rather than faithful physical simulation.

More critically, the relative explanatory power of the two physical prediction mechanisms varied across the planning trajectory. In early stages, when the scene was sparsely populated, IPE exhibited a modest advantage. As scene complexity grew with additional blocks, however, this advantage reversed: the visual-heuristic model decisively outperformed IPE ( $p < .001$ ).




Figure 4: **Physical prediction mechanism performance across structure complexity.** The left panel shows the log-likelihood of stability predicted by three models on human-rational states as construction progresses; larger values indicate a better fit to human judgments. The right panel shows the relative log-likelihood advantage of IPE vs. visual heuristics, computed as the residual with respect to the veridical simulator. As scene complexity grows over successive placements, the visual-heuristic model surpasses IPE.

This crossover reflects a fundamental asymmetry in how the two mechanisms scale with structural depth. In IPE, perceptual noise—modeled as Gaussian perturbations  $\mathcal{N}(0, \sigma^2)$ —propagates through forward simulation: positional uncertainty in lower blocks cascades into contact forces above, causing predictive variance to grow with stack height. Stability estimates thus become increasingly diffuse and uninformative for complex towers. If human mental simulation faces an analogous computational bottleneck, continuing to invest limited cognitive resources in a physical prediction mechanism whose reliability rapidly degrades would be resource-irrational.

The visual-heuristic mechanism plausibly sidesteps this complexity penalty by mapping static geometric features—such as vertical misalignment or unsupported overhang—directly onto stability estimates, without unrolling contact dynamics forward in time. Its representational cost therefore remains roughly constant regardless of stack depth. This asymmetry naturally motivates a hybrid physical prediction architecture: human planners invest cognitive effort in deliberate IPE-style simulation for simple structures where it is reliable, but shift adaptively to efficient visual heuristics when compounding uncertainty renders deep simulation prohibitive.

### Behavioral Evidence for Deliberative Lookahead

We first visualize the planning trajectories of participants in Fig. 3a. Participants under time pressure exhibited a stronger preference for vertical stacking—a conservative planning strategy that prioritizes immediate stability—whereas unconstrained participants explored more laterally extended configurations. This divergence suggests that time pressure suppresses multi-step lookahead. To test this hypothesis more rigorously, we developed a structural metric that operationalizes lookahead in terms of *order dependency*.

Our key insight is that deliberative planning strategies—counter-weighting and vertical anchoring—exploit physical dependencies that emerge only when blocks are placed in a precise causal sequence. A counterweight must be positioned *before* the cantilevered block it stabilizes; reversing this order causes collapse. Thus, if a tower’s final geometry can only be achieved through a small fraction of possible construction

Table 2: **Terminal reward across 20 tasks.** Averaged terminal reward for human participants and computational models. Stochastic models were run 40 times to reduce variance.  $D$  indicates lookahead search depth.

| Model                  | Terminal Reward  |
|------------------------|------------------|
| Human (w/ time limit)  | $0.913 \pm 0.03$ |
| Human (w/o time limit) | $0.979 \pm 0.03$ |
| Myopic                 | $0.52 \pm 0.12$  |
| Lookahead ( $D = 2$ )  | $0.912 \pm 0.10$ |
| Lookahead ( $D = 3$ )  | $1.180 \pm 0.24$ |

orders, this path dependency serves as a behavioral signature of deliberative lookahead. Formally, for a completed tower  $G_T$ , let  $\Pi(G_T)$  denote the set of all geometrically valid construction orders—permutations satisfying non-penetration and support constraints at each step. We define **Order Dependency** as:

$$\Gamma_{G_T} = 1 - \frac{|\{\pi \in \Pi(G_T) \mid \text{Stable}(\pi)\}|}{|\Pi(G_T)|}. \quad (6)$$

High  $\Gamma$  indicates that most feasible construction paths fail due to intermediate instability—the structure is path-dependent and requires deliberative sequencing.

As shown in Fig. 3b, time pressure significantly reduced Order Dependency (Unconstrained:  $\Gamma = 0.24 \pm 0.03$ ; Time-constrained:  $\Gamma = 0.15 \pm 0.02$ ;  $p = .022$ ). This difference reveals a qualitative shift in planning strategy: unconstrained participants constructed towers demanding precise sequencing to maintain stability, whereas time-constrained participants retreated to order-invariant heuristics—trading structural complexity for cognitive efficiency.

### Computational Evidence for Deliberative Lookahead

Building on prior results, we evaluate the joint architecture of sequential physical planning by coupling the stage-dependent hybrid IPE-heuristic physical prediction mechanism with varying depths of forward search. Because computing exact action-level likelihoods in our continuous, high-dimensional state space is computationally intractable, we focus on the aggregate structural achievements generated by these coupled models across 20 tasks. The results, alongside human performance, are summarized in Tab. 2.

The results reveal a clear dissociation between myopic and deliberative planning strategies. Guided by the same physical prediction mechanism, the myopic planner achieves a terminal reward drastically below human baselines. This performance gap arises because myopic planning greedily pursues immediate overhang, rapidly exhausting the stability budget—confirming that human physical reasoning cannot be reduced to reactive, single-step utility maximization.

Extending the planning horizon reliably recovers performance. The correspondence between human performance under varying time constraints and models of different search depths suggests that time pressure truncates, but does not eliminate, deliberative lookahead. Taken together, these results demonstrate that successful sequential physical planning requires a tightly coupled architecture: a physical prediction mechanism capable of estimating stability under uncertainty,

embedded within a planning strategy that extends deep enough to navigate the risk-reward trade-off.

## Discussion

**Coupling between prediction and planning** We primarily treated physical prediction mechanisms and planning strategies as independent and separable processes. However, resource-rational agents likely vary both dynamically (Callaway et al., 2018; Lieder & Griffiths, 2020), and their coupling may give rise to mixed or emergent mechanisms. Future work could examine how co-adaptation between these two capacities shapes behavior under graded resource constraints.

**Disentangling error sources** Human physical prediction is calibrated to noisy real-world perception and uncertain dynamics (Smith & Vul, 2013), whereas simulators implement idealized deterministic physics (Davis & Marcus, 2014). Consequently, some errors may reflect mismatches between human uncertainty and simulator assumptions, rather than planning limitations alone. An important future direction is to better dissociate errors from imperfect physical prediction and those from limited planning.

**Exploring diverse physical phenomena** Our study focuses on structural stability, a largely static physical domain. Whether the observed trade-offs between physical prediction mechanisms and planning strategies generalize to dynamic phenomena—such as collision or long-range causality (Wang et al., 2024)—where predictions must be propagated over longer temporal horizons and errors may accumulate over time, remains an open question.

**Impact of task exposure** Currently, our framework evaluates each task independently. However, humans actively adapt through trial-and-error. Future work could examine how cumulative task exposure shapes behavior. Beyond improving performance, environment familiarity and prior experience might systematically shift how participants weigh simulation against heuristics or adjust planning horizons across trials.

## Conclusion

We introduced *Overhang Tower*, a construction task in which we manipulated time pressure to test whether humans adaptively shift both prediction mechanisms and planning strategies. Our results support a resource-rational architecture: IPE-based simulation dominated early stages while visual heuristics prevailed as complexity grew; concurrently, time pressure truncated deliberative lookahead toward shallower horizons. These findings unify two previously isolated debates—simulation vs. heuristics and myopic vs. deliberative planning—as a dynamic repertoire reconfigured by cognitive budget.

By treating humans as active, capacity-limited agents, our framework offers a foundation for understanding physical problem-solving. Future work could extend these findings to dynamic environments and explore whether AI agents, which currently exhibit notable shortcomings in physical reasoning (Li et al., 2022), can leverage this flexible architecture to achieve human-like physical reasoning.



**Acknowledgement** The authors would like to thank Ruizhe Chen (Peking University), Bingrui Gong (Peking University), and Yuchen Lin (Peking University) for their helpful discussions. This work is supported in part by the Brain Science and Brain-like Intelligence Technology—National Science and Technology Major Project (2025ZD0219400), the National Natural Science Foundation of China (62376009), the PKU-BingJi Joint Laboratory for Artificial Intelligence, the Wuhan Major Scientific and Technological Special Program (2025060902020304), the Hubei Embodied Intelligence Foundation Model Research and Development Program, and the National Comprehensive Experimental Base for Governance of Intelligent Society, Wuhan East Lake High-Tech Development Zone.

## References

- Allen, K. R., Smith, K. A., & Tenenbaum, J. B. (2020). Rapid trial-and-error learning with simulation supports flexible tool use and physical reasoning. *Proceedings of the National Academy of Sciences (PNAS)*, 117(47), 29302–29310 (cit. on p. 1).
- Bakhtin, A., van der Maaten, L., Johnson, J., Gustafson, L., & Girshick, R. (2019). Phyre: A new benchmark for physical reasoning. *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)* (cit. on p. 1).
- Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences (PNAS)*, 110(45), 18327–18332 (cit. on pp. 1, 2).
- Bear, D. M., Wang, E., Mrowca, D., Binder, F. J., Tung, H.-Y. F., Pramod, R., Holdaway, C., Tao, S., Smith, K., Sun, F.-Y., et al. (2021). Physion: Evaluating physical prediction from vision in humans and machines. *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)* (cit. on p. 1).
- Binder, F. J., Mattar, M. G., J. Kirsh, D., & Fan, J. E. (2025). Humans select subgoals that balance immediate and future cognitive costs during physical assembly. *Cognitive Science*, 49(11), e70135 (cit. on p. 1).
- Calabro, R., Bhattacharyya, K., Bainbridge, W., & Leong, Y. C. (2025). Humans and convolutional neural networks prioritize similar visual features in intuitive physics judgments. *Proceedings of the Annual Meeting of the Cognitive Science Society (CogSci)* (cit. on pp. 1–4).
- Callaway, F., Hamrick, J. B., & Griffiths, T. L. (2017). Discovering simple heuristics from mental simulation. *Proceedings of the Annual Meeting of the Cognitive Science Society (CogSci)* (cit. on p. 2).
- Callaway, F., Lieder, F., Das, P., Gul, S., Krueger, P. M., & Griffiths, T. L. (2018). A resource-rational analysis of human planning. *Proceedings of the Annual Meeting of the Cognitive Science Society (CogSci)* (cit. on pp. 2, 6).
- Callaway, F., Van Opheusden, B., Gul, S., Das, P., Krueger, P. M., Griffiths, T. L., & Lieder, F. (2022). Rational use of cognitive resources in human planning. *Nature Human Behaviour*, 6(8), 1112–1125 (cit. on p. 2).
- Coumans, E. (2015). Bullet physics simulation. In *Acm SIGGRAPH courses* (p. 1). (Cit. on p. 3).
- Davis, E., & Marcus, G. (2014). The scope and limits of simulation in cognition. *arXiv preprint arXiv:1506.04956* (cit. on p. 6).
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans’ choices and striatal prediction errors. *Neuron*, 69(6), 1204–1215 (cit. on p. 2).
- Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, 7(2), 217–229 (cit. on p. 2).
- Groth, O., Fuchs, F. B., Posner, I., & Vedaldi, A. (2018). Shapestacks: Learning vision-based physical intuition for generalised object stacking. *Proceedings of European Conference on Computer Vision (ECCV)* (cit. on p. 2).
- Hamrick, J. B., Battaglia, P. W., & Tenenbaum, J. B. (2011). Internal physics models guide probabilistic judgments about object dynamics. *Proceedings of the Annual Meeting of the Cognitive Science Society (CogSci)* (cit. on pp. 1, 2).
- Holding, D. H. (1989). Counting backward during chess move choice. *Bulletin of the Psychonomic Society*, 27(5), 421–424 (cit. on p. 2).
- Huys, Q. J., Lally, N., Faulkner, P., Eshel, N., Seifritz, E., Gershman, S. J., Dayan, P., & Roiser, J. P. (2015). Interplay of approximate planning strategies. *Proceedings of the National Academy of Sciences (PNAS)*, 112(10), 3098–3103 (cit. on p. 1).
- Kahneman, D. (2011). *Thinking, fast and slow*. Allen Lane; Penguin Books, New York. (Cit. on pp. 1, 2).
- Keramati, M., Dezfouli, A., & Piray, P. (2011). Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLoS Computational Biology*, 7(5), e1002055 (cit. on p. 2).
- Kuperwajs, I., Russek, E. M., Mattar, M. G., Ma, W. J., & Griffiths, T. L. (2025). Looking deeper into the algorithms underlying human planning. *Trends in Cognitive Sciences* (cit. on p. 2).
- Li, S., Ma, Y., Yan, J., Dai, B., Peng, Y., Zhang, C., & Zhu, Y. (2025). A simulation-heuristics dual-process model for intuitive physics. *Proceedings of the Annual Meeting of the Cognitive Science Society (CogSci)* (cit. on pp. 2, 4).
- Li, S., Shen, R., Tao, Y., Zhang, C., & Zhu, Y. (2026). Neural force field: Few-shot learning of generalized physical reasoning. *Proceedings of International Conference on Learning Representations (ICLR)* (cit. on p. 1).
- Li, S., Wu, K., Zhang, C., & Zhu, Y. (2022). On the learning mechanisms in physical reasoning. *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)* (cit. on p. 6).
- Li, S., Wu, K., Zhang, C., & Zhu, Y. (2024). I-phyre: Interactive physical reasoning. *Proceedings of International Conference on Learning Representations (ICLR)* (cit. on p. 1).



- Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 43, e1 (cit. on p. 6).
- Lindsay, G. W. (2021). Convolutional neural networks as a model of the visual system: Past, present, and future. *Journal of Cognitive Neuroscience*, 33(10), 2017–2031 (cit. on p. 3).
- McCarthy, W. P., Kirsh, D., & Fan, J. E. (2020). Learning to build physical structures better over time. *Proceedings of the Annual Meeting of the Cognitive Science Society (CogSci)* (cit. on p. 1).
- McCloskey, M. (1983). Intuitive physics. *Scientific American*, 248(4), 122–131 (cit. on p. 1).
- Piloto, L. S., Weinstein, A., Battaglia, P., & Botvinick, M. (2022). Intuitive physics learning in a deep-learning model inspired by developmental psychology. *Nature Human Behaviour*, 6(9), 1257–1267 (cit. on p. 1).
- Russek, E. M., Acosta-Kane, D., van Opheusden, B., Mattar, M. G., & Griffiths, T. L. (2025). Time spent thinking in online chess reflects the value of computation. *Cognitive science*, 49(10), e70119 (cit. on p. 2).
- Sanborn, A. N., Mansinghka, V. K., & Griffiths, T. L. (2013). Reconciling intuitive physics and newtonian mechanics for colliding objects. *Psychological Review*, 120(2), 411 (cit. on p. 2).
- Smith, K. A., & Vul, E. (2013). Sources of uncertainty in intuitive physics. *Topics in Cognitive Science*, 5(1), 185–199 (cit. on pp. 2, 6).
- Snider, J., Lee, D., Poizner, H., & Gepshtein, S. (2015). Prospective optimization with limited resources. *PLoS computational biology*, 11(9), e1004501 (cit. on p. 2).
- Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, 10(1), 89–96 (cit. on p. 1).
- Szegedy, C., Ioffe, S., Vanhoucke, V., & Alemi, A. (2017). Inception-v4, inception-resnet and the impact of residual connections on learning. *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)* (cit. on pp. 3, 4).
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *Science*, 185(4157), 1124–1131 (cit. on p. 2).
- Ullman, T. D., Spelke, E., Battaglia, P., & Tenenbaum, J. B. (2017). Mind games: Game engines as an architecture for intuitive physics. *Trends in Cognitive Sciences*, 21(9), 649–665 (cit. on p. 2).
- Wang, H., Jedoui, K., Venkatesh, R., Binder, F. J., Tenenbaum, J., Fan, J. E., Yamins, D., & Smith, K. A. (2024). Probabilistic simulation supports generalizable intuitive physics. *Proceedings of the Annual Meeting of the Cognitive Science Society (CogSci)* (cit. on pp. 1, 4, 6).
- Zhang, R., Wu, J., Zhang, C., Freeman, W. T., & Tenenbaum, J. B. (2016). A comparative evaluation of approximate probabilistic simulation and deep neural networks as accounts of human physical scene understanding. *Proceedings of the Annual Meeting of the Cognitive Science Society (CogSci)* (cit. on p. 1).
- Zhou, L., Smith, K. A., Tenenbaum, J. B., & Gerstenberg, T. (2023). Mental jenga: A counterfactual simulation model of causal judgments about physical support. *Journal of Experimental Psychology: General*, 152(8), 2237 (cit. on p. 1).