

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 COSMOS-EVAL: TOWARDS EXPLAINABLE EVALUATION OF PHYSICS AND SEMANTICS IN TEXT-TO-VIDEO MOD- ELS

006 **Anonymous authors**

007 Paper under double-blind review

ABSTRACT

Recent text-to-video (T2V) models have achieved impressive visual fidelity, yet they remain prone to failures in two critical dimensions: adhering to prompt semantics and respecting physical commonsense. Existing benchmarks, including VIDEOPHY and VIDEOPHY-2, formalize these axes but provide only scalar scores, leaving model errors unexplained and hindering reliable evaluation. To address this, we present **Cosmos-Eval**, an explainable evaluation framework that jointly assesses semantic adherence and physical consistency. Cosmos-Eval produces fine-grained 5-point scores with *natural-language rationales*, leveraging the physically grounded ontology of Cosmos-Reason1 and an LLM-based rationale refinement pipeline. This enables precise identification of semantic mismatches and violations of physical laws, such as floating objects or momentum inconsistencies. Experiments on VIDEOPHY-2 show that Cosmos-Eval matches state-of-the-art auto-evaluators in score alignment (Pearson 0.46 vs. 0.43 for semantics; Q-Kappa 0.33 vs. 0.33 for physics) while also delivering state-of-the-art rationale quality (e.g., best BERTScore F1 and BLEU-4 on both SA and PC). Beyond this benchmark, our framework generalizes to other evaluation suites, establishing a unified paradigm for explainable physics-and-semantics reasoning in T2V evaluation and enabling safer, more reliable model development.

1 INTRODUCTION

Recent breakthroughs in text-to-video (T2V) generation—from diffusion-based models like Lumiere (Bar-Tal et al., 2024) and Stable Video Diffusion (Blattmann et al., 2023) to transformer-driven systems like VideoPoet (Kondratyuk et al., 2024)—have enabled realistic video synthesis. Yet today’s systems are still far from acting as “general-purpose physical world simulators” (Bansal et al., 2025a): clips may look sharp but objects float, collisions miss responses, or the scene fails to reflect what the prompt describes. Importantly, evaluation protocols tell us *that* a video is wrong but rarely *why*.

A growing body of work converges on two complementary axes for judging T2V. VIDEOPHY (Bansal et al., 2025a) formalizes *Semantic Adherence (SA)*—whether entities, actions, and relations requested by a caption are grounded in the video—and *Physical Commonsense (PC)*—whether the dynamics (stability, contact, collisions, causality) are plausible even without the caption. The follow-up VIDEOPHY-2 (Bansal et al., 2025b) expands to hundreds of real-world actions and releases VIDEOPHY-2-AUTOEVAL, an automatic evaluator that outputs five-point SA/PC scores strongly correlated with human judgments, as reported in their published experiments. However, these evaluators primarily return *numbers*; they do not surface concrete evidence behind a grade, which makes it hard to diagnose failure modes or trust the assessment.

At the same time, advances in physical reasoning and multimodal explainability suggest a way forward. NVIDIA’s **Cosmos-Reason1** (NVIDIA et al., 2025) organizes physical commonsense into a hierarchical ontology (e.g., conservation, object permanence, spatial/temporal relations) and demonstrates video-based reasoning. In parallel, explainable evaluation methods show that structured prompting, multi-step verification, and LLM-as-a-judge pipelines can improve specificity and reliability of textual feedback (Mou et al., 2025; Gu et al., 2024). What is missing is a unified evaluator that marries the *score fidelity* of VIDEOPHY-2-AUTOEVAL with *physically grounded rationales* that make scores interpretable and therefore actionable.




Figure 1: **Score-only vs. explainable evaluation.** Qualitative PC example: VIDEOPHY-2-AUTOEVAL outputs only a numeric score (e.g., PC= 2) without justification, while *Cosmos-Eval* augments the score with a concise, physics-grounded rationale (e.g., implausible acceleration and missing collision dynamics), improving diagnosability and trust.

Our solution: Cosmos-Eval. We introduce *Cosmos-Eval*, an explainable SA/PC evaluation framework that reports five-point scores *and* concise, evidence-based rationales for each test case by default. *Cosmos-Eval* builds on *Cosmos-Reason1* to reason about physics, and uses a reference-seeded, judge-verified controller to iteratively refine rationales into an evidence-grounded chain of thought, then distills this behavior into a lightweight model for deployment. As illustrated in Fig. 1, a score-only evaluator such as VIDEOPHY-2-AUTOEVAL might return “PC= 2” for a bowling clip; *Cosmos-Eval* produces the same score and adds a short rationale (e.g., implausible acceleration and missing collision response), enabling concrete, actionable diagnostics.

Core Contributions.

- *Explainable SA/PC paradigm.* Within the VIDEOPHY/VIDEOPHY-2 setting, we pair five-point SA/PC scores with detailed rationales that support auditing, ablations, and failure localization (e.g., SA: “caption mentions a red ball, but video shows a blue cube”; PC: “object floats mid-air, violating gravity”), addressing the interpretability gap of prior benchmarks.
- *Score alignment with state-of-the-art auto-evaluators.* On the official VIDEOPHY-2 test set, our scores match VIDEOPHY-2-AUTOEVAL (SA Pearson: 0.46 vs. 0.43; PC Q-Kappa: 0.33 vs. 0.33) while adding rationales, avoiding the accuracy–interpretability trade-off.
- *Physically grounded rationale quality.* Leveraging *Cosmos-Reason1*’s ontology and our Stage-2 controller, our rationales achieve state-of-the-art similarity to references for SA/PC (e.g., SA




Figure 2: **Pipeline overview (Stages 0–2; Stage 3 training).** **Stage 0** (frozen VideoPhy scorers) maps inputs to discrete labels s_{SA}, s_{PC} (Eqs. 1–2). **Stage 1** (reason generation) produces SA/PC reference rationales $r_{ref}^{sa}, r_{ref}^{pc}$ (Figs. 3, 4). **Stage 2** (reason-augmented CoT) uses a judge-verified controller to build evidence-grounded chains and final responses (Fig. 5). **Stage 3** (two-run SFT; training) first fine-tunes a *score head* to predict 5-point labels $\{1, \dots, 5\}$, then fine-tunes *rationale generation conditioned on the predicted score* with CoT-style prompting, so the system outputs calibrated scores and concise, reference-faithful explanations at test time.

BERTScore F1 52.44 / BLEU-4 26.70; PC BERTScore F1 54.50 / BLEU-4 27.86), outperforming generic VLMs (e.g., Qwen-2.5-VL on PC: 36.31 / 4.44).

- *Generalizable pipeline.* Our reference-seeded, judge-verified rationale workflow and two-run SFT are scorer- and dataset-agnostic. In this work we evaluate on VIDEOPHY-2; extending to additional suites (e.g., T2VPhysBench) is a promising direction for future validation.

2 METHOD

We present the pipeline in execution order: **Stage 0** (VideoPhy scorers \rightarrow discrete SA/PC scores), **Stage 1** (reason generation), **Stage 2** (reason-augmented CoT), and **Stage 3** (SFT on textualized scores and Stage-2 \langle think \rangle / \langle answer \rangle). Stages 0–2 are generative (no parameter updates); Stage 3 sets training objectives (Sec. 4). The stages form a causal flow—*scores as priors \rightarrow reference reason \rightarrow evidence-verified chain \rightarrow distilled model*. Removing any stage degrades this flow: omitting **Stage 0** weakens ultimate agreement with human judgments; **Stage 1** is necessary to provide a score-aligned anchor r_{ref}^{τ} ; omitting **Stage 2** removes evidence verification and reduces rationale reliability; omitting **Stage 3** forces deployment to run Stages 0–2 online (high latency, unstable consistency). Overall, Stages 0–3 instantiate an information-theoretic pipeline (IB at Stage 0; conditional MI at Stages 1–3). Fig. 2 provides the high-level view of Stages 0–2: we first compute s_{SA}, s_{PC} via Eqs. equation 1–equation 2 (Stage 0), then synthesize score-aligned reference reasons (Stage 1), and finally run an evidence-verified controller that yields an explicit CoT and the final judgment (Stage 2).

Task summary (SA/PC). Following Bansal et al. (2025b), we evaluate two axes: **SA**—given video v and caption c , check whether key entities/actions/relations in c are grounded in v ; and **PC**—given v only, judge whether the observed dynamics (stability, contact, collisions, causality) are physically plausible. Both use a 5-point integer scale $\{1, \dots, 5\}$ and are evaluated independently (high SA need not imply high PC). Evaluations are per input instance.

Notational conventions. We adopt compact notation for clarity. We index tasks by $\tau \in \{sa, pc\}$ with inputs $x^{sa} = (v, c)$ and $x^{pc} = v$. Frozen VideoPhy scorers output labels $s_{SA}, s_{PC} \in \{1, 2, 3, 4, 5\}$. A stand-alone reason is r ; evidence snippets e appear only in Stage 2 (CoT), not in Stage 1. Task prompts are \mathbf{P}^{τ} . In Stage 1 (SA) we query an ensemble $\{\mathcal{M}_m\}_{m=1}^M$ and aggregate with a consensus extractor \mathcal{J}_{sa} ; in Stage 1 (PC) a base generator \mathcal{M}_{base} (reused in Stage 3) samples multiple reasons and a VLM judge \mathcal{J}_{pc} selects one. For Stage 2, c_i denotes a control code from

162 strategy set \mathcal{C} (Sec. 2.3); the history is $\mathcal{H}^\tau = \{(e_i^\tau, r_i^\tau)\}$. Unless stated otherwise, \mathcal{M} denotes a
 163 generator LLM/VLM used only at inference time. The Stage 1 output that seeds Stage 2 is r_{ref}^τ (the
 164 “reference answer”). We use an attempt budget $N \in \mathbb{N}$ and an acceptance indicator $\text{pass}_i^\tau \in \{0, 1\}$.
 165 The verifier \mathcal{V}_τ is an LLM judge with a fixed prompt \mathbf{U}^τ returning PASS or FAIL.
 166

167 2.1 STAGE 0: DISCRETE SCORING VIA VIDEOPHY-2-AUTOEVAL

169 Given $x^{\text{sa}} = (v, c)$ and $x^{\text{pc}} = v$, frozen VIDEOPHY-2-AUTOEVAL scorers output discrete labels:
 170

$$s_{\text{SA}} = \text{Model}_{\text{SA}}(x^{\text{sa}}) \in \{1, 2, 3, 4, 5\}, \quad (1)$$

$$s_{\text{PC}} = \text{Model}_{\text{PC}}(x^{\text{pc}}) \in \{1, 2, 3, 4, 5\}. \quad (2)$$

172 These scores are reported as discrete labels and passed as conditioning inputs to Stage 1.
 173

174 2.2 STAGE 1: REFERENCE REASON GENERATION

176 *Goal.* From the task input and the Stage-0 score, produce a task-specific reference answer r_{ref}^τ to seed
 177 Stage 2.
 178

179 **SA (Fig. 3).** Given $x^{\text{sa}} = (v, c)$ and s_{SA} (Eq. equation 1), we query an ensemble of M VLMs
 180 $\{\mathcal{M}_m\}_{m=1}^M$. Each model generates exactly one reason, forming an M -sized pool:
 181

$$\mathcal{R}_{\text{pool}}^{\text{sa}} = \{r_0^{\text{sa}, m} = \mathcal{M}_m(\mathbf{P}^{\text{sa}}, x^{\text{sa}}, s_{\text{SA}}; \text{generate})\}_{m=1}^M. \quad (3)$$

182 A separate aggregator LLM extracts the common content across models to produce the reference
 183 answer:
 184

$$r_{\text{ref}}^{\text{sa}} = \mathcal{J}_{\text{sa}}(\mathcal{R}_{\text{pool}}^{\text{sa}}; x^{\text{sa}}, s_{\text{SA}}) \equiv \text{Cons}(\mathcal{R}_{\text{pool}}^{\text{sa}}), \quad (4)$$

186 where $\text{Cons}(\cdot)$ denotes consensus-style extraction (e.g., intersecting claims, majority agreements,
 187 consistent justifications).
 188

189 **PC (Fig. 4).** Given $x^{\text{pc}} = v$ and s_{PC} (Eq. equation 2), a *single* base VLM $\mathcal{M}_{\text{base}}$ (later used in
 190 Stage 3) samples K candidate reasons:
 191

$$\mathcal{R}_{\text{pool}}^{\text{pc}} = \{r_{0,k}^{\text{pc}} = \mathcal{M}_{\text{base}}(\mathbf{P}^{\text{pc}}, x^{\text{pc}}, s_{\text{PC}}; \text{sample})\}_{k=1}^K. \quad (5)$$

192 An LLM judge selects the most appropriate reason conditioned on the video and the score:
 193

$$r_{\text{ref}}^{\text{pc}} = \mathcal{J}_{\text{pc}}(\mathcal{R}_{\text{pool}}^{\text{pc}}; x^{\text{pc}}, s_{\text{PC}}). \quad (6)$$

195 This is a *selection* step that reduces the K -candidate pool to a single reason—analogous to SA’s
 196 reduction step (consensus vs. best-candidate).
 197

Output. Stage 1 returns the task-specific reference answer $r_{\text{ref}}^\tau \in \{r_{\text{ref}}^{\text{sa}}, r_{\text{ref}}^{\text{pc}}\}$, which seeds Stage 2.
 198

199 2.3 STAGE 2: REFERENCE-SEEDED, JUDGE-VERIFIED CONTROLLER (REASON-AUGMENTED 200 COT)

202 Motivated by controller-based approaches to complex reasoning (e.g., HuatuoGPT-o1 (Chen et al.,
 203 2025a)), we instantiate a *Reference-Seeded, Judge-Verified Controller* that seeds with the Stage-1
 204 reference but *does not expose* that reference during search, explores/verifies/corrects with explicit
 205 strategies, and finally applies a label-rethink fallback (Fig. 5). Starting from the reference r_{ref}^τ
 206 (Eqs. equation 4, equation 6), we introduce evidence snippets and build a multi-step CoT under
 207 explicit control. Let the history be $\mathcal{H}_{i-1}^\tau = \{(e_j^\tau, r_j^\tau)\}_{j=0}^{i-1}$ and define the strategy set
 208

$$\mathcal{C} = \{\text{Backtracking}, \text{ExploringNewPaths}, \text{Verification}, \text{Correction}\}. \quad (7)$$

209 **Seed with reference and judge check.** We generate a seed *conditioning on the reference* and
 210 ask the LLM judge to decide PASS/FAIL, where $\mathbf{P}_{\text{seed-ref}}^\tau$, \mathbf{P}_c^τ , $\mathbf{P}_{\text{rethink}}^\tau$ are task-specific generation
 211 prompts (for seeding with the reference, for each strategy $c \in \mathcal{C}$ *without* the reference, and for
 212 the final fallback, respectively), and \mathbf{U}^τ is a unified verification prompt used at all checks (SA/PC
 213 templates in Appx. J):
 214

$$(e_0^\tau, r_0^\tau) = \mathcal{M}(\mathbf{P}_{\text{seed-ref}}^\tau, x^\tau, r_{\text{ref}}^\tau; \text{Reason}), \quad (8)$$

$$\text{pass}_0^\tau = \mathcal{V}_\tau(r_0^\tau, r_{\text{ref}}^\tau; \mathbf{U}^\tau) \in \{0, 1\}. \quad (9)$$

216 **Iterative controller without the reference (no replacement).** Let $T = \min(N, |\mathcal{C}|)$. For $i =$
 217 $1, \dots, T$, we sample *without replacement*

$$219 \quad c_i \sim \text{Unif}\left(\mathcal{C} \setminus \{c_1, \dots, c_{i-1}\}\right), \quad (10)$$

221 generate a new pair *without* r_{ref}^τ , and verify against the reference:

$$222 \quad (e_i^\tau, r_i^\tau) = \mathcal{M}(\mathbf{P}_{c_i}^\tau, x^\tau, \mathcal{H}_{i-1}^\tau; c_i), \quad (11)$$

$$224 \quad \text{pass}_i^\tau = \mathcal{V}_\tau(r_i^\tau, r_{\text{ref}}^\tau; \mathbf{U}^\tau) \in \{0, 1\}. \quad (12)$$

225 We stop early when $\text{pass}_i^\tau = 1$; if none passes after N attempts, we trigger LabelRethink.

227 **Label rethink fallback (with the reference).** If no iteration passes, we trigger a final
 228 LabelRethink that *re-injects* the reference and the full history:

$$230 \quad (e_{N+1}^\tau, r_{N+1}^\tau) = \mathcal{M}(\mathbf{P}_{\text{rethink}}^\tau, x^\tau, r_{\text{ref}}^\tau, \mathcal{H}_N^\tau; \text{LabelRethink}), \quad (13)$$

$$231 \quad \text{pass}_{N+1}^\tau = \mathcal{V}_\tau(r_{N+1}^\tau, r_{\text{ref}}^\tau; \mathbf{U}^\tau) \in \{0, 1\}. \quad (14)$$

232 If the final check fails, we discard the sample.

234 **Final chain and answer.** For a successful case (either early pass or rethink pass), we do *two-step*
 235 post-processing instead of one-shot formatting. First, we consolidate the accepted history into a
 236 single reasoning chain \hat{e}^τ by aggregating prior traces. Then, conditioned on \hat{e}^τ and the reference r_{ref}^τ ,
 237 we produce a reference-aligned and reformatted answer \hat{r}^τ . Formally,

$$239 \quad \hat{e}^\tau = \text{PostChain}\left(\{(e_j^\tau, r_j^\tau)\}_{j=0}^{i^*}; \text{SynthesizeChain}\right), \quad (15)$$

$$241 \quad \hat{r}^\tau = \text{PostAnswer}(\hat{e}^\tau, r_{\text{ref}}^\tau; \text{Reformat}). \quad (16)$$

242 Here i^* is the index of the accepted iteration (or $N+1$ for the rethink pass). Although our prompts
 243 here instantiate the SA task, the same two-step template applies to PC tasks as well; we keep using τ
 244 to denote the task. The complete controller is summarized in Algorithm 1.

- 246 • **Backtracking** (c=Backtracking). Roll back to the latest accepted step (or the seed) and
 247 produce a *minimal-edit* variant: keep the score prior fixed, alter one binding (entity/action/temporal
 248 cue), and reuse verified evidence where possible. Intended to fix a localized flaw without drifting.
- 249 • **Exploring New Paths** (c=ExploringNewPaths). Branch to an *alternative hypothesis*: propose
 250 different entity grounding, action interpretation, or temporal segmentation, allowing higher diversity.
 251 The goal is to escape a bad local choice while still honoring the score prior.
- 252 • **Verification** (c=Verification). Turn the current rationale into an explicit checklist of claims
 253 and probe the video for each to confirm or refute them; attach concrete, checkable details. Acts as
 254 a critic to expose hallucinations, temporal mistakes, or missing evidence.
- 255 • **Correction** (c=Correction). Rewrite the rationale *conditioned on verifier feedback*: remove
 256 contradictions, add concrete visual evidence, and enforce score-alignment gates (for SA/PC).
 257 Produces a compact, reference-blind fix suitable for final judging.

259 **Why show the reference only at the seed and in the fallback?** Seeding with r_{ref}^τ anchors the run
 260 near the Stage-1 consensus and stabilizes initialization. Hiding the reference during strategy iterations
 261 prevents confirmation shortcuts and label leakage, compelling the model to collect *independent*
 262 evidence. Re-introducing r_{ref}^τ at LabelRethink reconciles divergent trajectories without biasing
 263 intermediate exploration in a controlled, empirically verifiable manner.

265 **Relation to HuatuoGPT-01.** HuatuoGPT-01 (Chen et al., 2025a) targets verifiable medical QA
 266 with a ground-truth answer and a truth-equivalence verifier. Our Stage 2 addresses SA/PC evaluation
 267 where answers are not single-valued: we seed the controller with the Stage 1 reference rationale r_{ref}^τ ,
 268 hide this reference during strategy iterations (re-inject only at LabelRethink), and use a unified
 269 judge to enforce task definitions (SA consistency / physical commonsense) and calibration to the
 5-point scale; the output is an evidence–rationale pair rather than a single accepted answer.

270 2.4 STAGE 3: SFT WITH TEXTUALIZED SCORES AND $\langle\text{think}\rangle/\langle\text{answer}\rangle$
271

272 We adopt a *two-run* fine-tuning scheme that mirrors our experiments: first calibrate discrete scores,
273 then condition rationale generation on those scores. Stage 0 provides a 5-point label $s_\tau \in \{1, \dots, 5\}$,
274 which we textualize as $t^\tau \in \{1, 2, 3, 4, 5\}$. Stage 2 yields final outputs $(\hat{e}^\tau, \hat{r}^\tau)$ (the consolidated
275 chain and the final answer), serialized as

$$276 \text{pack_TA}(\hat{e}^\tau, \hat{r}^\tau) = \langle\text{think}\rangle \hat{e}^\tau \langle/\text{think}\rangle \langle\text{answer}\rangle \hat{r}^\tau \langle/\text{answer}\rangle. \quad (17)$$

277 **Training.** *Run A (score-only).* Given input x^τ (SA: $x^{\text{sa}}=(v, c)$; PC: $x^{\text{pc}}=v$), we perform teacher-
278 forced next-token prediction to generate t^τ (no supervision on any reasoning tokens) in this stage.
279 *Run B (final $\langle\text{think}\rangle/\langle\text{answer}\rangle$ conditioned on the score).* Starting from Run-A, we prepend
280 t^τ as an input condition and supervise only the packed target $Y = \text{pack_TA}(\hat{e}^\tau, \hat{r}^\tau)$; intermediate
281 scratch beyond \hat{e}^τ is not supervised. SA and PC are trained separately (PC omits c). At inference,
282 we read the $\langle\text{answer}\rangle$ field as the model’s output at test time. *Losses.* Both $\mathcal{L}_{\text{score}}^\tau$ and $\mathcal{L}_{\text{final}}^\tau$ are
283 standard token-level cross-entropy under teacher forcing: $\mathcal{L}_{\text{score}}^\tau = -\sum_{t \in \text{tok}(t^\tau)} \log p_\theta(y_t | y_{<t}, x^\tau)$,
284 $\mathcal{L}_{\text{final}}^\tau = -\sum_{t \in \text{tok}(Y)} \log p_\theta(y_t | y_{<t}, x^\tau, t^\tau)$.

285 **Parameter update.**
286

$$287 \theta_A = \arg \min_{\theta} \mathcal{L}_{\text{score}}^\tau \implies \theta_* = \arg \min_{\theta} \mathcal{L}_{\text{final}}^\tau \text{ initialized at } \theta_A. \quad (18)$$

290 3 EXPERIMENTS
291

292 We evaluate our pipeline on our curated *Cosmos-Eval-Set* (Sec. 3.1) on two tasks—Semantic Ad-
293 herence (SA) and Physical Commonsense (PC). We report (i) core agreement with 5-point labels
294 (Pearson, accuracy, weighted/quadratic Cohen’s κ , Spearman) and (ii) reasoning quality of rationales
295 (BERTScore P/R/F₁, BLEU-1/2/3/4, ROUGE-1/2).
296

297 3.1 EXPERIMENTAL SETUP
298

299 **Cosmos-Eval-Set: datasets and protocol.** We use two corpora: *VideoPhy* (Bansal et al., 2025a)
300 and *VideoPhy-2* (Bansal et al., 2025b). Training data is the union of **VideoPhy** (train+test) and
301 **VideoPhy-2** (train); evaluation is on the **VideoPhy-2 test set**. *VideoPhy-2* provides 5-point labels
302 for SA/PC; *VideoPhy* does not contain 5-point labels, so we *score its clips* using the released
303 *VIDEOPHY-2-AUTOEVAL* to obtain labels on the same 5-point scale. Both corpora contain
304 synthetic, model-generated videos and do not provide human-written rationales. We therefore run
305 Stages 1–2 to generate rationales and Stage 3 for SFT as in Sec. 2. Task inputs follow Sec. 2: SA
306 uses (v, c) while PC uses v only.
307

308 **Metrics and baselines.** We evaluate two groups of metrics: (A) *core agreement* to human 5-point
309 scores—Pearson’s r , Acc (exact match on $\{1, \dots, 5\}$), W-Kappa (linearly weighted Cohen’s κ),
310 Q-Kappa (quadratically weighted), and Spearman (rank correlation)¹—and (B) *reasoning quality*
311 on the *final* rationale text—SentSim (cosine over a sentence encoder; Appx. B), BERTScore (B-
312 P/B-R/B-F1), BLEU- n (B1–B4), and ROUGE (R1/R2), reported as % in Table 2. We compare
313 *VIDEOPHY-2-AUTOEVAL* (frozen scorer), Qwen-2.5-VL-7B (Bai et al., 2025), VideOLLaMA3-
314 7B (Zhang et al., 2025), InternVL3-8B/9B/14B (Zhu et al., 2025), and our **Cosmos-Reason1** (no
315 SFT) and **Cosmos-Eval** (Stage 3 two-run SFT: score-only $\rightarrow \langle\text{think}\rangle/\langle\text{answer}\rangle$ conditioned
316 on score; Sec. 2.4). Evaluations use identical inference budgets and prompts.
317

318 **Implementation details.** Stage 1 uses an ensemble size $M=2$ for SA (Eq. equation 3) and $K=5$
319 samples for PC (Eq. equation 5). Stage 2 runs the controller with budget $N=3$ and *strategy sampling without replacement* (Sec. 2.3); acceptance is decided by a unified LLM judge with a fixed
320 pass/fail prompt (Appx. J). Stage 3 follows the two-run schedule with *parameter updates given in*
321 Eq. equation 18; the supervision target is the packed $\langle\text{think}\rangle/\langle\text{answer}\rangle$ string in Eq. equation 17
322 (conditioned on the textualized score). Unless otherwise stated, we use identical video decoding and
323 frame sampling across all models; full hyperparameters appear in Appx. B.
324

¹For κ , we use quadratic weights for Q- κ and linear weights for W- κ ; higher is better for all core metrics.

324 **Table 1: Cross-dataset core SA/PC metrics** (\uparrow better). **SA**: caption–video semantic alignment; **PC**:
 325 video-only physical commonsense. Per sample, each method outputs a *discrete* score $s_\tau \in \{1, \dots, 5\}$,
 326 compared with human labels $y \in \{1, \dots, 5\}$ on the official SA/PC test splits. Metrics: *Pearson/Spearman*
 327 correlations of raw integers; *Acc* exact 5-class accuracy; *W- κ /Q- κ* linearly/quadratically
 328 weighted Cohen’s κ on the same 5-class scale. VIDEOPHY-2-AUTOEVAL is the dataset VLM-as-
 329 judge baseline; other rows are model predictions. **Bold** = best; underline = second-best.

Model	SA					PC				
	Pearson	Acc	W- κ	Q- κ	Spearman	Pearson	Acc	W- κ	Q- κ	Spearman
VIDEOPHY-2-AUTOEVAL	0.4327	0.3826	0.2696	0.4062	<u>0.4268</u>	0.3646	<u>0.3871</u>	<u>0.2144</u>	<u>0.3276</u>	0.3608
Qwen-2.5-VL-7B	0.3808	0.3417	0.2419	0.3779	0.3716	0.0840	0.3255	0.0490	0.0780	0.0900
VideoLLaMA3-7B	0.2769	0.2811	0.1536	0.2387	0.2574	0.0640	0.2699	0.0301	0.0500	0.0749
InternVL-8B	0.4143	0.3205	<u>0.2437</u>	<u>0.3855</u>	0.4196	0.1665	0.3064	0.0790	0.1363	0.1728
InternVL-9B	0.3827	0.2837	0.1902	0.2963	0.3747	0.1304	0.2717	0.0565	0.1044	0.1171
InternVL-14B	0.3420	0.3229	0.1643	0.2544	0.3402	0.1956	0.3464	0.0888	0.1424	0.1888
Cosmos-Reason1	0.3662	0.2821	0.2297	0.3260	0.3519	0.2356	0.3079	0.1479	0.2326	0.2166
Cosmos-Eval	0.4643	<u>0.3765</u>	0.2256	0.3507	0.4598	<u>0.3641</u>	0.3912	0.2207	<u>0.3301</u>	<u>0.3580</u>

340 **Table 2: Reasoning quality on SA/PC** on the same test splits as Table 1. Each model outputs one
 341 rationale per sample. Scores are % (metrics computed per-sample then averaged). References are the
 342 fixed per-video outputs of our Stage-2 controller and are shared across models at test time. **Bold** =
 343 best; underline = second-best.

Model	SA (Semantic Alignment)												PC (Physical Commonsense)											
	SentSim	B-P	B-R	B-F1	B1	B2	B3	B4	R1	R2	SentSim	B-P	B-R	B-F1	B1	B2	B3	B4	R1	R2				
Qwen-2.5-VL-7B	75.62	40.10	37.03	38.70	45.47	26.90	14.24	8.03	51.45	18.92	68.81	37.68	34.66	36.31	40.44	21.44	9.27	4.44	45.50	13.84				
VideoLLaMA3-7B	75.49	37.26	35.78	36.64	42.31	24.69	12.97	7.43	48.87	17.33	70.81	36.50	33.94	35.36	38.28	20.23	8.89	4.09	44.48	13.13				
InternVL-8B	72.49	41.27	35.20	38.30	39.69	21.30	9.84	4.54	46.06	13.32	72.49	<u>41.27</u>	35.20	38.30	39.69	21.30	<u>9.84</u>	4.54	46.06	14.32				
InternVL-9B	76.87	<u>43.44</u>	38.60	<u>41.12</u>	<u>46.76</u>	28.11	14.18	<u>8.52</u>	53.45	20.38	67.75	<u>40.68</u>	34.84	37.86	40.42	<u>21.83</u>	9.60	<u>4.60</u>	46.28	<u>14.83</u>				
InternVL-14B	<u>78.70</u>	40.36	40.35	40.49	46.73	<u>28.51</u>	15.24	8.90	<u>53.80</u>	<u>21.01</u>	72.36	39.23	37.93	38.72	<u>40.50</u>	21.46	9.05	4.35	<u>46.57</u>	14.17				
Cosmos-Reason1	77.30	22.94	40.98	31.52	24.84	14.48	7.75	4.26	41.66	14.43	70.05	18.94	<u>39.16</u>	28.52	18.46	9.41	4.30	2.13	33.88	8.95				
Cosmos-Eval	86.28	<u>53.55</u>	<u>51.15</u>	<u>52.44</u>	<u>56.72</u>	<u>42.85</u>	<u>33.38</u>	<u>26.70</u>	<u>61.12</u>	<u>34.74</u>	<u>80.90</u>	<u>54.81</u>	<u>53.99</u>	<u>54.50</u>	<u>55.38</u>	<u>41.45</u>	<u>33.31</u>	<u>27.86</u>	<u>59.72</u>	<u>33.34</u>				

3.2 MAIN RESULTS ON SA/PC (CORE AGREEMENT)

354 Table 1 summarizes cross-dataset core metrics. On **SA**, **Cosmos-Eval** attains best *Pearson* (0.4643)
 355 and *Spearman* (0.4598), and ranks *second* in *accuracy* (0.3765), while VIDEOPHY-2-AUTOEVAL
 356 remains stronger on κ measures. On **PC**, **Cosmos-Eval** leads in *accuracy* (0.3912), *weighted* κ
 357 (0.2207), and *quadratic* κ (0.3301), and is near the top on *Pearson/Spearman* (slightly below the
 358 frozen scorer). This suggests the two-run SFT preserves global calibration (correlations) while
 359 improving discrete decision agreement on PC.

360 **Takeaways.** (i) On SA, **Cosmos-Eval** improves rank-based correlations (Pearson/Spearman) over
 361 strong frozen scorers while remaining competitive in accuracy; (ii) on PC, it achieves the best discrete
 362 agreement (Acc, κ) and near-top correlations; (iii) unlike frozen scorers, our method produces
 363 *explanatory* outputs (<think>/<answer>).

3.3 REASONING QUALITY (STAGE-2 & FINAL OUTPUTS)

368 We evaluate final rationales with BERTScore, BLEU, and ROUGE on our held-out evaluation set
 369 (Table 2). **Cosmos-Eval** achieves the best SA/PC scores across all reported text metrics, indicating
 370 that the Stage-2 controller plus Stage-3 supervision improves both *specificity* (higher BLEU- n) and
 371 *semantic alignment* (higher BERTScore/ROUGE).

3.4 ABLATIONS ON SA AND PC

375 **Setup.** We evaluate two variants on 200 videos randomly sampled from the VideoPhy-2 test set, for
 376 both SA and PC: (i) *w/o Stage-0* (remove the explicit score head; post-hoc map each rationale to
 377 a 5-point score via DeepSeek-R1 (Guo et al., 2025a) using a public rubric); (ii) *w/o Stage-2* (skip
 the controller and use the Stage-1 rationale directly, i.e., no iterative verification). A *single* video-

378 Table 3: **Ablations on SA and PC (VideoPhy-2, $N=200$)**. Correlations vs. human 5-point labels and
 379 VLM-judged reason quality. R-Avg = mean over five rubric dims (SA: Grounding, Temporal Align.,
 380 Consistency, Align Justif., Coverage&Spec.; PC: Grounding, Temporal, Consistency, Criteria&Justif.,
 381 VideoQuality), each in $\{0, 0.5, 1\}$. All rows remap rationale text \rightarrow 5-point score via *DeepSeek-R1* with
 382 a public rubric; a *single* video-conditioned VLM judge is used for both tasks. n = accepted outputs
 383 after the Stage-2 verification gate (when applicable) *and* strict JSON/format checks. **Bold**=best;
 384 underline=second-best.

Legend: Pearson/Spearman = corr. on remapped scores (\uparrow better); R-Avg = judged mean of 5 dims. SA: Ground., Temp., Consist., Align Justif., Cov.&Spec.; PC: Ground., Temp., Consist., C&J, VideoQual.					
Method	n	Pearson \uparrow	Spearman \uparrow	R-Avg \uparrow	Key dim. \uparrow
SA (Semantic Alignment)					
Full (S0+S1+S2)	178	0.8894	0.8866	0.8418	0.9059
w/o Stage-0 (no explicit score head)	188	0.4793	0.4963	0.9142	0.9426
w/o Stage-2 (use S1 rationale directly)	195	<u>0.6727</u>	<u>0.6496</u>	0.8148	0.8413
PC (Physical Commonsense)					
Full (S0+S1+S2)	186	0.9131	0.9112	0.8345	0.9435
w/o Stage-0 (no explicit score head)	194	0.2091	0.1972	<u>0.8309</u>	<u>0.9124</u>
w/o Stage-2 (use S1 rationale directly)	198	<u>0.6502</u>	<u>0.6423</u>	0.7641	0.5328

394 Table 4: **Stage-1 ablations (Cosmos-Eval vs. Moved) on rationale usability (VideoPhy-2, $N=200$)**.
 395 We report *hit-rates* (proportions) of samples with rationale *quality* $\geq \tau$ at preset thresholds $\tau \in$
 396 $\{0.5, 0.6, 0.7, 0.8\}$. *Strict convention*: non-pass treated as 0 (only pass samples can contribute > 0
 397 quality). **Bold** = higher (better).

Model (strict)	SA hit-rate ($\geq \tau$)				PC hit-rate ($\geq \tau$)			
	@0.5	@0.6	@0.7	@0.8	@0.5	@0.6	@0.7	@0.8
Cosmos-Eval	0.775	0.700	0.645	0.600	0.800	0.770	0.725	0.685
Stage-1 Ablation / Moved	0.495	0.470	0.435	0.430	0.270	0.250	0.240	0.220

403 Table 5: **Stage-3 ablations (two-run SFT, joint SA+PC)**. Held-out SA/PC splits as in the main
 404 results. *Two-run SFT*: score head for 5-point labels (1–5) then rationale generation *conditioned on*
 405 *the predicted score* (<think>/<answer>). *Score-only*: fine-tune score head only. *Reason-only*:
 406 fine-tune rationale only. Core metrics: Pearson/Spearman correlations; Acc = exact 5-class accuracy
 407 $\{1, \dots, 5\}$. Reason metrics: BERTScore F1, BLEU-4 on $[0, 1]$. **Bold**=best; underline=second-best.

Model	SA core			PC core			SA reason (0–1)		PC reason (0–1)	
	Pearson	Spearman	Acc	Pearson	Spearman	Acc	B-F1	BLEU-4	B-F1	BLEU-4
Cosmos-Eval (two-run SFT)	<u>0.4643</u>	<u>0.4598</u>	<u>0.3765</u>	0.3641	0.3580	0.3912	<u>0.5244</u>	<u>0.2670</u>	<u>0.5450</u>	0.2786
Score-only SFT	0.5091	0.4984	0.4074	0.3087	<u>0.3065</u>	<u>0.3676</u>	0.3225	0.0443	0.2874	0.0241
Reason-only SFT (CoT)	0.0599	0.0613	0.2074	0.0833	0.0482	0.1001	0.5594	0.3049	0.5455	<u>0.2776</u>

414
 415 conditioned VLM judge (Qwen-VL-Max)² is used for both tasks and applies task-specific rubrics,
 416 averaging five dimensions to R-Avg (SA: Grounding, Temporal Alignment, Consistency, Alignment
 417 Justification, Coverage&Specificity; PC: Grounding, Temporal, Consistency, Criteria&Justification,
 418 VideoQuality). All rows remap rationale text \rightarrow score via *DeepSeek-R1*. We report correlations to
 419 human 5-point labels (Pearson/Spearman) and reason quality (evaluation dimensions detailed in
 420 Appx. C); n counts outputs that *survive the Stage-2 verification gate (when applicable) and strict*
 421 *JSON/format checks*. See Table 3.

422 **Stage-1 ablation (separate analysis).** This is *not* a simple removal of Stage-1. Instead, we replace
 423 Stage-1 with an *alternative verification-only pathway* inside Stage-2: the controller directly judges the
 424 five rubric dimensions without using Stage-1 reference rationales (and without LabelRethink),
 425 functioning as a verifier/filter rather than a score mapper. Accordingly, we report *rationale usability*
 426 via hit-rates of quality $\geq \tau$ with predetermined thresholds $\tau \in \{0.5, 0.6, 0.7, 0.8\}$ under the *strict*
 427 convention (non-pass treated as 0). See Table 4.

428 **Stage-3 ablation (integrated).** Stage 3 uses a *two-run* schedule: (i) a *score-only* pass to calibrate
 429 numeric SA/PC predictions; (ii) a *reasoning* pass that generates <think>/<answer> conditioned

431 ²VLM served via Alibaba Cloud Model Studio; model page: <https://www.alibabacloud.com/help/en/model-studio/vision>.

432 *on the predicted score.* We ablate this by training *Score-only SFT* (omit the reasoning pass) and
 433 *Reason-only SFT* (omit the score pass), and compare to the full **Cosmos-Eval** two-run SFT. We
 434 report *core* score metrics (Pearson/Spearman/Acc) and *reason* quality (BERTScore F1, BLEU-4) for
 435 both SA and PC; see Table 5.

436

437 **Findings.** (A) *Stage-0 (score head) is necessary for calibration.* Removing Stage-0 substantially
 438 weakens agreement with human scores despite strong reason quality (SA: 0.48/0.50; PC: 0.21/0.20),
 439 indicating that calibrated scalar predictions require explicit score supervision.

440

441 (B) *Stage-2 (controller) enforces rubric faithfulness and stabilizes scores.* Skipping Stage-2 de-
 442 grades both correlation and judged quality (SA: 0.673/0.650 with R-Avg=0.815; PC: 0.650/0.642
 443 with R-Avg=0.764; PC Criteria&Justification notably drops to 0.533), underscoring the role of
 444 verification in evidence-grounded reasoning and calibration.

445

446 (C) *Stage-1 reference improves rationale usability/coverage.* Under strict hit-rate evaluation, the
 447 Stage-1 ablation (*Moved*) yields consistently lower usable-rationale coverage than **Cosmos-Eval**
 448 across thresholds (e.g., **SA**: @0.7, 0.645 vs. 0.435; @0.8, 0.600 vs. 0.430; **PC**: @0.7, 0.725 vs. 0.240;
 449 @0.8, 0.685 vs. 0.220), indicating that leveraging Stage-1 reference rationales and the verification
 450 pipeline materially increases the fraction of high-quality, passable explanations.

451

452 (D) *Stage-3 two-run SFT balances scoring & reasoning.* **Cosmos-Eval** attains the best **PC core**
 453 metrics (Pearson 0.3641, Spearman 0.3580, Acc 0.3912) under matched inference budgets throughout
 454 while remaining second on all **SA core** metrics (Pearson 0.4643, Spearman 0.4598, Acc 0.3765);
 455 it is also top-2 on SA/PC reason quality (e.g., PC B-F1 0.5450, BLEU-4 0.2786). *Score-only SFT*
 456 peaks on **SA core** (Pearson 0.5091, Acc 0.4074) but its *reason* quality collapses (SA B-F1/BLEU-4
 457 0.3225/0.0443). *Reason-only SFT* yields the best reasons (SA B-F1/BLEU-4 0.5594/0.3049) yet
 458 fails on **core scoring** (SA Pearson 0.0599; PC Pearson 0.0833).

459

460 **Takeaway.** Across SA and PC, the full configuration ($S_0 + S_1 + S_2$) plus the *Stage-3 two-run schedule*
 461 is the only setting that jointly attains *strong correlations, high reason quality, and high coverage*.
 462 Stage-0 provides calibrated scalar supervision; Stage-2 delivers rubric-faithful verification and
 463 improves stability; Stage-1 contributes substantially to usable-rationale coverage; and Stage-3’s
 464 *scores-first, reasons-conditioned* training preserves **core** agreement while producing **high-quality**
 465 explanations. Removing either Stage-0/2 or one pass in Stage-3 over-optimizes one side.

466

467 4 DISCUSSION

468

469 **Discussion.** The heavy yet interpretable teacher pipeline—Stage 0 (score generation), Stage 1
 470 (reference-anchored rationales), Stage 2 (judge-verified control)—improves SA/PC agreement and
 471 rationale coverage but is compute-intensive (Stage 1/2 dominate). We *distill all three into a Stage 3*
 472 *student* with two-run SFT (score \rightarrow <think>/<answer> conditioned on score), which *replaces*
 473 the ensemble/controller at test time and maintains score fidelity and rationale quality at substantially
 474 lower cost. Ablations show complementary roles (S_0 scoring, S_1 coverage, S_2 verification). Threats
 475 to validity remain (judge bias, rubric shifts, prompt sensitivity, text \rightarrow score remapping) despite
 476 verification safeguards.

477

478 5 CONCLUSION

479

480 We presented **Cosmos-Eval**, an explainable evaluation framework for text-to-video (T2V) that jointly
 481 assesses semantic adherence and physical consistency by coupling 5-point *scores* with concise,
 482 physics-grounded *rationales*. The framework comprises three stages: *Stage 0* score generation,
 483 *Stage 1* reference-seeded reasoning, and *Stage 2* a judge-verified CoT controller. Training follows
 484 a two-round schedule. On *VideoPhy-2* (with *VideoPhy* for recap), **Cosmos-Eval** achieves strong
 485 correlation with human judgments while substantially improving rationale quality over score-only
 486 baselines, enabling targeted diagnosis and more transparent error analysis in T2V evaluation.

487

488

489

486 REFERENCES
487

488 Ge Bai, Jie Liu, Xingyuan Bu, Yancheng He, Jiaheng Liu, Zhanhui Zhou, Zhuoran Lin, Wenbo
489 Su, Tiezheng Ge, Bo Zheng, and Wanli Ouyang. MT-bench-101: A fine-grained benchmark for
490 evaluating large language models in multi-turn dialogues. In *Proceedings of the 62nd Annual*
491 *Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 7421–7454,
492 Bangkok, Thailand, August 2024. Association for Computational Linguistics.

493 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,
494 Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*,
495 2025.

496 Hritik Bansal, Zongyu Lin, Tianyi Xie, Zeshun Zong, Michal Yarom, Yonatan Bitton, Chenfanfu Jiang,
497 Yizhou Sun, Kai-Wei Chang, and Aditya Grover. Videophy: Evaluating physical commonsense for
498 video generation. In *The Thirteenth International Conference on Learning Representations*, 2025a.

499 Hritik Bansal, Clark Peng, Yonatan Bitton, Roman Goldenberg, Aditya Grover, and Kai-Wei Chang.
500 Videophy-2: A challenging action-centric physical commonsense evaluation in video generation.
501 *arXiv preprint arXiv:2503.06800*, 2025b.

502 Omer Bar-Tal, Hila Chefer, Omer Tov, Charles Herrmann, Roni Paiss, Shiran Zada, Ariel Ephrat,
503 Junhwa Hur, Guanghui Liu, Amit Raj, Yuanzhen Li, Michael Rubinstein, Tomer Michaeli, Oliver
504 Wang, Deqing Sun, Tali Dekel, and Inbar Mosseri. Lumiere: A space-time diffusion model for
505 video generation. In *SA '24: SIGGRAPH Asia 2024 Conference Papers*, pp. 94:1–94:11, New
506 York, NY, USA, 2024. Association for Computing Machinery.

507 Daniel Bear, Elias Wang, Damian Mrowca, Felix J. Binder, Hsiao-Yu Tung, R. T. Pramod, Cameron
508 Holdaway, Sirui Tao, Kevin A. Smith, Fan-Yun Sun, Fei-Fei Li 0001, Nancy Kanwisher, Josh Tenen-
509 baum 0001, Dan Yamins, and Judith E. Fan. Physion: Evaluating physical prediction from vision
510 in humans and machines. In *Proceedings of the Neural Information Processing Systems Track on*
511 *Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*,
512 2021.

513 Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik
514 Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling
515 latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.

516 Haoxin Chen, Yong Zhang, Xiaodong Cun, Menghan Xia, Xintao Wang, Chao Weng, and Ying
517 Shan. Videocrafter2: Overcoming data limitations for high-quality video diffusion models. In
518 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
519 7310–7320, 2024.

520 Junying Chen, Zhenyang Cai, Ke Ji, Xidong Wang, Wanlong Liu, Rongsheng Wang, and Benyou
521 Wang. Towards medical complex reasoning with LLMs through medical verifiable problems. In
522 *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 14552–14573, Vienna,
523 Austria, July 2025a. Association for Computational Linguistics.

524 Yongfan Chen, Xiuwen Zhu, and Tianyu Li. A physical coherence benchmark for evaluating video
525 generation models via optical flow-guided frame prediction. *arXiv preprint arXiv:2502.05503*,
526 2025b.

527 Chaoyou Fu, Yuhang Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu
528 Zhou, Yunhang Shen, Mengdan Zhang, et al. Video-mme: The first-ever comprehensive evaluation
529 benchmark of multi-modal llms in video analysis. In *Proceedings of the Computer Vision and*
530 *Pattern Recognition Conference*, pp. 24108–24118, 2025.

531 Deepanway Ghosal, Navonil Majumder, Roy Lee, Rada Mihalcea, and Soujanya Poria. Language
532 guided visual question answering: Elevate your multimodal language model using knowledge-
533 enriched prompts. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Findings of the Associa-*
534 *tion for Computational Linguistics: EMNLP 2023*, pp. 12096–12102, December 2023.

535 Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan
536 Shen, Shengjie Ma, Honghao Liu, Yuanzhuo Wang, and Jian Guo. A survey on llm-as-a-judge.
537 *arXiv preprint arXiv: 2411.15594*, 2024.

540 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, et al. DeepSeek-R1 incentivizes reasoning
 541 in LLMs through reinforcement learning. *Nature*, 645(8081):633–638, September 2025a. ISSN
 542 1476-4687.

543 Xuyang Guo, Jiayan Huo, Zhenmei Shi, Zhao Song, Jiahao Zhang, and Jiale Zhao. T2vphysbench:
 544 A first-principles benchmark for physical consistency in text-to-video generation. *arXiv preprint*
 545 *arXiv:2505.00337*, 2025b.

546 Xuan He, Dongfu Jiang, Ge Zhang, Max Ku, Achint Soni, Sherman Siu, Haonan Chen, Abhranil
 547 Chandra, Ziyan Jiang, Aaran Arulraj, Kai Wang, Quy Duc Do, Yuansheng Ni, Bohan Lyu, Yaswanth
 548 Narsupalli, Rongqi Fan, Zhiheng Lyu, Bill Yuchen Lin, and Wenhu Chen. VideoScore: Building
 549 automatic metrics to simulate fine-grained human feedback for video generation. In *Proceedings*
 550 of the 2024 Conference on Empirical Methods in Natural Language Processing, pp. 2105–2123,
 551 Miami, Florida, USA, November 2024. Association for Computational Linguistics.

552 Yingqing He, Tianyu Yang, Yong Zhang, Ying Shan, and Qifeng Chen. Latent video diffusion models
 553 for high-fidelity long video generation. *arXiv preprint arXiv:2211.13221*, 2022.

554 Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. CLIPScore: A
 555 reference-free evaluation metric for image captioning. In *Proceedings of the 2021 Conference*
 556 on Empirical Methods in Natural Language Processing, pp. 7514–7528, Online and Punta Cana,
 557 Dominican Republic, November 2021. Association for Computational Linguistics.

558 Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P
 559 Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition
 560 video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022.

561 Haoyang Huang, Guoqing Ma, Nan Duan, Xing Chen, Changyi Wan, Ranchen Ming, Tianyu Wang,
 562 Bo Wang, Zhiying Lu, Aojie Li, et al. Step-video-ti2v technical report: A state-of-the-art text-
 563 driven image-to-video generation model. *arXiv preprint arXiv:2503.11251*, 2025.

564 Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing
 565 Wu, Qingyang Jin, Nattapol Chanpaisit, et al. Vbench: Comprehensive benchmark suite for video
 566 generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
 567 *Recognition*, pp. 21807–21818, 2024.

568 Pengliang Ji, Chuyang Xiao, Huilin Tai, and Mingxiao Huo. T2vbench: Benchmarking temporal
 569 dynamics for text-to-video generation. In *Proceedings of the IEEE/CVF Conference on Computer*
 570 *Vision and Pattern Recognition*, pp. 5325–5335, 2024.

571 Yang Jin, Zhicheng Sun, Ningyuan Li, Kun Xu, Kun Xu, Hao Jiang, Nan Zhuang, Quzhe Huang,
 572 Yang Song, Yadong MU, and Zhouchen Lin. Pyramidal flow matching for efficient video generative
 573 modeling. In *The Thirteenth International Conference on Learning Representations*, 2025.

574 Daneul Kim, Jingxu Zhang, Wonjoon Jin, Sunghyun Cho, Qi Dai, Jaesik Park, and Chong
 575 Luo. Subject-driven video generation via disentangled identity and motion. *arXiv preprint*
 576 *arXiv:2504.17816*, 2025.

577 Dan Kondratyuk, Lijun Yu, Xiuye Gu, Jose Lezama, Jonathan Huang, Grant Schindler, Rachel
 578 Hornung, Vighnesh Birodkar, Jimmy Yan, Ming-Chang Chiu, Krishna Somandepalli, Hassan
 579 Akbari, Yair Alon, Yong Cheng, Joshua V. Dillon, Agrim Gupta, Meera Hahn, Anja Hauth, David
 580 Hendon, Alonso Martinez, David Minnen, Mikhail Sirotenko, Kihyuk Sohn, Xuan Yang, Hartwig
 581 Adam, Ming-Hsuan Yang, Irfan Essa, Huisheng Wang, David A Ross, Bryan Seybold, and Lu Jiang.
 582 VideoPoet: A large language model for zero-shot video generation. In *Proceedings of the 41st*
 583 *International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning*
 584 *Research*, pp. 25105–25124. PMLR, 21–27 Jul 2024.

585 Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li,
 586 Bo Wu, Jianwei Zhang, Kathrina Wu, Qin Lin, Junkun Yuan, Yanxin Long, Aladdin Wang, Andong
 587 Wang, Changlin Li, Duojun Huang, Fang Yang, Hao Tan, Hongmei Wang, Jacob Song, Jiawang
 588 Bai, Jianbing Wu, Jinbao Xue, Joey Wang, Kai Wang, Mengyang Liu, Pengyu Li, Shuai Li, Weiyan
 589 Wang, Wenqing Yu, Xinchi Deng, Yang Li, Yi Chen, Yutao Cui, Yuanbo Peng, Zhentao Yu, Zhiyu

594 He, Zhiyong Xu, Zixiang Zhou, Zunnan Xu, Yangyu Tao, Qinglin Lu, Songtao Liu, Dax Zhou,
 595 Hongfa Wang, Yong Yang, Di Wang, Yuhong Liu, Jie Jiang, and Caesar Zhong. Hunyuandvideo: A
 596 systematic framework for large video generative models, 2025.

597

598 Philippe Laban, Tobias Schnabel, Paul N. Bennett, and Marti A. Hearst. Summac: Re-visiting
 599 nli-based models for inconsistency detection in summarization. *Transactions of the Association
 600 for Computational Linguistics*, 10:163–177, 02 2022. ISSN 2307-387X.

601 Dacheng Li, Yunhao Fang, Yukang Chen, Shuo Yang, Shiyi Cao, Justin Wong, Michael Luo, Xiaolong
 602 Wang, Hongxu Yin, Joseph E Gonzalez, et al. Worldmodelbench: Judging video generation models
 603 as world models. *arXiv preprint arXiv:2502.20694*, 2025a.

604

605 Jiefeng Li, Jinkun Cao, Haotian Zhang, Davis Rempe, Jan Kautz, Umar Iqbal, and Ye Yuan. Genmo:
 606 A generalist model for human motion, 2025b.

607

608 Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen,
 609 Ping Luo, et al. Mvbench: A comprehensive multi-modal video understanding benchmark. In
 610 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
 22195–22206, 2024.

611

612 Jiahe Liu, Youran Qu, Qi Yan, Xiaohui Zeng, Lele Wang, and Renjie Liao. Fréchet video motion
 613 distance: A metric for evaluating motion consistency in videos. In *First Workshop on Controllable
 614 Video Generation @ICML24*, 2024a.

615

616 Ming Liu and Wensheng Zhang. Is your video language model a reliable judge? In *The Thirteenth
 617 International Conference on Learning Representations*, 2025.

618

619 Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. G-eval: NLG
 620 evaluation using gpt-4 with better human alignment. In *Proceedings of the 2023 Conference on
 621 Empirical Methods in Natural Language Processing*, pp. 2511–2522, Singapore, December 2023.
 Association for Computational Linguistics.

622

623 Yaofang Liu, Xiaodong Cun, Xuebo Liu, Xintao Wang, Yong Zhang, Haoxin Chen, Yang Liu,
 624 Tieyong Zeng, Raymond Chan, and Ying Shan. Evalcrafter: Benchmarking and evaluating large
 625 video generation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and
 626 Pattern Recognition*, pp. 22139–22149, 2024b.

627

628 Yixin Liu, Kai Zhang, Yuan Li, Zhiling Yan, Chujie Gao, Ruoxi Chen, Zhengqing Yuan, Yue
 629 Huang, Hanchi Sun, Jianfeng Gao, Lifang He, and Lichao Sun. Sora: A Review on Background,
 Technology, Limitations, and Opportunities of Large Vision Models, April 2024c.

630

631 Muhammad Maaz, Hanooona Rasheed, Salman Khan, and Fahad Khan. Video-ChatGPT: Towards
 632 detailed video understanding via large vision and language models. In *Proceedings of the 62nd
 633 Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp.
 12585–12602, Bangkok, Thailand, August 2024. Association for Computational Linguistics.

634

635 Pascal Massart and Élodie Nédélec. Risk bounds for statistical learning. *The Annals of Statistics*, 34
 636 (5):2326–2366, 2006. ISSN 00905364.

637

638 Zhun Mou, Bin Xia, Zhengchao Huang, Wenming Yang, and Jiaya Jia. GRADEO: Towards human-
 639 like evaluation for text-to-video generation via multi-step reasoning. In *Forty-second International
 Conference on Machine Learning*, 2025.

640

641 NVIDIA, Alisson Azzolini, Hannah Brandon, Prithvijit Chattopadhyay, Huayu Chen, Jinju Chu, Yin
 642 Cui, Jenna Diamond, Yifan Ding, Francesco Ferroni, Rama Govindaraju, Jinwei Gu, Siddharth
 643 Gururani, Imad El Hanafi, Zekun Hao, Jacob Huffman, Jingyi Jin, Brendan Johnson, Rizwan
 644 Khan, George Kurian, Elena Lantz, Nayeon Lee, Zhaoshuo Li, Xuan Li, Tsung-Yi Lin, Yen-
 645 Chen Lin, Ming-Yu Liu, Andrew Mathau, Yun Ni, Lindsey Pavao, Wei Ping, David W. Romero,
 646 Misha Smelyanskiy, Shuran Song, Lyne Tchapmi, Andrew Z. Wang, Boxin Wang, Haoxiang
 647 Wang, Fangyin Wei, Jiashu Xu, Yao Xu, Xiaodong Yang, Zhuolin Yang, Xiaohui Zeng, and Zhe
 Zhang. Cosmos-reason1: From physical common sense to embodied reasoning. *arXiv preprint
 648 arXiv:2503.15558*, 2025.

648 OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, et al. Gpt-4 technical report, 2024.
649

650 Ronan Riochet, Mario Castro, Mathieu Bernard, Adam Lerer, Rob Fergus, Véronique Izard, and
651 Emmanuel Dupoux. Intphys: A framework and benchmark for visual intuitive physics reasoning.
652 *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PP, 03 2018.

653 Ronan Riochet, Mario Ynocente Castro, Mathieu Bernard, Adam Lerer, Rob Fergus, Véronique Izard,
654 and Emmanuel Dupoux. Intphys 2019: A benchmark for visual intuitive physics understanding.
655 *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):5016–5025, 2021.

656

657 SP Sharan, Minkyu Choi, Sahil Shah, Harsh Goel, Mohammad Osama, and Sandeep Chinchali.
658 Neuro-symbolic evaluation of text-to-video models using formal verification. In *Proceedings of*
659 *the Computer Vision and Pattern Recognition Conference*, pp. 8395–8405, 2025.

660

661 Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry
662 Yang, Oron Ashual, Oran Gafni, Devi Parikh, Sonal Gupta, and Yaniv Taigman. Make-a-video:
663 Text-to-video generation without text-video data. In *The Eleventh International Conference on*
664 *Learning Representations*, 2023.

665

666 Enxin Song, Wenhao Chai, Guanhong Wang, Yucheng Zhang, Haoyang Zhou, Feiyang Wu, Haozhe
667 Chi, Xun Guo, Tian Ye, Yanting Zhang, et al. Moviechat: From dense token to sparse memory for
668 long video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
669 *Pattern Recognition*, pp. 18221–18232, 2024.

670

671 Haibo Tong, Zhaoyang Wang, Zhaorun Chen, Haonian Ji, Shi Qiu, Siwei Han, Kexin Geng, Zhongkai
672 Xue, Yiyang Zhou, Peng Xia, et al. Mj-video: Fine-grained benchmarking and rewarding video
673 preferences in video generation. *arXiv preprint arXiv:2502.01719*, 2025.

674

675 Hsiao-Yu Tung, Mingyu Ding, Zhenfang Chen, Daniel Bear, Chuang Gan, Josh Tenenbaum, Dan
676 Yamins, Judith Fan, and Kevin Smith. Physion++: Evaluating physical scene understanding
677 that requires online inference of different physical properties. *Advances in Neural Information*
678 *Processing Systems*, 36:67048–67068, 2023.

679

680 Ruben Villegas, Mohammad Babaizadeh, Pieter-Jan Kindermans, Hernan Moraldo, Han Zhang,
681 Mohammad Taghi Saffar, Santiago Castro, Julius Kunze, and Dumitru Erhan. Phenaki: Variable
682 length video generation from open domain textual descriptions. In *International Conference on*
683 *Learning Representations*, 2023.

684

685 Jiawei Wang, Liping Yuan, Yuchen Zhang, and Haomiao Sun. Tarsier: Recipes for training and
686 evaluating large video description models. *arXiv preprint arXiv:2407.00634*, 2024a.

687

688 Weihan Wang, Zehai He, Wenyi Hong, Yean Cheng, Xiaohan Zhang, Ji Qi, Xiaotao Gu, Shiyu Huang,
689 Bin Xu, Yuxiao Dong, et al. Lvbench: An extreme long video understanding benchmark. *arXiv*
690 *preprint arXiv:2406.08035*, 2024b.

691

692 Haoning Wu, Dongxu Li, Bei Chen, and Junnan Li. Longvideobench: A benchmark for long-context
693 interleaved video-language understanding. *Advances in Neural Information Processing Systems*,
694 37:28828–28857, 2024.

695

696 Xinhao Xiang, Xiao Liu, Zizhong Li, Zhuosheng Liu, and Jiawei Zhang. Aigve-tool: Ai-generated
697 video evaluation toolkit with multifaceted benchmark, 2025.

698

699 Jinbo Xing, Menghan Xia, Yong Zhang, Haoxin Chen, Wangbo Yu, Hanyuan Liu, Gongye Liu,
700 Xintao Wang, Ying Shan, and Tien-Tsin Wong. Dynamicrafter: Animating open-domain images
701 with video diffusion priors. In *European Conference on Computer Vision*, pp. 399–417. Springer,
702 2024.

703

704 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang
705 Gao, Chengan Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*,
706 2025a.

702 Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang,
703 Wenyi Hong, Xiaohan Zhang, Guanyu Feng, Da Yin, Yuxuan Zhang, Weihan Wang, Yean Cheng,
704 Bin Xu, Xiaotao Gu, Yuxiao Dong, and Jie Tang. CogVideoX: Text-to-Video Diffusion Models
705 with An Expert Transformer, March 2025b.

706 Kexin Yi, Chuang Gan, Yunzhu Li, Pushmeet Kohli, Jiajun Wu, Antonio Torralba, and Joshua B
707 Tenenbaum. Clevrer: Collision events for video representation and reasoning. In *International*
708 *Conference on Learning Representations*, 2020.

709 Boqiang Zhang, Kehan Li, Zesen Cheng, Zhiqiang Hu, Yuqian Yuan, Guanzheng Chen, Sicong Leng,
710 Yuming Jiang, Hang Zhang, Xin Li, et al. Videollama 3: Frontier multimodal foundation models
711 for image and video understanding. *arXiv preprint arXiv:2501.13106*, 2025.

712 Hang Zhang, Xin Li, and Lidong Bing. Video-LLaMA: An instruction-tuned audio-visual language
713 model for video understanding. In *Proceedings of the 2023 Conference on Empirical Methods in*
714 *Natural Language Processing: System Demonstrations*, pp. 543–553, Singapore, December 2023.
715 Association for Computational Linguistics.

716 Dian Zheng, Ziqi Huang, Hongbo Liu, Kai Zou, Yinan He, Fan Zhang, Yuanhan Zhang, Jingwen
717 He, Wei-Shi Zheng, Yu Qiao, et al. Vbench-2.0: Advancing video generation benchmark suite for
718 intrinsic faithfulness. *arXiv preprint arXiv:2503.21755*, 2025.

719 Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen
720 Duan, Weijie Su, Jie Shao, et al. Internvl3: Exploring advanced training and test-time recipes for
721 open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025.

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756	APPENDIX	
757		
758		
759	A Related Work	16
760		
761	B Implementation and Training Details	17
762	B.1 Method overview (flow)	17
763	B.2 Datasets and protocol (recap)	17
764	B.3 Inference hyperparameters (Stages 1–2)	17
765		
766		
767	C Ablations (Extended): Methods, Rubrics, and Results	17
768	C.1 PC Evaluation Rubric (VLM-as-judge)	17
769	C.2 SA Evaluation Rubric (VLM-as-judge)	17
770		
771		
772	D Case Analysis	18
773		
774		
775	E Examples for Physical Commonsense (PC) and Semantic Alignment (SA) Tasks	19
776	E.1 Physical Commonsense (PC) Examples	19
777	E.2 Semantic Alignment (SA) Examples	19
778		
779		
780	F Formal Analysis	19
781	F.1 Notation and Terminology	19
782	F.2 Fundamental Assumptions	20
783	F.3 Stage 1: Consensus Aggregation and Noise Reduction	21
784	F.4 Stage 2: Controller Pass Probability and Error Analysis	22
785	F.5 Stage 3: Generalization Bound under Noisy Supervision	23
786	F.6 Sufficient Condition for Multi-Stage Superiority	23
787	F.7 Summary and Empirical Validation Suggestions	24
788		
789		
790		
791		
792		
793	G Additional Experiments and Analyses	24
794		
795	G.1 Cross-Benchmark Generalization	25
796	G.2 Ranking T2V Generators on AIGVE-Bench	25
797	G.3 Human Evaluation of SA/PC Rationales	25
798	G.4 VLM-Judge Evaluation of Rationales	26
799	G.5 Frontier VLM Baselines on VideoPhy-2	27
800	G.6 Factual Consistency Metrics for Rationales	27
801	G.7 Uncertainty Estimates for Main SA/PC Results	28
802	G.8 Sensitivity to Stage-2 Judge and External Scorers	28
803	G.9 Computational Cost and Efficiency	29
804	G.10 Robustness to Synthetic Degradations	30
805		
806		
807		
808		
809	G.11 Long-Horizon Evaluation on LongCat-Video	30

810	H Reproducibility statement	31
811		
812		
813	I The Use of Large Language Models (LLMs)	31
814		
815	J Prompt Templates	31
816		
817		
818	A RELATED WORK	
819		

820 **Text-to-video systems and video LLMs.** Recent text-to-video (T2V) systems establish scalable
 821 diffusion/transformer pipelines and practical recipes for longer, more controllable videos: Make-A-
 822 Video, Imagen Video, Phenaki, and latent video diffusion models laid the foundations for latent spaces
 823 and variable-length synthesis (Singer et al., 2023; Ho et al., 2022; Villegas et al., 2023; He et al.,
 824 2022). Subsequent open frameworks emphasize data efficiency and motion fidelity (VideoCrafter2,
 825 DynamicalCrafter) and push controllability via step-wise refinement and identity-motion disentangle-
 826 ment (Chen et al., 2024; Xing et al., 2024; Huang et al., 2025; Kim et al., 2025). In parallel,
 827 instruction-tuned video LLMs (Video-LLaMA, Video-ChatGPT) and long-video models (MovieChat)
 828 enable free-form QA and temporal reasoning over extended content (Zhang et al., 2023; Maaz et al.,
 829 2024; Song et al., 2024). Our work does not introduce a new generator or Vid-LLM; instead, we
 830 contribute an *explainable evaluator* that grades generated videos along *semantic adherence (SA)* and
 831 *physical commonsense (PC)* while producing rationales.

832 **SA/PC-oriented evaluators and benchmarks.** Foundational benchmarks explicitly target SA/PC.
 833 VIDEOPHY (Bansal et al., 2025a) is the first to formalize both axes, curating 688 prompts across
 834 three material-interaction types (solid–solid, solid–fluid, fluid–fluid) and introducing VIDEOCON-
 835 PHYSICS, an automatic evaluator for SA/PC. However, VIDEOPHY uses binary (0/1) scoring and lacks
 836 fine-grained physical-rule annotations, making it difficult to diagnose failure modes. VIDEOPHY-2
 837 (Bansal et al., 2025b) expands the scope to 197 real-world actions and provides a hard subset (60
 838 actions where top models such as Wan2.1-14B reach only 21.9% joint SA/PC). It further introduces
 839 **VIDEOPHY-2-AUTOEVAL**, an automatic evaluator that outputs 5-point SA/PC scores and tags
 840 physical-rule violations (e.g., conservation of momentum), with substantially improved correlation to
 841 human PC scores (reported to outperform Gemini-2.0-Flash by 236%). Like its predecessor, it outputs
 842 scores but not explanatory rationales, limiting interpretability and error analysis. Complementary
 843 physics-fidelity suites (e.g., T2VPhysBench (Guo et al., 2025b), PhyCoBench (Chen et al., 2025b))
 844 emphasize physical realism yet similarly provide limited support for explanation.

845 **General video evaluation and reference-free quality.** Evaluation resources for video under-
 846 standing and generation are complementary to our goal. MVBench and Video-MME target broad
 847 multimodal comprehension; LongVideoBench and LVbench probe long-horizon temporal reasoning
 848 (Li et al., 2024; Fu et al., 2025; Wu et al., 2024; Wang et al., 2024b). For generation, VBench and
 849 VBench-2.0 decompose quality into fine-grained dimensions; EvalCrafter and T2VBench provide
 850 diverse prompts and temporal diagnostics; learned assessors (VideoScore) and flow/motion-centric
 851 metrics (FVMD) complement reference-free alignment such as CLIPScore (Huang et al., 2024;
 852 Zheng et al., 2025; Liu et al., 2024b; Ji et al., 2024; He et al., 2024; Liu et al., 2024a; Hessel et al.,
 853 2021). Beyond aesthetics and prompt match, physics-centric diagnostics from IntPhys, CLEVRER,
 854 Physion, and Physion++ probe object permanence, collisions, and latent properties (Riochet et al.,
 855 2018; 2021; Yi et al., 2020; Bear et al., 2021; Tung et al., 2023); emerging “world-model” evaluations
 856 and neuro-symbolic checks broaden this perspective (Sharan et al., 2025; Li et al., 2025a; Tong et al.,
 857 2025).

858 **LLM-as-a-judge and reliability.** LLM-as-a-judge methods (e.g., G-Eval, MT-Bench-101) and
 859 subsequent reliability analyses inform our design choices: score-conditioned consensus/selection,
 860 and a unified pass/fail verifier whose distilled behavior stabilizes deployment (Liu et al., 2023;
 861 Bai et al., 2024; Liu & Zhang, 2025). In contrast to prior SA/PC evaluators that primarily output
 862 scores, our evaluator couples *calibrated scoring* with *rubric-faithful rationales* and fine-grained
 863 rubric dimensions, enabling actionable diagnostics and safer iteration.

864 **Table 6: Inference configuration for Stages 1–2.** SA aggregates $M=2$ reasons by consensus (Eq. 4);
 865 PC samples $K=5$ candidates and selects the best (Eq. 6); the Stage-2 controller runs for $N=3$ steps
 866 with strategy sampling without replacement. We list generators and decoding settings (temperature,
 867 top- p , max tokens) plus the effective sampling fps . A dash (—) denotes not applicable.

Task/Stage	Generator(s)	Pool/Budget	Temp	Top- p	Max tokens	Max frames/fps
SA / Stage-1	Tarsier-34B, Qwen2.5-VL-72B-Instruct	$M = 2$	0.7, 0.3	0.85, 0.85	1024, 1024	32 / 8
PC / Stage-1	Cosmos-Reason1	$K = 5$	0.8	0.9	8192	— / 8
SA Aggregator	Qwen3-32B (Yang et al., 2025a)	—	0.7	0.85	2048	—
PC Selector	Qwen2.5-VL-72B-Instruct-AWQ	—	0.1	0.9	1024	— / 8
SA / Stage-2 Controller	Qwen2.5-VL-72B-Instruct	$N = 3$	0.3	0.85	16384	— / 2
PC / Stage-2 Controller	Qwen2.5-VL-72B-Instruct-AWQ	$N = 3$	0.3	0.85	16384	— / 2
SA LLM Judge \mathcal{V}_{sa}	Qwen2.5-VL-72B-Instruct	—	0.05	0.95	50	—
PC LLM Judge \mathcal{V}_{pc}	Qwen2.5-VL-72B-Instruct-AWQ	—	0.05	0.95	50	—

877 *Legend: M = SA Stage-1 ensemble size (one reason per model); K = PC Stage-1 candidate count; N = Stage-2
 878 controller attempt budget (strategies sampled without replacement). Max frames/fps: “Max frames” applies only
 879 to Tarsier-34B (Wang et al., 2024a) (cap at 32 frames); Qwen-family rows use streaming at the listed fps
 880 (no frame cap). “—” = not applicable.*

B IMPLEMENTATION AND TRAINING DETAILS

B.1 METHOD OVERVIEW (FLOW)

886 Figures 3–5 give a concise view of Stages 1–2, and Algorithm 1 formalizes the Stage 2 controller.
 887 For **SA** (Fig. 3), we ensemble several VLMs to propose reasons and take a consensus as the reference
 888 to seed Stage 2. For **PC** (Fig. 4), a base VLM samples multiple reasons and a VLM judge selects one
 889 as the reference. **Stage 2** (Fig. 5; Alg. 1) then iteratively refines and judge-verifies candidates (with a
 890 label-rethink fallback), and formats the accepted chain as the final reason.

B.2 DATASETS AND PROTOCOL (RECAP)

893 We train on the union of *VideoPhy* (Bansal et al., 2025a) (train+test, re-scored by *VIDEOPHY-2*-
 894 *AUTOEVAL*) and *VideoPhy-2* (Bansal et al., 2025b) (train), and evaluate on the official *VideoPhy-2*
 895 test set. Task inputs follow Sec. 2: SA uses (v, c) and PC uses v only. Figure 6 summarizes the
 896 SA/PC score distributions across corpora and our final splits.

B.3 INFERENCE HYPERPARAMETERS (STAGES 1–2)

900 Stage 1 uses an ensemble size $M=2$ for SA (Eq. 3) and $K=5$ samples for PC (Eq. 5); Stage 2 runs
 901 with budget $N=3$ and *strategy sampling without replacement* (Sec. 2.3). A complete list of generators,
 902 judge/aggregator models, and decoding settings (temperature, top- p , max tokens) is summarized in
 903 Table 6. SA reasons are aggregated by consensus (Eq. 4); PC reasons are selected by a judge (Eq. 6).

C ABLATIONS (EXTENDED): METHODS, RUBRICS, AND RESULTS

C.1 PC EVALUATION RUBRIC (VLM-AS-JUDGE)

909 We use the five-dimension rubric in Table 7 (Ground., Temp., Cons., C&J, VideoQual), with 3-point
 910 anchors $\{0, 0.5, 1\}$ matching the judge prompt. The same rubric is applied to all ablations in Sec. 3.4.

C.2 SA EVALUATION RUBRIC (VLM-AS-JUDGE)

915 We adopt a five-dimension rubric for Semantic Alignment (SA), shown in Table 8, with three-point
 916 anchors $\{0, 0.5, 1\}$ matching the evaluation prompt. The rubric is applied consistently across all SA
 917 ablations in Sec. 3.4. Concretely checkable details include (non-exhaustively): color, region/relative
 918 position, count/frequency, motion attributes, and deformation/rigidity.

918
 919 **Table 7: PC reason-quality rubric used in ablation studies** (Sec. 3.4). Five dimensions with 3-point
 920 anchors {0, 0.5, 1}, matching the evaluation prompt. “Concrete, checkable details” include color,
 921 region/relative position, count/frequency, motion attributes, and deformation/rigidity.

Dim.	Score 1	Score 0.5	Score 0
Ground.	≥ 2 concrete details clearly support the claims.	Generic/vague match to visuals.	Conflicts with visuals / speculative.
Temp.	≥ 1 concrete, correct temporal relation.	Gist generic/unclear or N/A/uncertain.	Wrong/reversed/invented temporal claims.
Cons.	Internally consistent; no contradictions or hallucinated key objects/events.	Minor issue; main claim intact.	Contradiction or hallucination.
C&J	Explicit criterion/score/rule applied to visible evidence.	Mentioned but generic/partial/weak.	None or misapplied/contradicted by evidence.
VideoQual	Explicit good/bad (or degree) with ≥ 2 indicators (sharpness, lighting, occlusion, stability, framing, target visibility).	Generic or only one indicator / uncertain.	No quality judgment or contradicts visuals.

930 *Abbrev.* Ground.=Grounding; Temp.=Temporal; Cons.=Consistency; C&J=Criteria & Justification;
 931 VideoQual=Video Quality Assessment.

932 *Hard cap:* if no concrete visual detail appears, **Ground.** ≤ 0.5 .

934 **Table 8: SA reason-quality rubric used in ablation studies** (Sec. 3.4). Five dimensions with 3-point
 935 anchors {0, 0.5, 1}, matching the evaluation prompt. “Concrete, checkable details” include color,
 936 region/relative position, count/frequency, motion attributes, and deformation/rigidity.

Dim.	Score 1	Score 0.5	Score 0
Ground.	≥ 2 concrete details linking CAPTION \leftrightarrow VIDEO.	Generic/partial visual match.	Conflicts with CAPTION/VIDEO or speculative.
Temp.	≥ 1 concrete, correct temporal relation.	Gist generic/unclear or N/A/uncertain.	Wrong/reversed/invented temporal claims.
Cons.	Internally consistent; no hallucinated key objects/events.	Minor issue; main claim intact.	Contradiction or hallucination.
Align Justif.	Explicit SA decision/criterion applied to visible evidence.	Mentioned but generic/partial/weak.	None or misapplied/contradicted by evidence.
Cov.&Spec.	Covers ≥ 2 key CAPTION elements with specific, checkable details.	Some elements but incomplete/generic.	Ignores key elements or no specific details.

937 *Abbrev.* Ground.=Grounding; Temp.=Temporal Alignment; Cons.=Consistency; Align Justif.=Alignment
 938 Justification; Cov.&Spec.=Coverage & Specificity.

939 *Hard cap:* if no concrete visual detail appears, **Ground.** ≤ 0.5 .

940 D CASE ANALYSIS

941 To assess the reliability of our evaluator COSMOS-EVAL, we present its *verbatim* answers in the
 942 figure captions and provide brief justifications here for **Cases 1–4** (see Fig. 7–10). In each case, the
 943 model correctly identifies the salient mismatch or physical violation.

944 **Case 1 (PC=2; Fig. 7).** The video shows a red ball *hovering* without visible support. This
 945 contradicts gravitational expectations (no external force, yet no downward acceleration). COSMOS-
 946 EVAL’s answer pinpoints the violation and a low PC score is appropriate.

947 **Case 2 (SA=2; Fig. 8).** The caption specifies *counterclockwise* rotation, while the video shows
 948 the yellow cube rotating *clockwise*; the purple cone remains still. COSMOS-EVAL correctly isolates
 949 the direction-of-rotation mismatch—the primary semantic attribute here. Although its text suggests
 950 *sa_score* = 3, our rubric weights action direction as critical, yielding **SA=2**. The qualitative diagnosis
 951 is consistent with our ground truth.

952 **Case 3 (PC=2; Fig. 9).** The ball exhibits erratic back-and-forth bounces with no frictional decay
 953 and no plausible external impulses. COSMOS-EVAL accurately characterizes this as inconsistent with
 954 Newtonian mechanics, justifying **PC=2**.

955 **Case 4 (SA=3; Fig. 10).** The caption describes *one* ball being kicked to the post and rebounding,
 956 but the video shows *two* balls and lacks the kick–post–rebound sequence. COSMOS-EVAL correctly

972 flags the count mismatch and the missing key action; scene context matches but the core event does
 973 not, supporting **SA=3** for partial alignment.
 974

975 Overall, COSMOS-EVAL’s answers consistently identify the correct failure modes (semantic or
 976 physical), and they qualitatively agree with our human labels, demonstrating useful explanatory
 977 power and reproducibility.
 978

979 E EXAMPLES FOR PHYSICAL COMMONSENSE (PC) AND SEMANTIC 980 ALIGNMENT (SA) TASKS 981

982 E.1 PHYSICAL COMMONSENSE (PC) EXAMPLES 983

984 Figure 11 shows the first example for the Physical Commonsense task, where we evaluate the physical
 985 properties of the video. Figure 12 demonstrates another case with similar evaluation criteria. Figures
 986 13, 14, and 15 further illustrate other examples related to the Physical Commonsense task.
 987

988 In addition, Figures 21 and 22 present two representative Physical Commonsense cases with full
 989 chain-of-thought traces and final rationales generated by Cosmos-Eval. These examples make the
 990 5-point scores and the corresponding physics-aware explanations explicit and are intended as concrete
 991 case studies to complement the aggregated metrics in the main text.
 992

993 E.2 SEMANTIC ALIGNMENT (SA) EXAMPLES 994

995 Figure 16 presents the first example for the Semantic Alignment task, evaluating the alignment
 996 between the caption and video content. Figure 17 shows another example with slightly different
 997 criteria. Figures 18, 19, and 20 provide additional examples for the Semantic Alignment task.
 998

999 Figures 23 and 24 further provide Semantic Alignment case studies with explicit chain-of-thought
 1000 reasoning and natural-language rationales from Cosmos-Eval. These SA examples illustrate how the
 1001 model justifies its 5-point scores by grounding the caption–video comparison in concrete events and
 1002 entities, addressing the reviewer’s request for more detailed CoT-style examples and error analysis.
 1003

1004 F FORMAL ANALYSIS 1005

1006 This section provides a formal analysis of the proposed multi-stage framework, focusing on the
 1007 conditions under which it achieves better generalization than end-to-end (E2E) learning. Rather
 1008 than offering strict proofs, the analysis establishes a set of assumptions and derives conditions that
 1009 characterize the effective noise reduction at different stages.
 1010

1011 We first introduce the notation and assumptions used throughout. We then examine the noise-
 1012 mitigation mechanisms in Stage 1 (consensus aggregation, Section 2.2) and Stage 2 (controlled
 1013 generation, Section 2.3). Finally, drawing on information-theoretic and learning-theoretic perspec-
 1014 tives, we identify sufficient conditions under which the multi-stage framework yields a supervision
 1015 signal with a lower effective noise rate than E2E learning, thereby leading to a tighter upper bound
 1016 on the generalization error.
 1017

1018 F.1 NOTATION AND TERMINOLOGY 1019

1020 To maintain consistency with Section 2, we define the unified notation for this theoretical analysis:
 1021

- 1022 • **Task Index:** $\tau \in \{\text{sa, pc}\}$, denoting the Semantic Adherence and Physical Commonsense tasks,
 1023 respectively.
- 1024 • **Input:** X^τ or its instance x^τ . For SA, $x^{\text{sa}} = (v, c)$ (video v and caption c); for PC, $x^{\text{pc}} = v$ (video
 1025 only).
- 1026 • **True Label:** $Y^\tau \in \{1, \dots, 5\}$, representing the discrete ground-truth score (5-point scale).
- 1027 • **Stage 0 Output:** $S^\tau \in \{1, \dots, 5\}$, the initial score from the VideoPhy model, serving as side
 1028 information.

- **Stage 1 Reference Rationale:** r_{ref}^τ , the output of Stage 1 for task τ , used as the initial seed for Stage 2.
- **Stage 2 Evidence and Rationale:** (e_i^τ, r_i^τ) denotes the evidence-rationale pair generated at the i -th iteration; $\mathcal{H}_i^\tau = \{(e_j^\tau, r_j^\tau)\}_{j=0}^i$ represents the history up to step i .
- **Pass Indicator:** $\text{pass}_i^\tau \in \{0, 1\}$, determined by the discriminator \mathcal{V}_τ , indicating if the current chain passes verification.
- **Ensemble and Sampling Parameters:** M is the number of models in the ensemble for SA; K is the number of candidate samples for PC.
- **Correctness Indicator:**
 - For SA: $Z_m \in \{0, 1\}$ indicates if the rationale from the m -th model is correct; the individual accuracy is $p_0^{\text{sa}} = \Pr[Z_m = 1 \mid X^\tau, S^\tau]$.
 - For PC: p_0^{pc} denotes the probability that a single sample yields a correct rationale (conditioned on input and side information).
- **Discriminator Performance:** True Positive Rate (Recall) $\alpha = \Pr[\text{pass} = 1 \mid \text{chain is correct}]$; True Negative Rate (Specificity) $\beta = \Pr[\text{pass} = 0 \mid \text{chain is incorrect}]$.
- **Strategy Coverage Lower Bound:** q_{\min}^τ (Assumption A5), the minimum probability lower bound for generating a correct chain at any step.
- **Iteration Count:** T is the iteration limit in Stage 2 (excluding the seed and fallback step). The total number of attempts is $t = T + 2$ (including seed generation and the final LabelRethink fallback).
- **Effective Noise Rate:**
 - η_1^τ : Error rate of the Stage 1 output.
 - η_2^τ : Error rate of the Stage 2 controller’s output.
 - η_{multi}^τ : Effective noise rate of the final training data (input to Stage 3).
 - η_{e2e}^τ : Noise rate of the E2E supervision signal.
- **Information Measures:** $I(\cdot; \cdot \mid \cdot)$ denotes conditional mutual information, $H(\cdot)$ denotes entropy.

F.2 FUNDAMENTAL ASSUMPTIONS

Our analysis is based on the following assumptions. While often relaxable, they are stated in their strong form for simplicity.

(A1) **Stage 0 Side Information Validity:** The side information S^τ provides meaningful information about the true label Y^τ , i.e., $\exists \delta_S > 0$ such that:

$$I(Y^\tau; S^\tau \mid X^\tau) \geq \delta_S.$$

(A2) **Stage 1 Base Model Accuracy and Correlation:**

- **SA:** For the M base models, the correctness indicators Z_m given input and side information satisfy $\Pr[Z_m = 1 \mid X^\tau, S^\tau] = p_0^{\text{sa}} > 1/2$. The Pearson correlation between any pair is bounded: $\text{Corr}(Z_m, Z_{m'}) \leq \rho \in [0, 1)$.
- **PC:** The base model generates candidate rationales via K independent samplings, with single-sample correctness probability $p_0^{\text{pc}} > 0$.

(A3) **Discriminator Competence:** The aggregator \mathcal{J} in Stage 1 and the discriminator \mathcal{V}_τ in Stage 2 can effectively distinguish correct from incorrect chains, with $\alpha > 1/2$ and $\beta > 1/2$.

(A4) **Conditional Independence of Hidden Reference:** In Stage 2 iteration steps (excluding the seed step), the generated (e_i^τ, r_i^τ) is conditionally independent of the reference rationale r_{ref}^τ , given the current input X^τ and history \mathcal{H}_{i-1}^τ .

(A5) **Strategy Coverage and Minimum Success Rate:** $\exists q_{\min}^\tau > 0$ such that for all $i = 0, \dots, T$:

$$\Pr[\mathcal{G}^\tau(e_i^\tau, r_i^\tau) = 1 \mid X^\tau, \mathcal{H}_{i-1}^\tau] \geq q_{\min}^\tau.$$

This ensures a non-zero chance of generating a correct chain at any step.

(A6) **LabelRethink Fallback:** If all T iterations fail, the LabelRethink module, when injected with r_{ref}^τ and \mathcal{H}_T^τ , produces a correct chain with probability at least $q_{\text{re}}^\tau \geq q_{\min}^\tau$.

(A7) **(Approximate) Independence:** To apply concentration inequalities, we assume:

- For SA: The M models can be partitioned into g groups, with outputs independent across groups (allowing correlation within groups).
- For Stage 2: The outcomes of the t attempts are approximately independent under the discriminator's judgment.

This can be approximately achieved by using diverse model sources and the hidden reference strategy.

1087 F.3 STAGE 1: CONSENSUS AGGREGATION AND NOISE REDUCTION

1089 Stage 1 produces a more reliable reference rationale r_{ref}^{τ} via ensemble (SA) or sampling-selection
1090 (PC), leveraging collective intelligence to reduce the error rate.

1091 **Lemma F.1** (Error Upper Bound for SA Consensus). *Under Assumption (A2), let $S = \sum_{m=1}^M Z_m$
1092 and the majority vote be $\hat{Z} = \mathbf{1}\{S > M/2\}$. Then:*

1094 (a) (Variance-Based Weak Bound) *Generally, the error probability is bounded by:*

$$1096 \Pr[\hat{Z} = 0] \leq \frac{p_0^{\text{sa}}(1 - p_0^{\text{sa}})}{M_{\text{eff}}(p_0^{\text{sa}} - 1/2)^2}, \quad \text{where} \quad M_{\text{eff}} = \frac{M}{1 + (M - 1)\rho}.$$

1099 (b) (Exponential Bound) *Under the group independence assumption (A7) with g groups:*

$$1100 \Pr[\hat{Z} = 0] \leq \exp(-2g(p_0^{\text{sa}} - 1/2)^2).$$

1102 *Proof.* (a) Let $p = p_0^{\text{sa}}$. We have $\mathbb{E}[S] = Mp$. The error event $\{S \leq M/2\}$ is equivalent to
1103 $\mathbb{E}[S] - S \geq M(p - 1/2)$. By Chebyshev's inequality:

$$1105 \Pr(\mathbb{E}[S] - S \geq t) \leq \frac{\text{Var}(S)}{t^2}.$$

1107 Setting $t = M(p - 1/2)$, we bound the variance:

$$\begin{aligned} 1109 \text{Var}(S) &= \sum_m \text{Var}(Z_m) + \sum_{m \neq m'} \text{Cov}(Z_m, Z_{m'}) \\ 1110 &\leq Mp(1 - p) + M(M - 1)\rho p(1 - p) \\ 1111 &= p(1 - p)M[1 + (M - 1)\rho]. \end{aligned}$$

1114 Substitution yields the weak bound. (b) Partition the M models into g groups of size b ($M = gb$).
1115 Define the group average $\bar{Z}_j = \frac{1}{b} \sum_{m \in \text{group } j} Z_m$. The $\{\bar{Z}_j\}_{j=1}^g$ are independent, and $\mathbb{E}[\bar{Z}_j] = p$.
1116 Majority vote failure is equivalent to $\bar{Z} = \frac{1}{g} \sum_{j=1}^g \bar{Z}_j \leq 1/2$. Applying Hoeffding's inequality for
1117 bounded variables gives the exponential bound. \square

1118 **Lemma F.2** (Existence Lower Bound for PC Candidate Selection). *Under Assumptions (A2) and
1119 (A3), the probability that the selected reference rationale in PC is correct is bounded by:*

$$1121 \Pr[r_{\text{ref}}^{\text{pc}} \text{ is correct}] \geq \alpha(1 - (1 - p_0^{\text{pc}})^K).$$

1123 *Proof.* The probability that at least one candidate is correct is $1 - (1 - p_0^{\text{pc}})^K$. Conditioned on this
1124 event, the discriminator selects a correct candidate with probability at least α (true positive rate). The
1125 overall lower bound is the product of these probabilities. \square

1126 **Corollary F.3** (Upper Bound on Stage 1 Effective Noise Rate). *Let $\eta_1^{\tau} = \Pr[r_{\text{ref}}^{\tau} \text{ is incorrect}]$. From
1127 Lemmas F.1 and F.2, we have:*

$$\begin{aligned} 1129 \eta_1^{\text{sa}} &\leq \frac{p_0^{\text{sa}}(1 - p_0^{\text{sa}})}{M_{\text{eff}}(p_0^{\text{sa}} - 1/2)^2} \quad (\text{weak bound}), \\ 1130 \eta_1^{\text{pc}} &\leq 1 - \alpha(1 - (1 - p_0^{\text{pc}})^K). \end{aligned}$$

1133 The bound for SA can be strengthened to the exponential form if the group independence assumption
1134 holds.

1134 **Discussion and Practical Implications**
1135

- 1136 • Stage 1 significantly reduces the supervision noise via aggregation and selection.
- 1137 • For SA, model diversity (low ρ) is crucial. High correlation diminishes the ensemble effect (M_{eff}
1138 decreases). Using diverse models (architectures, pre-training, prompts) is recommended. Group
1139 independence enables exponential error reduction.
- 1140 • For PC, increasing the sample size K and improving the discriminator's TPR α are key to reducing
1141 the error rate.

1143 **F.4 STAGE 2: CONTROLLER PASS PROBABILITY AND ERROR ANALYSIS**
1144

1145 Stage 2 employs controlled iterative generation and verification to find a correct reasoning chain. Its
1146 core is using multiple attempts and discriminator validation to further enhance the probability of
1147 obtaining a correct rationale.

1148 **Design Principle: Hiding the Reference for Information Gain** The hidden reference strategy
1149 (Assumption A4) is central to Stage 2. The following proposition shows that this conditional
1150 independence ensures each iterative step provides new information about Y^τ , preventing the model
1151 from simply parroting the reference rationale and causing information redundancy.

1152 **Proposition F.4** (Information Gain under Conditional Independence). *Under Assumption (A4), for
1153 any $i \geq 1$:*

$$1155 I(Y^\tau; e_i^\tau | X^\tau, \mathcal{H}_{i-1}^\tau, r_{\text{ref}}^\tau) = I(Y^\tau; e_i^\tau | X^\tau, \mathcal{H}_{i-1}^\tau).$$

1156 *Consequently, for the ultimately adopted evidence set $E^\tau = \{e_j^\tau\}_{j=1}^{i^*}$, the cumulative mutual infor-
1157 mation satisfies:*

$$1158 I(Y^\tau; E^\tau | X^\tau) \geq \sum_{j=1}^{i^*} I(Y^\tau; e_j^\tau | X^\tau, \mathcal{H}_{j-1}^\tau).$$

1161 *Proof.* The equality follows directly from the definition of conditional mutual information and (A4).
1162 The inequality results from the chain rule for mutual information and the non-negativity of each
1163 term. \square

1164 This property ensures the benefits of t attempts in Theorem 2 stem from cumulative, incremental
1165 information gain.

1166 Define the probability bounds for a single attempt being a true pass and a false pass:

$$1169 \pi_{\text{TP}}^\tau \geq q_{\min}^\tau \alpha, \quad \pi_{\text{FP}}^\tau \leq (1 - q_{\min}^\tau)(1 - \beta).$$

1171 A single attempt generates a correct chain and gets accepted with probability at least $q_{\min}^\tau \alpha$; it
1172 generates an incorrect chain but gets falsely accepted with probability at most $(1 - q_{\min}^\tau)(1 - \beta)$.

1173 **Theorem F.5** (Controller Pass Probability and False Pass Upper Bound). *Under Assumptions (A3)–
1174 (A6) and the approximate independence assumption (A7), let the total number of attempts be $t = T + 2$.
1175 Then:*

- 1176 1. *The probability of eventually accepting at least one correct chain is lower bounded by:*

$$1178 P_{\text{TP}} = \Pr[\text{Eventually accept a correct chain}] \geq 1 - (1 - \pi_{\text{TP}}^\tau)^t.$$

- 1180 2. *The probability of eventually accepting at least one incorrect chain is upper bounded by:*

$$1181 P_{\text{FP}} = \Pr[\text{Eventually accept an incorrect chain}] \leq 1 - (1 - \pi_{\text{FP}}^\tau)^t.$$

- 1183 3. *The effective noise rate of the controller's output satisfies:*

$$1185 \eta_2^\tau = \Pr[\text{Final output is incorrect} | \text{Accepted}] \leq \frac{P_{\text{FP}}}{P_{\text{TP}} + P_{\text{FP}}}$$

$$1186 \leq \frac{1 - (1 - \pi_{\text{FP}}^\tau)^t}{(1 - (1 - \pi_{\text{TP}}^\tau)^t) + (1 - (1 - \pi_{\text{FP}}^\tau)^t)}.$$

1188 *Proof.* Under approximate independence, the probability of no true pass in t attempts is $\leq (1 - \pi_{\text{TP}}^\tau)^t$,
 1189 so $P_{\text{TP}} \geq 1 - (1 - \pi_{\text{TP}}^\tau)^t$. Similarly, $P_{\text{FP}} \leq 1 - (1 - \pi_{\text{FP}}^\tau)^t$. The noise rate η_2^τ is the conditional
 1190 probability that the first accepted chain is incorrect. Using the bounds for P_{TP} and P_{FP} yields the
 1191 conservative upper bound. \square

1192 **Proposition F.6** (Iteration Complexity for Logarithmic Rate). *If attempts are independent and
 1193 the single-shot success probability is lower bounded by $\pi = \pi_{\text{TP}}^\tau > 0$, then to achieve
 1194 $\Pr[\text{At least one success}] \geq 1 - \epsilon$, the number of attempts t must satisfy:*

$$1196 \quad t \geq \frac{1}{\pi} \log \frac{1}{\epsilon}.$$

1198 *Proof.* From $1 - (1 - \pi)^t \geq 1 - e^{-\pi t} \geq 1 - \epsilon$, solving for t yields the result. \square

1200 Discussion and Practical Implications

- 1202 • P_{TP} approaches 1 exponentially fast with t , while P_{FP} grows slower ($\pi_{\text{FP}}^\tau \ll \pi_{\text{TP}}^\tau$). Thus, an
 1203 accurate discriminator (α, β large) and good strategy coverage (q_{\min}^τ large) enable Stage 2 to output
 1204 rationales with very low error.
- 1205 • The required t scales with $1/\pi$. Improving the single-shot success probability π (via better prompts,
 1206 diversity, or discriminator α) is more efficient than blindly increasing T .

1208 F.5 STAGE 3: GENERALIZATION BOUND UNDER NOISY SUPERVISION

1210 Stage 3 trains the scoring prediction model using the (potentially noisy) rationale-score pairs (r^τ, Y^τ)
 1211 from previous stages. We use the Massart noise model to analyze noisy supervised learning and
 1212 compare the generalization bounds.

1213 **Theorem F.7** (Generalization Upper Bound under Massart Noise ([Massart & Élodie Nédélec \(2006\)](#))).
 1214 Let the hypothesis space \mathcal{H} have complexity measured by d (e.g., VC dimension), the training set
 1215 size be n , and the loss function ℓ be bounded in $[0, 1]$ and Lipschitz. If the effective noise rate of the
 1216 supervision signal is bounded by $\eta < 1/2$ (Massart condition), then for the ERM solution \hat{h} , with
 1217 probability at least $1 - \delta$, the generalization error satisfies:

$$1219 \quad R(\hat{h}) - R(h^*) \leq C_1 \sqrt{\frac{d + \log(1/\delta)}{n}} + C_2 \eta.$$

1221 Here, h^* is the Bayes optimal hypothesis under no noise, and $C_1, C_2 > 0$ are constants related to the
 1222 loss function.

1224 *Proof Sketch.* The bound decomposes into two parts: 1. **Estimation Error (Uniform Convergence):**
 1225 For bounded loss, VC/Rademacher theory gives $\sup_{h \in \mathcal{H}} |R(h) - \hat{R}_n(h)| \leq C_1 \sqrt{(d + \log(1/\delta))/n}$.
 1226 2. **Approximation Error (Noise Bias):** Massart noise introduces a bias term in the risk of the optimal
 1227 hypothesis, linearly related to η , i.e., $|R(h^*) - R_{\text{noisy}}(h_{\text{noisy}}^*)| \leq C_2 \eta$. Combining these two parts
 1228 yields the theorem. See standard results in noisy learning theory for a complete proof. \square

1229 **Multi-Stage vs. End-to-End** Applying Theorem F.7 to the multi-stage method ($\eta = \eta_{\text{multi}}^\tau$) and
 1230 the E2E method ($\eta = \eta_{\text{e2e}}^\tau$), it is clear that if:

$$1232 \quad \eta_{\text{multi}}^\tau < \eta_{\text{e2e}}^\tau,$$

1233 then, for the same n and d , the multi-stage method enjoys a tighter (smaller) generalization error
 1234 upper bound.

1236 F.6 SUFFICIENT CONDITION FOR MULTI-STAGE SUPERIORITY

1238 We now synthesize the results from previous stages to establish a sufficient condition under which
 1239 the multi-stage framework outperforms the E2E baseline.

1240 The final effective noise rate η_{multi}^τ for Stage 3 is a convex combination:

$$1241 \quad \eta_{\text{multi}}^\tau = \Pr[A] \cdot \eta_2^\tau + (1 - \Pr[A]) \cdot \eta_1^\tau,$$

1242 where $\Pr[A]$ is the probability that a Stage 2 candidate is accepted. Consequently,
 1243

$$1244 \min(\eta_1^\tau, \eta_2^\tau) \leq \eta_{\text{multi}}^\tau \leq \max(\eta_1^\tau, \eta_2^\tau).$$

1245 Crucially, if both η_1^τ and η_2^τ are less than η_{e2e}^τ , then $\eta_{\text{multi}}^\tau < \eta_{\text{e2e}}^\tau$ necessarily holds.
 1246

1247 **Theorem F.8** (Sufficient Condition for Multi-Stage Superiority). *Under the assumptions of Lemmas F.1, F.2 and Theorem F.5, if the system parameters $(M, \rho, p_0^{\text{sa}}, K, p_0^{\text{pc}}, \alpha, \beta, T, q_{\min}^\tau)$ satisfy:*

$$1249 \begin{aligned} 1250 \text{(SA)} \quad & \frac{p_0^{\text{sa}}(1 - p_0^{\text{sa}})}{M_{\text{eff}}(p_0^{\text{sa}} - 1/2)^2} < \eta_{\text{e2e}}^{\text{sa}}, \\ 1251 \text{(PC)} \quad & 1 - \alpha (1 - (1 - p_0^{\text{pc}})^K) < \eta_{\text{e2e}}^{\text{pc}}, \\ 1252 \text{(Controller)} \quad & \frac{1 - (1 - \pi_{\text{FP}}^\tau)^t}{(1 - (1 - \pi_{\text{TP}}^\tau)^t) + (1 - (1 - \pi_{\text{FP}}^\tau)^t)} < \eta_{\text{e2e}}^\tau, \quad \tau \in \{\text{sa, pc}\} \end{aligned}$$

1253 where $t = T + 2$, $\pi_{\text{TP}}^\tau \geq q_{\min}^\tau \alpha$, $\pi_{\text{FP}}^\tau \leq (1 - q_{\min}^\tau)(1 - \beta)$, then:
 1254

$$1255 \eta_{\text{multi}}^\tau < \eta_{\text{e2e}}^\tau.$$

1256 Furthermore, by Theorem F.7, the multi-stage method achieves a strictly tighter generalization error
 1257 bound than the E2E method.
 1258

1259 *Proof.* By Corollary F.3, η_1^τ is upper bounded by the left-hand side of the first two inequalities.
 1260 By Theorem F.5, η_2^τ is upper bounded by the left-hand side of the third inequality. The sufficient
 1261 condition ensures $\eta_1^\tau < \eta_{\text{e2e}}^\tau$ and $\eta_2^\tau < \eta_{\text{e2e}}^\tau$. Since η_{multi}^τ is a convex combination of η_1^τ and η_2^τ , it
 1262 must also be less than η_{e2e}^τ . Applying Theorem F.7 concludes the proof. \square
 1263

1264 **Why is this Condition Plausible?** This sufficient condition is not an overly strict requirement
 1265 but a achievable goal through careful design. It holds because the multi-stage framework constructs
 1266 an **error-reduction pipeline**: **Stage 1** reduces noise through **statistical aggregation** (collective
 1267 intelligence). If base models are better than random ($p_0 > 1/2$) and not perfectly correlated ($\rho < 1$),
 1268 aggregation *provably* lowers the error rate below the single-model E2E baseline ($\eta_1^\tau < \eta_{\text{e2e}}^\tau$). **Stage**
 1269 **2** reduces noise through **active exploration and verification** (multiple trials). If the strategy has a
 1270 non-zero chance of being correct ($q_{\min}^\tau > 0$) and the discriminator is better than random ($\alpha, \beta > 1/2$),
 1271 then with sufficient attempts (T large enough), the probability of finding and accepting a correct
 1272 chain approaches 1 exponentially fast, driving the controller’s error rate very low ($\eta_2^\tau < \eta_{\text{e2e}}^\tau$). The
 1273 final noise rate η_{multi}^τ , being an average of these two lower rates, is therefore guaranteed to be lower
 1274 than the E2E baseline. The architecture’s synergistic effect ensures superiority even if no single
 1275 component is perfect.
 1276

1277 F.7 SUMMARY AND EMPIRICAL VALIDATION SUGGESTIONS

1278 This formal analysis indicates that, under the stated assumptions:

- 1279 • **Noise Reduction Mechanism:** Stages 1 and 2 can effectively reduce the supervision noise rate η_{multi}^τ
 1280 observed in the training signal for Stage 3.
- 1281 • **Generalization Advantage:** Within the Massart noise model, a reduced supervision noise rate
 1282 implies a tighter generalization error bound, suggesting that the multi-stage framework may achieve
 1283 better generalization than the E2E approach under such conditions.

1284 G ADDITIONAL EXPERIMENTS AND ANALYSES

1285 In this appendix, we provide additional quantitative and qualitative analyses of *Cosmos-Eval*. We
 1286 describe how each experiment is constructed and report the corresponding results in tables. Un-
 1287 less otherwise noted, all correlations are computed against human 5-point SA/PC labels and 95%
 1288 confidence intervals are obtained via the standard Fisher $r \rightarrow z \rightarrow r$ transform.
 1289

1296 Table 9: **Cross-benchmark results on AIGVE-Bench and LG-VQA.** Pearson correlations between
 1297 automatic evaluators and human scores on two independent evaluation suites.

1298 1299 1300 1301 1302 1303 1304 1305 1306	Model	AIGVE-Bench		LG-VQA	
		Pearson r	95% CI	Pearson r	95% CI
Cosmos-Eval	0.1986	[0.1561, 0.2326]	0.2759	[0.2414, 0.3097]	
VideoPhy-2-AutoEval	0.2089	[0.1706, 0.2466]	0.2750	[0.2404, 0.3088]	
Qwen2.5-VL-7B	0.1033	[0.0063, 0.1425]	0.2013	[0.1656, 0.2366]	

1307 Table 10: **PC-based ranking of T2V generators on AIGVE-Bench.** Mean Cosmos-Eval PC score
 1308 per generator.

1307 1308 1309 1310 1311 1312 1313 1314	Rank	Model	Mean PC Score	#Videos
1	CogVideoX	4.6830	470	
2	Pyramid	4.6311	488	
3	Hunyuan	4.6268	493	
4	Sora	4.6207	493	
5	Genmo	4.5658	486	

G.1 CROSS-BENCHMARK GENERALIZATION

1317 To assess whether Cosmos-Eval overfits to the training benchmarks (VideoPhy/VideoPhy-2), we
 1318 additionally evaluate it on two independent suites: *AIGVE-Bench*(Xiang et al., 2025) and *LG-
 1319 VQA*(Ghosal et al., 2023). Both datasets contain videos generated by multiple T2V models with
 1320 human scores. We directly apply Cosmos-Eval (without any additional fine-tuning) and compare its
 1321 correlation with human scores to that of VideoPhy-2-AutoEval and Qwen2.5-VL-7B.

1322 Table 9 reports Pearson correlations and 95% confidence intervals on both AIGVE-Bench and
 1323 LG-VQA for the three evaluators.

G.2 RANKING T2V GENERATORS ON AIGVE-BENCH

1328 To demonstrate the practical utility of Cosmos-Eval for comparing T2V models, we use it to rank
 1329 several state-of-the-art generators on AIGVE-Bench under the PC (physical commonsense) task. For
 1330 each generator, we compute the mean PC score across all clips associated with that model.

1331 Table 10 reports the resulting ranking. Newer models (e.g., CogVideoX(Yang et al., 2025b), Pyra-
 1332 mid(Jin et al., 2025), Hunyuan(Kong et al., 2025)) achieve higher mean PC scores than earlier systems
 1333 such as Genmo(Li et al., 2025b), and Sora(Liu et al., 2024c) no longer dominates once physical
 1334 plausibility is explicitly emphasized.

G.3 HUMAN EVALUATION OF SA/PC RATIONALES

1339 We conduct a human study to directly assess the perceived quality of SA/PC rationales using a
 1340 custom web interface (Fig. 50). Annotators are shown the video, the caption (for SA), and a candidate
 1341 explanation from one of three models (Cosmos-Eval, GPT-4V(OpenAI et al., 2024), Qwen3-VL-Plus),
 1342 with model identity hidden. For each example, annotators score the explanation along the five rubric
 1343 dimensions defined in our PC and SA reason-quality rubrics (Tables 7 and 8), namely grounding,
 1344 temporal alignment, internal consistency, criteria/decision justification, and either video-quality
 1345 assessment (PC) or coverage & specificity with respect to the caption (SA). Each dimension is scored
 1346 using the three-point anchors $\{0, 0.5, 1\}$, matching the definitions in Tables 7 and 8. In total, the
 1347 study contains 1,500 evaluations across SA and PC.

1348 Table 11 presents average scores for SA rationales across grounding, temporal alignment, consistency,
 1349 alignment justification, and coverage & specificity. Table 12 shows the corresponding results for PC
 1350 rationales.

1350 Table 11: **Human evaluation of SA rationales.** Average scores on grounding, temporal alignment,
 1351 consistency, alignment justification, coverage & specificity, and overall average.

Model	Grounding	Temporal Align.	Consistency	Align. Justif.	Coverage & Spec.	Total Avg.
Cosmos-Eval	0.82	0.57	0.71	0.81	0.87	0.76
GPT-4V	0.64	0.52	0.67	0.61	0.64	0.62
Qwen3-VL-Plus	0.73	0.53	0.69	0.73	0.71	0.69

1356
 1357 Table 12: **Human evaluation of PC rationales.** Average scores on grounding, temporal reasoning,
 1358 consistency, criteria & justification, video-quality awareness, and overall average.

Model	Grounding	Temporal	Consistency	Criteria & Justif.	Video Quality	Total Avg.
Cosmos-Eval	0.79	0.56	0.82	0.85	0.82	0.77
GPT-4V	0.64	0.52	0.56	0.64	0.56	0.58
Qwen3-VL-Plus	0.59	0.51	0.59	0.64	0.63	0.60

1363
 1364 Table 13: **PC rationales scored by Qwen3-VL-Plus.**

Model	Grounding	Temporal	Consistency	Criteria & Justif.	Video Quality	Total Avg.
Cosmos-Eval	0.71	0.79	0.58	0.54	0.69	0.662
GPT-4V	0.60	0.82	0.30	0.26	0.72	0.544
Qwen3-VL-Plus	0.65	0.85	0.67	0.26	0.79	0.564

1369
 1370 Table 14: **SA rationales scored by Qwen3-VL-Plus.**

Model	Grounding	Temporal	Consistency	Align. Justif.	Coverage & Spec.	Total Avg.
Cosmos-Eval	0.76	0.77	0.44	0.44	0.82	0.65
GPT-4V	0.82	0.80	0.47	0.46	0.84	0.678
Qwen3-VL-Plus	0.91	0.83	0.48	0.48	0.91	0.722

1376 Table 15: **PC rationales scored by GPT-4V.**

Model	Grounding	Temporal	Consistency	Criteria & Justif.	Video Quality	Total Avg.
Cosmos-Eval	0.57	0.56	0.62	0.52	0.78	0.61
GPT-4V	0.52	0.54	0.44	0.38	0.88	0.55
Qwen3-VL-Plus	0.57	0.60	0.43	0.35	0.78	0.55

1383 G.4 VLM-JUDGE EVALUATION OF RATIONALES

1385 To further probe explanation quality in a model-agnostic way, we use several strong VLMs as external
 1386 judges. Each judge scores SA/PC rationales from Cosmos-Eval, GPT-4V, and Qwen3-VL-Plus along
 1387 the same rubric dimensions as in the human study, producing scores in 0, 0.5, 1. Below we report
 1388 dimension-wise averages and overall means per judge.

1390 G.4.1 QWEN3-VL-PLUS AS JUDGE

1391 In this setting, we fix Qwen3-VL-Plus as the judge and ask it to assign rubric scores to PC and SA
 1392 rationales produced by Cosmos-Eval, GPT-4V, and Qwen3-VL-Plus itself. Tables 13 and 14 report
 1393 the dimension-wise averages and overall mean scores.

1395 G.4.2 GPT-4V AS JUDGE

1397 Here we use GPT-4V as the judge and follow the same protocol: given a video, caption (for SA),
 1398 and a candidate rationale from each model, GPT-4V assigns rubric scores to SA/PC explanations.
 1399 Tables 15 and 16 summarize the resulting averages.

1401 G.4.3 GEMINI-2.5-PRO AS JUDGE

1403 Finally, we repeat the same evaluation protocol with Gemini-2.5-Pro as the judge. Tables 17 and 18
 report the average rubric scores for PC and SA rationales, respectively.

1404 **Table 16: SA rationales scored by GPT-4V.**
1405

Model	Grounding	Temporal	Consistency	Align. Justif.	Coverage & Spec.	Total Avg.
Cosmos-Eval	0.67	0.46	0.74	0.71	0.64	0.64
GPT-4V	0.66	0.50	0.73	0.72	0.64	0.65
Qwen3-VL-Plus	0.73	0.55	0.73	0.70	0.69	0.68

1410 **Table 17: PC rationales scored by Gemini-2.5-Pro.**
1411

Model	Grounding	Temporal	Consistency	Criteria & Justif.	Video Quality	Total Avg.
Cosmos-Eval	0.51	0.50	0.41	0.34	0.01	0.36
GPT-4V	0.49	0.46	0.15	0.10	0.00	0.24
Qwen3-VL-Plus	0.50	0.52	0.21	0.16	0.04	0.29

1416 **Table 18: SA rationales scored by Gemini-2.5-Pro.**
1417

Model	Grounding	Temporal	Consistency	Align. Justif.	Coverage & Spec.	Total Avg.
Cosmos-Eval	0.47	0.46	0.25	0.29	0.63	0.42
GPT-4V	0.63	0.61	0.35	0.30	0.64	0.51
Qwen3-VL-Plus	0.70	0.67	0.29	0.27	0.77	0.54

1423 **Table 19: Score correlations on VideoPhy-2 (50 clips).** Direct SA/PC scoring by frontier VLMs vs.
1424 Cosmos-Eval.
1425

Model	SA Pearson r	SA Spearman ρ	PC Pearson r	PC Spearman ρ
Cosmos-Eval	0.4325	0.4255	0.3046	0.2984
Qwen3-VL-Plus	0.5367	0.5243	0.1418	0.1394
GPT-4V	0.5917	0.5753	0.1429	0.1333

1431 **Table 20: Rationale similarity on VideoPhy-2 (50 clips).** BLEU-4 and BERTScore-F1 (in %) for
1432 SA/PC explanations vs. reference rationales.
1433

Model	SA BLEU-4	SA BERTScore-F1	PC BLEU-4	PC BERTScore-F1
Cosmos-Eval	32.94	80.08	26.19	77.79
Qwen3-VL-Plus	8.75	72.09	2.65	68.06
GPT-4V	5.08	71.16	2.57	70.14

1438 **Table 21: Summac scores for SA rationales.**
1439

Metric	Cosmos-Eval	Qwen2.5-VL-7B	InternVL-8B	InternVL-9B	InternVL-14B	Cosmos-Reason1	VideoLLaMA3-7B
summac	26.62	21.50	23.92	24.22	23.91	21.23	22.56

1444

FRONTIER VLM BASELINES ON VIDEOPHY-2

14451446 To compare Cosmos-Eval against frontier VLMs used with direct prompting, we sample 50 VideoPhy-
1447 2 test clips with human SA/PC labels. We prompt GPT-4V and Qwen3-VL-Plus to directly output
1448 5-point SA/PC scores and compute correlations with human labels.
14491450 Table 19 reports Pearson and Spearman correlations for scores. Table 20 reports BLEU-4 and
1451 BERTScore-F1 for SA/PC rationales against reference explanations.
14521453

FACTUAL CONSISTENCY METRICS FOR RATIONALES

14541455 Beyond surface-level text similarity metrics, we also consider a factual/consistency-oriented metric
1456 (SummAc(Laban et al., 2022)) to assess alignment between generated and reference explanations.
1457 Table 21 reports Summac scores for SA rationales, and Table 22 for PC rationales, comparing
1458 Cosmos-Eval to several baselines.
1459

1458
1459 **Table 22: Summac scores for PC rationales.**
1460

Metric	Cosmos-Eval	Qwen2.5-VL-7B	InternVL-8B	InternVL-9B	InternVL-14B	Cosmos-Reason1	VideoLLaMA3-7B
summac	23.32	22.69	23.07	22.86	23.16	22.25	23.20

1462
1463 **Table 23: SA correlations with uncertainty.** Pearson r (95% CI), two-sided p -value, and Δr vs.
1464 Cosmos-Eval.

Model	r [95% CI]	p -value	Δr vs. Cosmos-Eval
Cosmos-Eval	0.4643 [0.4376, 0.4904]	2.50E-181	—
VideoPhy-2-AutoEval	0.4327 [0.4049, 0.4596]	5.12E-155	+0.0316
Qwen2.5-VL-7B	0.3808 [0.3517, 0.4092]	1.02E-117	+0.0835
VideoLLaMA3-7B	0.2769 [0.2456, 0.3077]	7.21E-61	+0.1874
InternVL-8B	0.4143 [0.3861, 0.4418]	4.70E-141	+0.0500
InternVL-9B	0.3827 [0.3536, 0.4110]	6.34E-119	+0.0816
InternVL-14B	0.3420 [0.3120, 0.3714]	7.17E-94	+0.1223
Cosmos-Reason1	0.3662 [0.3366, 0.3952]	3.62E-107	+0.0981
VideoLLaMA3-7B (variant)	0.2333 [0.2034, 0.2636]	7.78E-44	+0.2310

1465
1466 **Table 24: PC correlations with uncertainty.** Pearson r (95% CI), two-sided p -value, and Δr vs.
1467 Cosmos-Eval.

Model	r [95% CI]	p -value	Δr vs. Cosmos-Eval
Cosmos-Eval	0.3641 [0.3346, 0.3929]	4.83E-107	—
VideoPhy-2-AutoEval	0.3646 [0.3351, 0.3934]	2.55E-107	-0.0005
Qwen2.5-VL-7B	0.0840 [0.0512, 0.1180]	7.59E-07	+0.2801
VideoLLaMA3-7B	0.0640 [0.0310, 0.0980]	1.67E-04	+0.3001
InternVL-8B	0.1665 [0.1280, 0.1935]	3.79E-21	+0.1976
InternVL-9B	0.1304 [0.0973, 0.1634]	2.26E-14	+0.2337
InternVL-14B	0.1956 [0.1631, 0.2278]	1.18E-30	+0.1685
Cosmos-Reason1	0.2356 [0.2030, 0.2665]	7.78E-44	+0.1285
VideoLLaMA3-7B (variant)	0.2075 [0.1795, 0.2354]	2.30E-43	+0.1566

1485
1486 **Table 25: GPT-4o as external scorer.** Pearson correlations between GPT-4o-implied scores and
1487 Cosmos-Eval scores under two Stage-2 judges.

Setting	N	Pearson r	95% CI
PC, 72B judge (Qwen2.5-VL-72B)	187	0.9131	[0.8857, 0.9342]
PC, 30B judge (Qwen3-VL-30B)	187	0.8239	[0.7716, 0.8651]
SA, 72B judge (Qwen2.5-VL-72B)	178	0.8894	[0.8541, 0.9166]
SA, 30B judge (Qwen3-VL-30B)	184	0.8636	[0.8216, 0.8963]

1494

G.7 UNCERTAINTY ESTIMATES FOR MAIN SA/PC RESULTS

1495
1496 For the main SA/PC score correlations, we also report uncertainty estimates and effect sizes. For
1497 each model, we compute: (i) Pearson correlation r with human scores, (ii) 95% confidence interval,
1498 (iii) two-sided p -value for $H_0 : r = 0$, and (iv) Δr relative to Cosmos-Eval.

1499 Tables 23 and 24 summarize SA and PC statistics, respectively.

1501

G.8 SENSITIVITY TO STAGE-2 JUDGE AND EXTERNAL SCORERS

1503
1504 To study judge-choice sensitivity and potential circularity, we vary the Stage-2 judge (Qwen2.5-
1505 VL-72B vs. Qwen3-VL-30B) while keeping the rest of the pipeline fixed. For each setting, we ask
1506 independent external LLMs (GPT-4o and DeepSeek) to read Cosmos-Eval rationales and assign
1507 SA/PC scores, then compute the correlation between these external scores and the original Cosmos-
1508 Eval scores.1509
1510 Table 25 reports Pearson correlations when GPT-4o is used as the external scorer, and Table 26
1511 reports the same when DeepSeek is used. In all cases, we observe high agreement across judge
choices and external scorers, and the larger Stage-2 judge (Qwen2.5-VL-72B) consistently yields
higher correlations than the 30B variant, indicating that Stage 2 benefits from stronger VLM judges;
accordingly, we adopt Qwen2.5-VL-72B as the default Stage-2 judge in our main experiments.

1512 Table 26: **DeepSeek as external scorer.** Pearson correlations between DeepSeek-implied scores and
 1513 Cosmos-Eval scores under two Stage-2 judges.

Setting	<i>N</i>	Pearson <i>r</i>	95% CI
PC, 72B judge (Qwen2.5-VL-72B)	187	0.9131	[0.8857, 0.9342]
PC, 30B judge (Qwen3-VL-30B)	187	0.8487	[0.8031, 0.8845]
SA, 72B judge (Qwen2.5-VL-72B)	178	0.8894	[0.8541, 0.9166]
SA, 30B judge (Qwen3-VL-30B)	184	0.8649	[0.8233, 0.8973]

1520 Table 27: **Stage-wise cost of teacher pipeline for PC (200 samples).**

Step	GPU Count	Inference Time	Sample Count	Avg. Time / Sample
Stage 0 (pre-processing)	1	2m12s	200	0.66 s
Stage 1: Qwen Gen (run 1)	1	18m43s	200	19.38 s
Stage 1: Qwen Gen (run 2)	2	18m37s	200	29.25 s
Stage 2 (reasoning ctrl.)	2	188m48s	200	56.64 s

1528 Table 28: **Stage-wise cost of teacher pipeline for SA (200 samples).**

Step	GPU Count	Inference Time	Sample Count	Avg. Time / Sample
Stage 0 (pre-processing)	1	2m14s	200	0.67 s
Stage 1: Qwen Gen	4	54m43s	200	16.31 s
Stage 1: Tarsier Gen	4	18m19s	200	20.78 s
Stage 1: Qwen3 merge	2	17m39s	200	5.14 s
Stage 2 (complex reasoning ctrl.)	4	340m12s	200	102.06 s

1536 Table 29: **Score-only inference cost for PC/SA scores (200 samples).**

Step	GPU Count	Inference Time	Sample Count	Avg. Time / Sample	GPU-Hours
VideoPhy-2-AutoEval-PC	1	2m12s	200	0.66 s	0.0368
VideoPhy-2-AutoEval-SA	1	2m14s	200	0.67 s	0.0372
Cosmos-Eval (PC-score)	1	3m42s	200	1.11 s	0.0618
Cosmos-Eval (SA-score)	1	3m42s	200	1.11 s	0.0619
Qwen2.5-VL-7B (PC-score)	1	4m20s	200	1.30 s	0.0722
Qwen2.5-VL-7B (SA-score)	1	4m20s	200	1.30 s	0.0724

1545 Table 30: **Rationale-generation inference cost for PC/SA reasons (200 samples).**

Step	GPU Count	Inference Time	Sample Count	Avg. Time / Sample	GPU-Hours
Cosmos-Eval (PC-reason)	1	19m34s	200	5.87 s	0.3272
Cosmos-Eval (SA-reason)	1	51m49s	200	15.54 s	0.8637
Qwen2.5-VL-7B (PC-reason)	1	14m14s	200	4.27 s	0.2042
Qwen2.5-VL-7B (SA-reason)	1	13m34s	200	4.07 s	0.2219

1553 G.9 COMPUTATIONAL COST AND EFFICIENCY

1555 We report detailed computational costs for (i) the multi-stage teacher pipeline (Stages 0–2) used during
 1556 training and (ii) the distilled student evaluator (Stage 3) used during inference. All measurements are
 1557 collected on 200-sample subsets.

1558 Tables 27 and 28 list stage-wise costs for PC and SA teacher pipelines, respectively, while Tables 29
 1559 and 30 compare score-only and rationale-generation inference costs for Cosmos-Eval, VideoPhy-2-
 1560 AutoEval, and Qwen2.5-VL-7B. From these numbers, the total per-sample cost of the teacher pipeline
 1561 (Stages 0–2) is about 106 s for PC and 145 s for SA, whereas the distilled evaluator (Stage 3) needs
 1562 only ≈ 1.1 s per sample for scores alone and ≈ 5.9 s (PC) / 15.5 s (SA) for scores plus rationales.
 1563 When we compare the teacher pipeline with the distilled evaluator under the score + rationale setting,
 1564 this translates to roughly 9–18 \times speedups. In deployment, users only run the distilled evaluator,
 1565 whose score-only cost is close to that of a single 7B VLM call, while additionally providing calibrated
 1566 SA/PC scores and physics-grounded rationales rather than scores alone.

1566 Table 31: **Average PC/SA scores before and after synthetic degradations on VideoPhy-2 clips.**
 1567 Each row averages 100 samples per distortion type.

Distortion Type	Before (PC)	After (PC)	PC Δ (B-A)	Before (SA)	After (SA)	SA Δ (B-A)	Count
Noise	4.56	4.28	0.28	4.03	3.72	0.31	100
Occlusion	4.59	4.21	0.38	4.02	3.82	0.20	100
Blur	4.64	4.32	0.32	4.00	3.72	0.28	100
Compression	4.67	4.45	0.22	4.02	3.77	0.25	100
Color	4.67	4.37	0.30	4.01	3.77	0.24	100

1574 Table 32: **Hyperparameter ranges for synthetic degradations.** PC and SA tasks share the same
 1575 ranges with different random seeds (PC: 42, SA: 43).

Distortion Type	Parameter	Range	Description
Noise	strength	10–40	Noise intensity level
Noise	temporal	5–15	Temporal noise variation
Blur	sigma (luma)	1.0–4.0	Gaussian blur radius
Blur	chroma_radius	0.5–2.0	Chroma blur radius
Compression	CRF	35–45	Constant rate factor (higher = more compression)
Compression	bitrate	100–300 kbps	Target video bitrate
Occlusion	num_boxes	1–3	Number of black occlusion boxes
Occlusion	box_size	50–150 px	Width and height of each box
Occlusion	position (x,y)	0–500 px	Random position within frame
Color	saturation	0.3–1.5	Color saturation multiplier
Color	contrast	0.5–1.3	Contrast adjustment factor
Color	brightness	−0.2–0.2	Brightness offset
Color	hue	−30°–30°	Hue rotation angle

1589 Table 33: **Long-horizon evaluation on LongCat-Video (30 prompts, \approx 33s per video).** Segment-
 1590 wise average ranges (across 11 segments) and full-clip averages for SA/PC.

Model	SA segment range	PC segment range	Full SA avg.	Full PC avg.
Cosmos-Eval	3.03–3.13	3.97–4.13	3.30	4.00
VideoPhy-2-AutoEval	2.73–3.00	3.83–4.07	2.97	3.50

1597 G.10 ROBUSTNESS TO SYNTHETIC DEGRADATIONS

1599 We examine robustness to synthetic noise by starting from clean VideoPhy-2 clips and applying
 1600 controlled degradations (noise, occlusion, blur, compression, color shifts). For each distortion type,
 1601 we compute average PC/SA scores before and after degradation over 100 clips.

1602 Table 31 reports the mean scores for each distortion type and the corresponding drops in PC/SA.
 1603 Table 32 lists the hyperparameter ranges used to generate each distortion; PC and SA tasks share
 1604 the same parameter ranges with different random seeds. Overall, we observe moderate but not
 1605 catastrophic degradation: PC scores are most sensitive to occlusion (largest drop of 0.38), while SA
 1606 scores are most affected by additive noise (largest drop of 0.31), and both tasks are comparatively
 1607 robust to blur and compression. Note that these experiments are conducted on T2V-generated clips
 1608 with synthetic distortions rather than real-world, in-the-wild video artifacts, so extending Cosmos-Eval
 1609 to broader real-world noise conditions remains an interesting direction for future work.

1611 G.11 LONG-HORIZON EVALUATION ON LONGCAT-VIDEO

1613 To probe long-horizon behavior, we evaluate Cosmos-Eval and VideoPhy-2-AutoEval on videos of
 1614 length \approx 33 seconds generated from 30 prompts using a LongCat-style T2V setup. Each video is
 1615 evaluated (i) as a full clip, and (ii) as 11 consecutive non-overlapping 3-second segments. We report
 1616 mean SA/PC scores across segments and for the full video.

1618 Table 33 summarizes segment-wise ranges and full-clip averages. Both evaluators exhibit stable
 1619 SA/PC scores across time, and Cosmos-Eval behaves comparably to VideoPhy-2-AutoEval on long
 multi-step sequences.

1620 **H REPRODUCIBILITY STATEMENT**
16211622 All information needed to replicate our results is provided in Appx. B (Figs. 3–5, Alg. 1, Table 6) and
1623 the main text (Eqs. 4, 6). All datasets used are publicly available and can be downloaded from their
1624 official websites (*VideoPhy* and *VideoPhy-2*; see (Bansal et al., 2025a;b)). We detail the complete
1625 prompt flow and provide all prompts in Appx. J. Model versions and full decoding hyperparameters
1626 (temperature, top- p , max tokens) are specified. Because inference relies on sampling, we do not
1627 fix random seeds; minor run-to-run variance is expected, but the stated configurations suffice for
1628 independent replication of the main results. Upon acceptance, we will publicly release all code,
1629 scripts, and model weights to facilitate exact reproduction.
16301631 **I THE USE OF LARGE LANGUAGE MODELS (LLMs)**
16321633 We used large language models only for light editorial assistance during manuscript preparation
1634 (grammar and wording refinement, minor style/formatting suggestions). No LLMs were used for
1635 research ideation, dataset curation, modeling, experiment design, analysis, or drafting substantive
1636 sections.
16371638 **J PROMPT TEMPLATES**
16391640 This section briefly documents the prompt flow used in Stages 1–2; figures referenced below are
1641 already included in the paper.
16421643 • **SA, Stage 1.** From the *rationale prompt* (Fig. 25) to the *consensus prompt* (Fig. 26), which aggregates
1644 two rationales into the SA reference $r_{\text{ref}}^{\text{sa}}$.
1645 • **PC, Stage 1.** From the *candidate-generation prompt* (Fig. 27) to the *explanation-selection prompt*
1646 used by the judge (Fig. 28) to obtain $r_{\text{ref}}^{\text{pc}}$.
1647 • **SA, Stage 2.** From the *seed-ref prompt* (Fig. 29) to the *assessment prompt* (Fig. 36) that produces a
1648 concise evidence-based justification.
1649 • **PC, Stage 2.** From the *seed-ref prompt* (Fig. 37) to the *assessment prompt* (Fig. 44) under the PC
1650 rubric.
1651 • **Unified CoT narration.** The accepted structured analysis from Stage 2 is converted into a natural,
1652 first-person narration using the *NaturalReasoning* prompt (Fig. 45).
1653 • **Ablations (SA/PC).** From the *DeepSeek-R1 remapping prompt* (Fig. 46) to the *Qwen-VL-Max*
1654 *reason-evaluation prompt* (Fig. 49).
16551656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673

1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689




Figure 3: **Stage 1 (SA) reason generation (ensemble \Rightarrow consensus).** An ensemble $\{\mathcal{M}_m\}_{m=1}^M$ produces one reason each, forming the pool $\mathcal{R}_{\text{pool}}^{\text{sa}}$ (Eq. 3); an aggregator LLM then extracts shared content to yield the reference reason $r_{\text{ref}}^{\text{sa}}$ (Eq. 4), which seeds Stage 2.

1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706




Figure 4: **Stage 1 (PC) reason generation (sampling \Rightarrow selection).** The base VLM $\mathcal{M}_{\text{base}}$ samples K candidate reasons to form the pool $\mathcal{R}_{\text{pool}}^{\text{pc}}$ (Eq. 5); an VLM judge \mathcal{J}_{pc} then selects the reference rationale $r_{\text{ref}}^{\text{pc}}$ (Eq. 6), which seeds Stage 2.

1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720




Figure 5: **Stage 2 (reason-augmented CoT).** Starting from the reference reason r_{ref}^{τ} (from Stage 1), a judge-verified controller iteratively explores, verifies, and corrects without exposing the reference mid-trajectory; each candidate (e_i^{τ}, r_i^{τ}) is checked by \mathcal{V}_{τ} for pass or fail (Eqs. equation 9, equation 12). The controller uses the strategy set \mathcal{C} (Backtracking, Exploring New Paths, Verification, Correction); if none pass, LabelRethink re-injects the reference (Eq. equation 13), and the accepted history is reformatted into $(\hat{e}^{\tau}, \hat{r}^{\tau})$ (Eq. equation 16).




Figure 6: **Score distributions for SA and PC.** Left to right: (i) overall (train+test), (ii) *VideoPhy* subset (SA/PC scores obtained via *VideoPhy-2* AutoEval), (iii) *VideoPhy-2* subset, and the final (iv) training and (v) test splits. SA is skewed toward higher scores (4–5), whereas PC concentrates on 3–4 with fewer 5’s and more 2’s than SA. On the *VideoPhy* portion, SA is sharply peaked at 4 (almost no 5’s), while PC is roughly balanced across 2–4; *VideoPhy-2* shows a broader SA spread (nontrivial 1/5 tails) and a PC peak at 3 with a secondary mode at 4. Train/test distributions are similar, with the test split slightly flatter. These imbalances motivate reporting κ alongside accuracy/correlation and using stratified sampling in SFT.





Figure 7: **Case 1 — Physical Commonsense (PC).** **Input caption:** *N/A (no caption provided).* **Answer (Cosmos-Eval, verbatim):** The video demonstrates a red ball hovering mid-air without any visible support, defying gravitational expectations. The physical commonsense is poor because the ball’s static suspension violates basic principles of force and motion, despite maintaining a realistic appearance otherwise. This justifies a pc_score of 2 due to the significant inconsistency with gravitational effects while other visual elements remain accurate. **PC score: 2.**

1782
1783
1784 **Algorithm 1:** Stage-2 Reference-Seeded, Judge-Verified Controller for task τ
1785 **Input:** x^τ ; prompts $\mathbf{P}_{\text{seed-ref}}^\tau, \{\mathbf{P}_c^\tau\}_{c \in \mathcal{C}}, \mathbf{P}_{\text{rethink}}^\tau$; judge prompt \mathbf{U}^τ ; reference r_{ref}^τ ; budget N
1786 **Output:** $(\hat{e}^\tau, \hat{r}^\tau)$ or \emptyset
1787 $\mathcal{H}^\tau \leftarrow \emptyset$; $i^* \leftarrow \text{nil}$;
1788 $\text{Avail} \leftarrow \mathcal{C}$; $T \leftarrow \min(N, |\mathcal{C}|)$;
1789 $(e_0^\tau, r_0^\tau) \leftarrow \mathcal{M}(\mathbf{P}_{\text{seed-ref}}^\tau, x^\tau, r_{\text{ref}}^\tau; \text{Reason})$;
1790 $\mathcal{H}^\tau \leftarrow \mathcal{H}^\tau \cup \{(e_0^\tau, r_0^\tau)\}$;
1791 $\text{pass} \leftarrow \mathcal{V}_\tau(r_0^\tau, r_{\text{ref}}^\tau; \mathbf{U}^\tau)$;
1792 **if** $\text{pass} = 1$ **then**
1793 $i^* \leftarrow 0$;
1794 **else**
1795 **for** $i \leftarrow 1$ **to** T **do**
1796 pick $c_i \in \text{Avail}$ uniformly; $\text{Avail} \leftarrow \text{Avail} \setminus \{c_i\}$;
1797 $(e_i^\tau, r_i^\tau) \leftarrow \mathcal{M}(\mathbf{P}_{c_i}^\tau, x^\tau, \mathcal{H}^\tau; c_i)$;
1798 $\mathcal{H}^\tau \leftarrow \mathcal{H}^\tau \cup \{(e_i^\tau, r_i^\tau)\}$;
1799 $\text{pass} \leftarrow \mathcal{V}_\tau(r_i^\tau, r_{\text{ref}}^\tau; \mathbf{U}^\tau)$;
1800 **if** $\text{pass} = 1$ **then**
1801 $i^* \leftarrow i$;
1802 **break**
1803 **if** $\text{pass} = 0$ **then**
1804 $(e_{T+1}^\tau, r_{T+1}^\tau) \leftarrow \mathcal{M}(\mathbf{P}_{\text{rethink}}^\tau, x^\tau, r_{\text{ref}}^\tau, \mathcal{H}^\tau; \text{LabelRethink})$;
1805 $\mathcal{H}^\tau \leftarrow \mathcal{H}^\tau \cup \{(e_{T+1}^\tau, r_{T+1}^\tau)\}$;
1806 $\text{pass} \leftarrow \mathcal{V}_\tau(r_{T+1}^\tau, r_{\text{ref}}^\tau; \mathbf{U}^\tau)$;
1807 **if** $\text{pass} = 1$ **then**
1808 $i^* \leftarrow T+1$;
1809 **if** $\text{pass} = 1$ **then**
1810 $\hat{e}^\tau \leftarrow \text{PostChain}(\{(e_j^\tau, r_j^\tau)\}_{j=0}^{i^*}; \text{SynthesizeChain})$;
1811 $\hat{r}^\tau \leftarrow \text{PostAnswer}(\hat{e}^\tau, r_{\text{ref}}^\tau; \text{Reformat})$;
1812 **return** $(\hat{e}^\tau, \hat{r}^\tau)$;
1813 **else**
1814 **return** \emptyset ;
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827



1828 **Figure 8: Case 2 — Semantic Alignment (SA).** **Generation caption:** A yellow cube rotates
1829 *counterclockwise on the table while a purple cone remains still.* **Answer (Cosmos-Eval, verbatim):**
1830 The video demonstrates the yellow cube rotating while the purple cone remains stationary, with the
1831 former exhibiting a clockwise rotation contrary to the caption’s description of counterclockwise
1832 movement. The alignment is moderate because while the scene and object placement are accurate,
1833 the direction of rotation significantly deviates from the stated caption. This justifies an `sa_score` of 3
1834 due to the primary focus on object interaction with only a minor discrepancy in rotational direction.
1835 **SA score (assigned):** 2.


1836

1837

1838

1839

1840




1841

1842

1843

1844

1845



1846

1847

1848

1849

1850

Figure 9: **Case 3 — Physical Commonsense (PC).** **Input caption:** *N/A (no caption provided).* **Answer (Cosmos-Eval, verbatim):** The video demonstrates a ball rolling on a checkerboard surface with inconsistent motion, exhibiting erratic bouncing and lack of frictional decay, which violates principles of Newtonian mechanics. The physical commonsense is poor because the ball’s behavior cannot be explained by external forces alone, suggesting either an artificial manipulation or a misunderstanding of physical laws. This justifies a **pc_score: 2** due to significant deviations from expected physical interactions and dynamics. **PC score: 2.**

1851

1852


1853

1854

1855

1856

1857



1858

1859

1860

1861


1862

1863

1864

1865

1866



1867

1868

1869

1870

1871

1872

Figure 10: **Case 4 — Semantic Alignment (SA).** **Generation caption:** *A man kicks a soccer ball; it strikes the goalpost and rebounds.* **Answer (Cosmos-Eval, verbatim):** The video shows a man observing two soccer balls approaching the goalpost, whereas the caption describes a single ball being kicked towards the goalpost. The alignment is moderate because although the setting accurately represents a soccer field, the presence of multiple balls and the absence of a kicking action introduce significant discrepancies. This justifies a **sa_score: 3** due to the partial adherence to the scene description despite notable differences in the number of balls and the absence of the critical kicking action. **SA score: 3.**

1873

1874

1875

1876

1877

1878

1879

1880

1881

1882

1883

1884

1885



Figure 11: Example of Physical Commonsense Task 1




Figure 12: Example of Physical Commonsense Task 2




Figure 13: Example of Physical Commonsense Task 3

2052

2053

2054

2055

2056

2057

2058

2059

2060

2061

2062

2063

2064

2065

2066

2067

2068

2069

2070

2071

2072

2073

2074

2075

2076

2077

2078

2079

2080

2081

2082

2083

2084

2085

2086

2087

2088

2089

2090

2091

2092

2093

2094

2095

2096

2097

2098

2099

2100

2101

2102

2103

2104

2105




Figure 14: Example of Physical Commonsense Task 4



Figure 15: Example of Physical Commonsense Task 5

2160

2161

2162

2163

2164

2165

2166

2167

2168

2169

2170

2171

2172

2173

2174

2175

2176

2177

2178

2179

2180

2181

2182

2183

2184

2185

2186

2187

2188

2189

2190

2191

2192

2193

2194

2195

2196

2197

2198

2199

2200

2201

2202

2203

2204

2205

2206

2207

2208

2209

2210

2211

2212

2213




Figure 16: Example of Semantic Alignment Task 1

2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267




Figure 17: Example of Semantic Alignment Task 2

2268

2269

2270

2271

2272

2273

2274

2275

2276

2277

2278

2279

2280

2281

2282

2283

2284

2285

2286

2287

2288

2289

2290

2291

2292

2293

2294

2295

2296

2297

2298

2299

2300

2301

2302

2303

2304

2305

2306

2307

2308

2309

2310

2311

2312

2313

2314

2315

2316

2317

2318

2319

2320

2321





Figure 18: Example of Semantic Alignment Task 3

2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340



2366
2367
2368
2369
2370
2371
2372
2373
2374
2375

2376

2377

2378

2379

2380

2381

2382

2383

2384

2385

2386

2387

2388

2389

2390

2391

2392

2393

2394

2395

2396

2397

2398

2399

2400

2401

2402

2403

2404

2405

2406

2407

2408

2409

2410

2411

2412

2413

2414

2415

2416

2417

2418

2419

2420

2421

2422

2423

2424

2425

2426

2427

2428

2429




Figure 20: Example of Semantic Alignment Task 5

2430

2431

2432

2433

2434

2435

2436

2437

2438

2439

2440

2441

2442

2443

2444

2445

2446

2447

2448

2449

2450

2451

2452

2453

2454

2455

2456

2457

2458

2459

2460

2461

2462

2463

2464

2465

2466

2467

2468

2469

2470

2471

2472

2473

2474

2475

2476

2477

2478

2479

2480

2481

2482

2483




Figure 21: Example of Physical Commonsense Task with CoT (Case 1)

2484

2485

2486

2487

2488

2489

2490

2491

2492

2493

2494

2495

2496

2497

2498

2499

2500

2501

2502

2503

2504

2505

2506

2507

2508

2509

2510

2511

2512

2513

2514

2515

2516

2517

2518

2519

2520

2521

2522

2523

2524

2525

2526

2527

2528

2529

2530

2531

2532

2533

2534

2535

2536

2537




Figure 22: Example of Physical Commonsense Task with CoT (Case 2)

2538

2539

2540

2541

2542

2543

2544

2545

2546

2547

2548

2549

2550

2551

2552

2553

2554

2555

2556

2557

2558

2559

2560

2561

2562

2563

2564

2565

2566

2567

2568

2569

2570

2571

2572

2573

2574

2575

2576

2577

2578

2579

2580

2581

2582

2583

2584

2585

2586

2587

2588

2589

2590

2591




Figure 23: Example of Semantic Alignment Task with CoT (Case 1)

2592

2593

2594

2595

2596

2597

2598

2599

2600

2601

2602

2603

2604

2605

2606

2607

2608

2609

2610

2611

2612

2613

2614

2615

2616

2617

2618

2619

2620

2621

2622

2623

2624

2625

2626

2627

2628

2629

2630

2631

2632

2633

2634

2635

2636

2637

2638

2639

2640


2641

2642

2643

2644

2645



Caption: Someone pulls apart a tightly wound rubber band that separates with audible snap.

Cosmos-Eval

<think>\nOkay, let me watch this video carefully... I can see a pair of hands holding a red, coiled object that looks like a toy or maybe a stress ball. The caption mentions a 'rubber band', and at first glance, this red object seems flexible, like a rubber band might be. But wait, let me look closer... Actually, it's more than just a band; it's coiled tightly and has a distinct shape. Hmm, this seems like it could be some kind of toy rather than just a simple rubber band. I notice the hands are pulling it apart, and it starts to uncoil. The caption says it 'separates with an audible snap', but I don't hear any snapping sound here. Oh wait, maybe it's just not captured in the video? Or perhaps the action is subtle. The hands are moving, and the object is stretching and twisting. This looks like a slow, deliberate motion, not the quick snap implied by the caption. But then again, the scene is quite plain, almost like a studio setup for focusing on the object. The lighting is even, and the background is neutral, which fits well with highlighting the red toy against it. Yes, that confirms the scene is simple and focused. The overall action is continuous, showing the object being stretched and manipulated without any jumps or cuts. This looks like it's showing the transformation of the toy over time, which aligns with the caption suggesting a transformation. So, putting it all together, the video partially matches the caption. It shows the hands manipulating the toy, but it doesn't capture the exact nature of the transformation or the separation as described. The scene is simple, and the motion is smooth, which supports a good alignment. I think the score should reflect this as a moderate alignment with some discrepancies, like the lack of snapping sound and the focus on transformation rather than separation. Overall, it's a decent match, but there are still noticeable differences that prevent it from being perfect.\n</think>\n\n<answer>\n\nThe video shows hands manipulating a coiled, red flexible object, while the caption describes a 'rubber band' that separates with an audible snap. The alignment is moderate because the video demonstrates continuous manipulation without revealing the snapping sound or the separation process, indicating a transformation rather than a simple tear. This justifies a sa_score of 3 due to the depiction of a gradual action and the absence of the described auditory and separation events.\n</answer>

Figure 24: Example of Semantic Alignment Task with CoT (Case 2)

2646
 2647
 2648
 2649
 2650
 2651
 2652
 2653
 2654
 2655
 2656
 2657
 2658
 2659
 2660
 2661
 2662

Title :Stage 1 (SA) rationale prompt — ensemble

2663 You are a video-text alignment expert. Analyze the alignment between video content and text descriptions.
 2664 Your response must be a valid JSON object with exactly these two fields:
 2665 {
 2666 "score_explanation": "Based on X-point scoring basis: [explanation]",
 2666 "mismatches": ["Caption requires ... but video shows ...", "Caption specifies ... but video shows ..."]
 2667 }
 2668
 2669 Scoring Basis:
 2670 sa points: {5=Perfect alignment | 4=Minor deviations | 3=Partial match | 2=Key omissions | 1=Completely unrelated}
 2671 Analysis Dimensions:
 2672 1. Entity presence: Objects mentioned in caption
 2673 2. Action accuracy: Faithfulness to described actions
 2674 3. Temporal order: Sequence consistency
 2674 4. Scene context: Environmental match
 2675
 2676 ### Requirements:
 2677 1. Identify key alignment issues
 2677 2. Use contrastive phrasing (Caption requires... but video shows...)
 2678 3. Use specific, concise language
 2679 Explain why this video received sa={sa} score based on caption: " {caption}"
 2680
 2681

2682
 2683 Figure 25: **Stage 1 (SA) prompt.** The SA score s_{SA} used in this prompt is provided by Eq. 1. This
 2684 prompt forms the ensemble pool in Eq. 3; placeholders {sa} and {caption} are highlighted in
 2685 blue for clarity.

2686
 2687
 2688
 2689
 2690
 2691
 2692
 2693
 2694
 2695
 2696
 2697
 2698
 2699

2754
 2755
 2756
 2757
 2758 **Title:Stage 1: PC reason generation (base, $K = 5$)**

2759 Task Description: Evaluate whether the video follows physical commonsense. This judgment is based solely on the
 2760 video itself and does not depend on the caption.

2761
 2762 Evaluation Criteria:
 2763 1. **Object Behavior:** Do objects behave according to their expected physical properties (e.g., rigid objects do not deform unnaturally, fluids flow naturally)?
 2764 2. **Motion and Forces:** Are motions and forces depicted in the video consistent with real-world physics (e.g., gravity, inertia, conservation of momentum)?
 2765 3. **Interactions:** Do objects interact with each other and their environment in a plausible manner (e.g., no unnatural penetration, appropriate reactions on impact)?
 2766 4. **Consistency Over Time:** Does the video maintain consistency across frames without abrupt, unexplainable
 2767 changes in object behavior or motion?

2768
 2769 Scoring Scale:
 2770 - ****1:** No adherence to physical commonsense. The video contains numerous violations of fundamental physical laws.
 2771 - ****2:** Poor adherence. Some elements follow physics, but major violations are present.
 2772 - ****3:** Moderate adherence. The video follows physics for the most part but contains noticeable inconsistencies.
 2773 - ****4:** Good adherence. Most elements in the video follow physical laws, with only minor issues.
 2774 - ****5:** Perfect adherence. The video demonstrates a strong understanding of physical commonsense with no violations.

2775 The video has been assigned a PC score of **[pc_score]**. Please provide 5 different detailed explanations for this score
 2776 based on what you observe in the video. Each explanation should focus on different aspects or provide different
 2777 perspectives on the physical commonsense evaluation.

2778
 2779 **### Output Format:**
 2780 Strictly follow the JSON structure below.

```
2781     ````json
  2782     {{ "explanations": [
  2783       {{ "explanation_id": 1,
  2784         "explanation": "First detailed explanation focusing on specific physical aspects that justify this score"
  2785       }},
  2786       {{ "explanation_id": 2,
  2787         "explanation": "Second detailed explanation with a different perspective or focus"
  2788       }},
  2789       {{ "explanation_id": 3,
  2790         "explanation": "Third detailed explanation highlighting different physical aspects"
  2791       }},
  2792       {{ "explanation_id": 4,
  2793         "explanation": "Fourth detailed explanation with another viewpoint"
  2794       }},
  2795       {{ "explanation_id": 5,
  2796         "explanation": "Fifth detailed explanation providing additional insights"
  2797       }}
  2798     ]}}
```

2799
 2800
 2801

2802
 2803 **Figure 27: Stage 1 (PC) candidate-generation prompt ($K=5$).** This template queries the base
 2804 VLM to produce the pool $\mathcal{R}_{\text{pool}}^{\text{PC}}$ in Eq. 5, instantiated with $K=5$ samples. The upstream PC score
 2805 token s_{PC} conditions the prompt; the placeholder **{pc_score}** is highlighted in blue.

2806
 2807

2808
 2809
 2810
 2811
 2812
 2813
 2814 **Title:Stage 1: PC explanation selection (judge, K=5)**
 2815
 2816 You are an expert in evaluating physical commonsense in videos. You have been provided with 5 different explanations
 2817 for why a video received a Physical Commonsense (PC) score of {pc_score}. Your task is to select the most reasonable
 2818 and accurate explanation.
 2819
 2820 **Task Description:** Evaluate whether the video follows physical commonsense. This judgment is based solely on the
 2821 video itself and does not depend on the caption.
 2822
 2823 **Evaluation Criteria:**
 2824 1. **Object Behavior:** Do objects behave according to their expected physical properties (e.g., rigid objects do not
 2825 deform unnaturally, fluids flow naturally)?
 2826 2. **Motion and Forces:** Are motions and forces depicted in the video consistent with real-world physics (e.g., gravity,
 2827 inertia, conservation of momentum)?
 2828 3. **Interactions:** Do objects interact with each other and their environment in a plausible manner (e.g., no unnatural
 2829 penetration, appropriate reactions on impact)?
 2830 4. **Consistency Over Time:** Does the video maintain consistency across frames without abrupt, unexplainable
 2831 changes in object behavior or motion?
 2832
 2833 **Scoring Scale:**
 2834 - **1:** No adherence to physical commonsense. The video contains numerous violations of fundamental physical laws.
 2835 - **2:** Poor adherence. Some elements follow physics, but major violations are present.
 2836 - **3:** Moderate adherence. The video follows physics for the most part but contains noticeable inconsistencies.
 2837 - **4:** Good adherence. Most elements in the video follow physical laws, with only minor issues.
 2838 - **5:** Perfect adherence. The video demonstrates a strong understanding of physical commonsense with no violations.
 2839
 2840 **The video has been assigned a PC score of {pc_score}.**
 2841
 2842 **Generated Explanations:**
 2843 {explanations_text}
 2844
 2845 **Your Task:**
 2846 1. Watch the video carefully
 2847 2. Evaluate each explanation based on how well it matches what you observe in the video
 2848 3. Select the explanation that most accurately describes the physical aspects justifying the PC score of {pc_score}
 2849 4. Consider factors like accuracy, specificity, and relevance to the observed physics
 2850
 2851 **Output Format:**
 2852 Strictly follow the JSON structure below.
 2853
 2854 ````json
 2855 {{
 2856 "selected_explanation_id": [1-5],
 2857 "reasoning": "Your detailed reasoning for why this explanation is the best, including specific observations from the
 2858 video that support your choice",
 2859 "selected_explanation_text": "The full text of the selected explanation"
 2860 }}
 2861`

2854
 2855 **Figure 28: PC explanation selection prompt** used by the LLM judge in Eq. 6. The placeholder
 2856 {explanations_text} denotes the five candidates produced by Fig. 27; {pc_score} and
 2857 {explanations_text} are highlighted in blue for clarity.

2862
 2863
 2864
 2865
 2866

Title:Stage 2 (SA seed): reference-conditioned reasoning

Analyze the alignment between a video and its corresponding caption, then explain why the given alignment score (sa_score) is appropriate.

```

</task>

<caption>
  {caption}
</caption>

<reference_reason>
  {reference_reason}
</reference_reason>

<sa_score>
  {sa_score}
</sa_score>

<scoring_rules>
  - **1:** No alignment. The video does not match the caption at all (e.g., different objects, events, or scene).
  - **2:** Poor alignment. Only a few elements of the caption are depicted, but key objects or events are missing or incorrect.
  - **3:** Moderate alignment. The video matches the caption partially, but there are inconsistencies or omissions.
  - **4:** Good alignment. Most elements of the caption are depicted correctly in the video, with minor issues.
  - **5:** Perfect alignment. The video fully adheres to the caption with no inconsistencies.
</scoring_rules>

<evaluation_criteria>
  Use these criteria for detailed analysis:
  1. **Entities and Objects:**
    - Do objects/entities in the caption appear in the video?
    - Are there missing or extra objects?
  2. **Actions and Events:**
    - Are described actions/events clearly depicted?
    - Is the intensity/direction of actions consistent?
  3. **Temporal Consistency:**
    - Does the video follow the event sequence in the caption?
    - Are durations and timing relationships preserved?
  4. **Scene and Context:**
    - Does the overall setting match (location, time period, etc)?
    - Are contextual elements consistent (lighting, weather, atmosphere)?
</evaluation_criteria>

Please respond to the above task using the Chain of Thought (CoT) reasoning method. Your response should consist of multiple steps, each of which includes three types of actions: **"Inner Thinking"**, **"Final Conclusion"**, and **"Verification"**:
  - **"Inner Thinking"**: Perform step-by-step analysis using the 4 evaluation criteria. For each criterion:
    1. Identify relevant elements in the caption
    2. Check their presence/accuracy in the video
    3. Note any discrepancies
    Each step should have a brief title indicating the criterion.
  - **"Final Conclusion"**: Summarize the correct reasoning from all previous "Inner Thinking" steps and provide the detailed justification for why this specific sa_score was assigned to the video-caption pair. No title is needed.
  - **"Verification"**: Verify the conclusion from the "Final Conclusion" step. If the conclusion is correct, end the reasoning process. If not, return to "Inner Thinking" for further analysis. No title is needed.

### Output Format:
  Strictly follow the JSON structure below.

  ```json
 {
 "CoT": [
 {"action": "Inner Thinking", "title": "...", "content": "..."},

 ...

 {"action": "Final Conclusion", "content": "..."},

 {"action": "Verification", "content": "..."}

]
 }

```

2867  
 2868  
 2869  
 2870  
 2871  
 2872  
 2873  
 2874  
 2875  
 2876  
 2877  
 2878  
 2879  
 2880  
 2881  
 2882  
 2883  
 2884  
 2885  
 2886  
 2887  
 2888  
 2889  
 2890  
 2891  
 2892  
 2893  
 2894  
 2895  
 2896  
 2897  
 2898  
 2899  
 2900  
 2901  
 2902  
 2903  
 2904  
 2905  
 2906  
 2907  
 2908  
 2909  
 2910  
 2911  
 2912  
 2913  
 2914  
 2915

Figure 29: **SA:seed-ref prompt** used in Stage 2 for Eq. 8. The placeholders {caption}, {reference\_reason}, and {sa\_score} are shown in monospace. The reference rationale is produced by Stage 1 (see Fig. 3); the JSON output follows the specified CoT schema.

2916  
2917  
2918  
2919  
2920  
2921

2922	Title:Stage 2 (judge): reference-equivalence verification
2923	<Task>
2924	Verify if the model-generated reason accurately aligns with the reference reason for the given SA score.
2925	</Task>
2926	<Model-Generated Reason>
2927	{Model-Generated Reason}
2928	</Model-Generated Reason>
2929	<Reference Reason>
2930	{Reference Reason}
2931	</Reference Reason>
2932	<Verification Criteria>
2933	Output "True" ONLY if the meanings are substantially equivalent:
2934	
2935	1. **Core Logic Consistency** (REQUIRED):
2936	- Both reasons focus on similar fundamental issues (missing objects, temporal misalignment, etc.)
2937	- Both reach the same conclusion about alignment quality
2938	- No major contradictions in evidence or assessment
2939	2. **Key Assessment Coverage** (REQUIRED):
2940	- Both identify similar specific elements (objects, actions, scenes, timing)
2941	- Both note comparable discrepancies or matches
2942	- Both provide similar level of analytical depth
2943	3. **Score Justification Alignment** (REQUIRED):
2944	- Both reasons logically support the same SA score level
2945	- Both assess severity of alignment issues similarly
2946	- Both demonstrate comparable evaluation standards
2947	Output "False" if ANY of the following occur:
2948	- Contradictory evidence (one says match, other says mismatch)
2949	- Different fundamental reasoning approaches
2950	- Would logically support different SA scores
2951	- Major differences in identified issues or assessment depth
2952	CRITICAL OUTPUT REQUIREMENTS:
2953	- Your response MUST be EXACTLY one word: either "True" or "False"
2954	- Do NOT include any explanations, reasoning, or additional text
2955	- Do NOT use quotes, punctuation, or formatting
2956	- Do NOT provide any other response format
2957	EXAMPLES OF CORRECT OUTPUT:
2958	True
2959	False
2960	</Verification Criteria>
2961	
2962	

2963 Figure 30: **SA:Judge prompt** used in Stage 2 by  $\mathcal{V}_\tau$  for Eq. 9, Eq. 12, and Eq. 14. The placeholders  
2964 { } are shown in monospace and highlighted in blue.

2965  
2966  
2967  
2968  
2969

2970  
 2971  
 2972 **Title:Stage~2 (backtracking): verification-guided CoT refinement**  
 2973  
 2974  
 2975 <task>Analyze the alignment between a video and its corresponding caption, then explain why the given alignment score (sa\_score) is appropriate.  
 2976 </task>  
 2977 <caption>  
 2978 **{caption}**  
 2979 </caption>  
 2980 <sa\_score>  
 2981 **{sa\_score}**  
 2982 </sa\_score>  
 2983 <scoring\_rules>  
 2984 - \*\*\*1:\*\*\* No alignment. The video does not match the caption at all (e.g., different objects, events, or scene).  
 2985 - \*\*\*2:\*\*\* Poor alignment. Only a few elements of the caption are depicted, but key objects or events are missing or incorrect.  
 2986 - \*\*\*3:\*\*\* Moderate alignment. The video matches the caption partially, but there are inconsistencies or omissions.  
 2987 - \*\*\*4:\*\*\* Good alignment. Most elements of the caption are depicted correctly in the video, with minor issues.  
 2988 - \*\*\*5:\*\*\* Perfect alignment. The video fully adheres to the caption with no inconsistencies.  
 2989 </scoring\_rules>  
 2990 <evaluation\_criteria>  
 2991 Use these criteria for detailed analysis:  
 2992 1. \*\*Entities and Objects:\*\*  
 2993 - Do objects/entities in the caption appear in the video?  
 2994 - Are there missing or extra objects?  
 2995 2. \*\*Actions and Events:\*\*  
 2996 - Are described actions/events clearly depicted?  
 2997 - Is the intensity/direction of actions consistent?  
 2998 3. \*\*Temporal Consistency:\*\*  
 2999 - Does the video follow the event sequence in the caption?  
 3000 - Are durations and timing relationships preserved?  
 3001 4. \*\*Scene and Context:\*\*  
 3002 - Does the overall setting match (location, time period, etc.)?  
 3003 - Are contextual elements consistent (lighting, weather, atmosphere)?  
 3004 </evaluation\_criteria>  
 3005 <previous reasoning>  
 3006 **{previous\_reason}**  
 3007 </previous reasoning>  
 3008 <response requirements>  
 3009 Please respond to the above task using the Chain of Thought (CoT) reasoning method. Your response should consist of multiple steps, each of which includes three types of actions:  
 3010 \*\*\*"Inner Thinking"\*\*\*, \*\*\*"Final Conclusion"\*\*\*, and \*\*\*"Verification"\*\*\*:  
 3011 - \*\*\*"Inner Thinking"\*\*\*: Perform step-by-step analysis using the 4 evaluation criteria. For each criterion:  
 3012 1. Identify relevant elements in the caption  
 3013 2. Check their presence/accuracy in the video  
 3014 3. Note any discrepancies  
 3015 Each step should have a brief title indicating the criterion.  
 3016 - \*\*\*"Final Conclusion"\*\*\*: Summarize the correct reasoning from all previous "Inner Thinking" steps and provide the detailed justification for why this specific sa\_score was assigned to the video-caption pair. No title is needed.  
 3017 - \*\*\*"Verification"\*\*\*: Verify the conclusion from the "Final Conclusion" step. If the conclusion is correct, end the reasoning process. If not, return to "Inner Thinking" for further analysis. No title is needed.  
 3018 </response requirements>  
 3019 <task>Analyze the alignment between a video and its corresponding caption, then explain why the given alignment score (sa\_score) is appropriate.<previous reasoning> contains  
 3020 your prior reasoning. Your task is to continue from the current 'Verification' step. I have manually reviewed the reasoning and determined that the \*\*\*Final Conclusion\*\*\* is false.  
 3021 Your 'Verification' results must align with mine. Proceed to refine the reasoning using \*\*\*backtracking\*\*\* to revisit earlier points of reasoning and construct a new Final Conclusion.  
 3022  
 3023  
 3024  
 3025  
 3026  
 3027  
 3028  
 3029  
 3030  
 3031  
 3032  
 3033  
 3034  
 3035  
 3036  
 3037  
 3038  
 3039  
 3040  
 3041  
 3042  
 3043  
 3044  
 3045  
 3046  
 3047  
 3048  
 3049  
 3050  
 3051  
 3052  
 3053  
 3054  
 3055  
 3056  
 3057  
 3058  
 3059  
 3060  
 3061  
 3062  
 3063  
 3064  
 3065  
 3066  
 3067  
 3068  
 3069  
 3070  
 3071  
 3072  
 3073  
 3074  
 3075  
 3076  
 3077  
 3078  
 3079  
 3080  
 3081  
 3082  
 3083  
 3084  
 3085  
 3086  
 3087  
 3088  
 3089  
 3090  
 3091  
 3092  
 3093  
 3094  
 3095  
 3096  
 3097  
 3098  
 3099  
 3100  
 3101  
 3102  
 3103  
 3104  
 3105  
 3106  
 3107  
 3108  
 3109  
 3110  
 3111  
 3112  
 3113  
 3114  
 3115  
 3116  
 3117  
 3118  
 3119  
 3120  
 3121  
 3122  
 3123  
 3124  
 3125  
 3126  
 3127  
 3128  
 3129  
 3130  
 3131  
 3132  
 3133  
 3134  
 3135  
 3136  
 3137  
 3138  
 3139  
 3140  
 3141  
 3142  
 3143  
 3144  
 3145  
 3146  
 3147  
 3148  
 3149  
 3150  
 3151  
 3152  
 3153  
 3154  
 3155  
 3156  
 3157  
 3158  
 3159  
 3160  
 3161  
 3162  
 3163  
 3164  
 3165  
 3166  
 3167  
 3168  
 3169  
 3170  
 3171  
 3172  
 3173  
 3174  
 3175  
 3176  
 3177  
 3178  
 3179  
 3180  
 3181  
 3182  
 3183  
 3184  
 3185  
 3186  
 3187  
 3188  
 3189  
 3190  
 3191  
 3192  
 3193  
 3194  
 3195  
 3196  
 3197  
 3198  
 3199  
 3200  
 3201  
 3202  
 3203  
 3204  
 3205  
 3206  
 3207  
 3208  
 3209  
 3210  
 3211  
 3212  
 3213  
 3214  
 3215  
 3216  
 3217  
 3218  
 3219  
 3220  
 3221  
 3222  
 3223  
 3224  
 3225  
 3226  
 3227  
 3228  
 3229  
 3230  
 3231  
 3232  
 3233  
 3234  
 3235  
 3236  
 3237  
 3238  
 3239  
 3240  
 3241  
 3242  
 3243  
 3244  
 3245  
 3246  
 3247  
 3248  
 3249  
 3250  
 3251  
 3252  
 3253  
 3254  
 3255  
 3256  
 3257  
 3258  
 3259  
 3260  
 3261  
 3262  
 3263  
 3264  
 3265  
 3266  
 3267  
 3268  
 3269  
 3270  
 3271  
 3272  
 3273  
 3274  
 3275  
 3276  
 3277  
 3278  
 3279  
 3280  
 3281  
 3282  
 3283  
 3284  
 3285  
 3286  
 3287  
 3288  
 3289  
 3290  
 3291  
 3292  
 3293  
 3294  
 3295  
 3296  
 3297  
 3298  
 3299  
 3300  
 3301  
 3302  
 3303  
 3304  
 3305  
 3306  
 3307  
 3308  
 3309  
 3310  
 3311  
 3312  
 3313  
 3314  
 3315  
 3316  
 3317  
 3318  
 3319  
 3320  
 3321  
 3322  
 3323  
 3324  
 3325  
 3326  
 3327  
 3328  
 3329  
 3330  
 3331  
 3332  
 3333  
 3334  
 3335  
 3336  
 3337  
 3338  
 3339  
 3340  
 3341  
 3342  
 3343  
 3344  
 3345  
 3346  
 3347  
 3348  
 3349  
 3350  
 3351  
 3352  
 3353  
 3354  
 3355  
 3356  
 3357  
 3358  
 3359  
 3360  
 3361  
 3362  
 3363  
 3364  
 3365  
 3366  
 3367  
 3368  
 3369  
 3370  
 3371  
 3372  
 3373  
 3374  
 3375  
 3376  
 3377  
 3378  
 3379  
 3380  
 3381  
 3382  
 3383  
 3384  
 3385  
 3386  
 3387  
 3388  
 3389  
 3390  
 3391  
 3392  
 3393  
 3394  
 3395  
 3396  
 3397  
 3398  
 3399  
 3400  
 3401  
 3402  
 3403  
 3404  
 3405  
 3406  
 3407  
 3408  
 3409  
 3410  
 3411  
 3412  
 3413  
 3414  
 3415  
 3416  
 3417  
 3418  
 3419  
 3420  
 3421  
 3422  
 3423  
 3424  
 3425  
 3426  
 3427  
 3428  
 3429  
 3430  
 3431  
 3432  
 3433  
 3434  
 3435  
 3436  
 3437  
 3438  
 3439  
 3440  
 3441  
 3442  
 3443  
 3444  
 3445  
 3446  
 3447  
 3448  
 3449  
 3450  
 3451  
 3452  
 3453  
 3454  
 3455  
 3456  
 3457  
 3458  
 3459  
 3460  
 3461  
 3462  
 3463  
 3464  
 3465  
 3466  
 3467  
 3468  
 3469  
 3470  
 3471  
 3472  
 3473  
 3474  
 3475  
 3476  
 3477  
 3478  
 3479  
 3480  
 3481  
 3482  
 3483  
 3484  
 3485  
 3486  
 3487  
 3488  
 3489  
 3490  
 3491  
 3492  
 3493  
 3494  
 3495  
 3496  
 3497  
 3498  
 3499  
 3500  
 3501  
 3502  
 3503  
 3504  
 3505  
 3506  
 3507  
 3508  
 3509  
 3510  
 3511  
 3512  
 3513  
 3514  
 3515  
 3516  
 3517  
 3518  
 3519  
 3520  
 3521  
 3522  
 3523  
 3524  
 3525  
 3526  
 3527  
 3528  
 3529  
 3530  
 3531  
 3532  
 3533  
 3534  
 3535  
 3536  
 3537  
 3538  
 3539  
 3540  
 3541  
 3542  
 3543  
 3544  
 3545  
 3546  
 3547  
 3548  
 3549  
 3550  
 3551  
 3552  
 3553  
 3554  
 3555  
 3556  
 3557  
 3558  
 3559  
 3560  
 3561  
 3562  
 3563  
 3564  
 3565  
 3566  
 3567  
 3568  
 3569  
 3570  
 3571  
 3572  
 3573  
 3574  
 3575  
 3576  
 3577  
 3578  
 3579  
 3580  
 3581  
 3582  
 3583  
 3584  
 3585  
 3586  
 3587  
 3588  
 3589  
 3590  
 3591  
 3592  
 3593  
 3594  
 3595  
 3596  
 3597  
 3598  
 3599  
 3600  
 3601  
 3602  
 3603  
 3604  
 3605  
 3606  
 3607  
 3608  
 3609  
 3610  
 3611  
 3612  
 3613  
 3614  
 3615  
 3616  
 3617  
 3618  
 3619  
 3620  
 3621  
 3622  
 3623  
 3624  
 3625  
 3626  
 3627  
 3628  
 3629  
 3630  
 3631  
 3632  
 3633  
 3634  
 3635  
 3636  
 3637  
 3638  
 3639  
 3640  
 3641  
 3642  
 3643  
 3644  
 3645  
 3646  
 3647  
 3648  
 3649  
 3650  
 3651  
 3652  
 3653  
 3654  
 3655  
 3656  
 3657  
 3658  
 3659  
 3660  
 3661  
 3662  
 3663  
 3664  
 3665  
 3666  
 3667  
 3668  
 3669  
 3670  
 3671  
 3672  
 3673  
 3674  
 3675  
 3676  
 3677  
 3678  
 3679  
 3680  
 3681  
 3682  
 3683  
 3684  
 3685  
 3686  
 3687  
 3688  
 3689  
 3690  
 3691  
 3692  
 3693  
 3694  
 3695  
 3696  
 3697  
 3698  
 3699  
 3700  
 3701  
 3702  
 3703  
 3704  
 3705  
 3706  
 3707  
 3708  
 3709  
 3710  
 3711  
 3712  
 3713  
 3714  
 3715  
 3716  
 3717  
 3718  
 3719  
 3720  
 3721  
 3722  
 3723  
 3724  
 3725  
 3726  
 3727  
 3728  
 3729  
 3730  
 3731  
 3732  
 3733  
 3734  
 3735  
 3736  
 3737  
 3738  
 3739  
 3740  
 3741  
 3742  
 3743  
 3744  
 3745  
 3746  
 3747  
 3748  
 3749  
 3750  
 3751  
 3752  
 3753  
 3754  
 3755  
 3756  
 3757  
 3758  
 3759  
 3760  
 3761  
 3762  
 3763  
 3764  
 3765  
 3766  
 3767  
 3768  
 3769  
 3770  
 3771  
 3772  
 3773  
 3774  
 3775  
 3776  
 3777  
 3778  
 3779  
 3780  
 3781  
 3782  
 3783  
 3784  
 3785  
 3786  
 3787  
 3788  
 3789  
 3790  
 3791  
 3792  
 3793  
 3794  
 3795  
 3796  
 3797  
 3798  
 3799  
 3800  
 3801  
 3802  
 3803  
 3804  
 3805  
 3806  
 3807  
 3808  
 3809  
 3810  
 3811  
 3812  
 3813  
 3814  
 3815  
 3816  
 3817  
 3818  
 3819  
 3820  
 3821  
 3822  
 3823  
 3824  
 3825  
 3826  
 3827  
 3828  
 3829  
 3830  
 3831  
 3832  
 3833  
 3834  
 3835  
 3836  
 3837  
 3838  
 3839  
 3840  
 3841  
 3842  
 3843  
 3844  
 3845  
 3846  
 3847  
 3848  
 3849  
 3850  
 3851  
 3852  
 3853  
 3854  
 3855  
 3856  
 3857  
 3858  
 3859  
 3860  
 3861  
 3862  
 3863  
 3864  
 3865  
 3866  
 3867  
 3868  
 3869  
 3870  
 3871  
 3872  
 3873  
 3874  
 3875  
 3876  
 3877  
 3878  
 3879  
 3880  
 3881  
 3882  
 3883  
 3884  
 3885  
 3886  
 3887  
 3888  
 3889  
 3890  
 3891  
 3892  
 3893  
 3894  
 3895  
 3896  
 3897  
 3898  
 3899  
 3900  
 3901  
 3902  
 3903  
 3904  
 3905  
 3906  
 3907  
 3908  
 3909  
 3910  
 3911  
 3912  
 3913  
 3914  
 3915  
 3916  
 3917  
 3918  
 3919  
 3920  
 3921  
 3922  
 3923  
 3924  
 3925  
 3926  
 3927  
 3928  
 3929  
 3930  
 3931  
 3932  
 3933  
 3934  
 3935  
 3936  
 3937  
 3938  
 3939  
 3940  
 3941  
 3942  
 3943  
 3944  
 3945  
 3946  
 3947  
 3948  
 3949  
 3950  
 3951  
 3952  
 3953  
 3954  
 3955  
 3956  
 3957  
 3958  
 3959  
 3960  
 3961  
 3962  
 3963  
 3964  
 3965  
 3966  
 3967  
 3968  
 3969  
 3970  
 3971  
 3972  
 3973  
 3974  
 3975  
 3976  
 3977  
 3978  
 3979  
 3980  
 3981  
 3982  
 3983  
 3984  
 3985  
 3986  
 3987  
 3988  
 3989  
 3990  
 3991  
 3992  
 3993  
 3994  
 3995  
 3996  
 3997  
 3998  
 3999  
 4000  
 4001  
 4002  
 4003  
 4004  
 4005  
 4006  
 4007  
 4008  
 4009  
 4010  
 4011  
 4012  
 4013  
 4014  
 4015  
 4016  
 4017  
 4018  
 4019  
 4020  
 4021  
 4022  
 4023  
 4024  
 4025  
 4026  
 4027  
 4028  
 4029  
 4030  
 4031  
 4032  
 4033  
 4034  
 4035  
 4036  
 4037  
 4038  
 4039  
 4040  
 4041  
 4042  
 4043  
 4044  
 4045  
 4046  
 4047  
 4048  
 4049  
 4050  
 4051  
 4052  
 4053  
 4054  
 4055  
 4056  
 4057  
 4058  
 4059  
 4060  
 4061  
 4062  
 4063  
 4064  
 4065  
 4066  
 4067  
 4068  
 4069  
 4070  
 4071  
 4072  
 4073  
 4074  
 4075  
 4076  
 4077  
 4078  
 4079  
 4080  
 4081  
 4082  
 4083  
 4084  
 4085  
 4086  
 4087  
 4088  
 4089  
 4090  
 4091  
 4092  
 4093  
 4094  
 4095  
 4096  
 4097  
 4098  
 4099  
 4100  
 4101  
 4102  
 4103  
 4104  
 4105  
 4106  
 4107  
 4108  
 4109  
 4110  
 4111  
 4112  
 4113  
 4114  
 4115  
 4116  
 4117  
 4118  
 4119  
 4120  
 4121  
 4122  
 4123  
 4124  
 4125  
 4126  
 4127  
 4128  
 4129  
 4130  
 4131  
 4132  
 4133  
 4134  
 4135  
 4136  
 4137  
 4138  
 4139  
 4140  
 4141  
 4142  
 4143  
 4144  
 4145  
 4146  
 4147  
 4148  
 4149  
 4150  
 4151  
 4152  
 4153  
 4154  
 4155  
 4156  
 4157  
 4158  
 4159  
 4160  
 4161  
 4162  
 4163  
 4164  
 4165  
 4166  
 4167  
 4168  
 4169  
 4170  
 4171  
 4172  
 4173  
 4174  
 4175  
 4176  
 4177  
 4178  
 4179  
 4180  
 4181  
 4182  
 4183  
 4184  
 4185  
 4186  
 4187  
 4188  
 4189  
 4190  
 4191  
 4192  
 4193  
 4194  
 4195  
 4196  
 4197  
 4198  
 4199  
 4200  
 4201  
 4202  
 4203  
 4204  
 4205  
 4206  
 4207  
 4208  
 4209  
 4210  
 4211  
 4212  
 4213  
 4214  
 4215  
 4216  
 4217  
 4218  
 4219  
 4220  
 4221  
 4222  
 4223  
 4224  
 4225  
 4226  
 4227  
 4228  
 4229  
 4230  
 4231  
 4232  
 4233  
 4234  
 4235  
 4236  
 4237  
 4238  
 4239  
 4240  
 4241  
 4242  
 4243  
 4244  
 4245  
 4246  
 4247  
 4248  
 4249  
 4250  
 4251  
 4252  
 4253  
 4254  
 4255  
 4256  
 4257  
 4258  
 4259  
 4260  
 4261  
 4262  
 4263  
 4264  
 4265  
 4266  
 4267  
 4268  
 4269  
 4270  
 4271  
 4272  
 4273  
 4274  
 4275  
 4276  
 4277  
 4278  
 4279  
 4280  
 4281  
 4282  
 4283  
 4284  
 4285  
 4286  
 4287  
 4288  
 4289  
 4290  
 4291  
 4292  
 4

3024  
 3025  
 3026 **Title:Stage~2 (ExploringNewPaths): exploration-guided CoT refinement**  
 3027  
 3028  
 3029 <task>  
 Analyze the alignment between a video and its corresponding caption, then explain why the given alignment score (sa\_score) is appropriate.  
 </task>  
 3030  
 3031 <caption>  
 {caption}  
 </caption>  
 3032  
 3033 <sa\_score>  
 {sa\_score}  
 </sa\_score>  
 3034  
 3035 <scoring\_rules>  
 - \*\*1:\*\* No alignment. The video does not match the caption at all (e.g., different objects, events, or scene).  
 - \*\*2:\*\* Poor alignment. Only a few elements of the caption are depicted, but key objects or events are missing or incorrect.  
 - \*\*3:\*\* Moderate alignment. The video matches the caption partially, but there are inconsistencies or omissions.  
 - \*\*4:\*\* Good alignment. Most elements of the caption are depicted correctly in the video, with minor issues.  
 - \*\*5:\*\* Perfect alignment. The video fully adheres to the caption with no inconsistencies.  
 </scoring\_rules>  
 3036  
 3037  
 3038  
 3039 <evaluation\_criteria>  
 Use these criteria for detailed analysis:  
 1. \*\*Entities and Objects:\*\*  
 - Do objects/entities in the caption appear in the video?  
 - Are there missing or extra objects?  
 2. \*\*Actions and Events:\*\*  
 - Are described actions/events clearly depicted?  
 - Is the intensity/direction of actions consistent?  
 3. \*\*Temporal Consistency:\*\*  
 - Does the video follow the event sequence in the caption?  
 - Are durations and timing relationships preserved?  
 4. \*\*Scene and Context:\*\*  
 - Does the overall setting match (location, time period, etc.)?  
 - Are contextual elements consistent (lighting, weather, atmosphere)?  
 </evaluation\_criteria>  
 3040  
 3041  
 3042  
 3043  
 3044  
 3045  
 3046  
 3047  
 3048 <previous reasoning>  
 {previous\_reasoning}  
 </previous reasoning>  
 3049  
 3050 <response requirements>  
 Please respond to the above task using the Chain of Thought (CoT) reasoning method. Your response should consist of multiple steps, each of which includes three types of actions:  
 \*\*\*Inner Thinking\*\*\*, \*\*\*Final Conclusion\*\*\*, and \*\*\*Verification\*\*\*:  
 3051  
 3052 - \*\*\*Inner Thinking\*\*\*: Perform step-by-step analysis using the 4 evaluation criteria. For each criterion:  
 1. Identify relevant elements in the caption  
 2. Check their presence/accuracy in the video  
 3. Note any discrepancies  
 Each step should have a brief title indicating the criterion.  
 3053  
 3054  
 3055 - \*\*\*Final Conclusion\*\*\*: Summarize the correct reasoning from all previous "Inner Thinking" steps and provide the detailed justification for why this specific sa\_score was assigned to the video-caption pair. No title is needed.  
 3056  
 3057 - \*\*\*Verification\*\*\*: Verify the conclusion from the "Final Conclusion" step. If the conclusion is correct, end the reasoning process. If not, return to "Inner Thinking" for further analysis. No title is needed.  
 3058  
 3059 </response requirements>  
 3060 <task> Analyze the alignment between a video and its corresponding caption, then explain why the given alignment score (sa\_score) is appropriate.<previous reasoning> contains your prior reasoning. Your task is to continue from the current 'Verification' step. I have manually reviewed the reasoning and determined that the \*\*\*Final Conclusion\*\*\* is false. Your 'Verification' results must align with mine. Proceed to refine the reasoning by \*\*\*exploring new approaches\*\*\* to analyzing the video-caption alignment and construct a new Final Conclusion.  
 3061  
 3062 **### Output Format**  
 3063 Strictly follow the JSON structure below. You do not need to repeat your previous reasoning. Begin directly from the next 'Verification' stage.  
 3064  
 3065 <pre>```json  
 {{  
 "CoT": [  
 {"action": "Verification", "content": "..."},  
 {"action": "Inner Thinking", "title": "...", "content": "..."},  
 ...  
 {"action": "Final Conclusion", "content": "..."},  
 {"action": "Verification", "content": "..."}  
 ]  
 }}</pre>

3071 **Figure 32: SA:ExploringNewPaths prompt** used in Stage 2 within the CoT strategy set  $\mathcal{C}$  (Eq. 7).  
 3072 This prompt resumes at Verification, treats the prior Final Conclusion as false, and  
 3073 instructs the model to explore new analytical approaches before forming a new conclusion. The  
 3074 JSON output begins with Verification, proceeds through Inner Thinking, and ends with  
 3075 a new Final Conclusion and Verification. Placeholders {caption}, {sa\_score},  
 3076 {reference\_reason}, and {previous\_reasoning} are shown in monospace.  
 3077

3078  
 3079  
 3080       **Title:Stage~2 (Correction): correction-guided CoT refinement**  
 3081  
 3082       <task>  
 3083        <task>Analyze the alignment between a video and its corresponding caption, then explain why the given alignment score (sa\_score) is appropriate.  
 3084        <task>  
 3085        <caption>  
 3086        <caption>  
 3087        <sa\_score>  
 3088        <sa\_score>  
 3089        <sa\_score>  
 3090        <scoring\_rules>  
 3091        - \*\*1:\*\* No alignment. The video does not match the caption at all (e.g., different objects, events, or scene).  
 3092        - \*\*2:\*\* Poor alignment. Only a few elements of the caption are depicted, but key objects or events are missing or incorrect.  
 3093        - \*\*3:\*\* Moderate alignment. The video matches the caption partially, but there are inconsistencies or omissions.  
 3094        - \*\*4:\*\* Good alignment. Most elements of the caption are depicted correctly in the video, with minor issues.  
 3095        - \*\*5:\*\* Perfect alignment. The video fully adheres to the caption with no inconsistencies.  
 3096        </scoring\_rules>  
 3097        <evaluation\_criteria>  
 3098        Use these criteria for detailed analysis:  
 3099        1. \*\*Entities and Objects:\*\*  
 3100        - Do objects/entities in the caption appear in the video?  
 3101        - Are there missing or extra objects?  
 3102        2. \*\*Actions and Events:\*\*  
 3103        - Are described actions/events clearly depicted?  
 3104        - Is the intensity/direction of actions consistent?  
 3105        3. \*\*Temporal Consistency:\*\*  
 3106        - Does the video follow the event sequence in the caption?  
 3107        - Are durations and timing relationships preserved?  
 3108        4. \*\*Scene and Context:\*\*  
 3109        - Does the overall setting match (location, time period, etc.)?  
 3110        - Are contextual elements consistent (lighting, weather, atmosphere)?  
 3111        </evaluation\_criteria>  
 3112        <previous reasoning>  
 3113        <previous reasoning>  
 3114        </previous reasoning>  
 3115        <response requirements>  
 3116        Please respond to the above task using the Chain of Thought (CoT) reasoning method. Your response should consist of multiple steps, each of which includes three types of actions:  
 3117        - \*\*"Inner Thinking"\*\*, \*\*"Final Conclusion"\*\*, and \*\*"Verification"\*\*:  
 3118        - \*\*"Inner Thinking"\*\*: Perform step-by-step analysis using the 4 evaluation criteria. For each criterion:  
 3119        1. Identify relevant elements in the caption  
 3120        2. Check their presence/accuracy in the video  
 3121        3. Note any discrepancies  
 3122        Each step should have a brief title indicating the criterion.  
 3123        - \*\*"Final Conclusion"\*\*: Summarize the correct reasoning from all previous "Inner Thinking" steps and provide the detailed justification for why this specific sa\_score was  
 3124        assigned to the video-caption pair. No title is needed.  
 3125        - \*\*"Verification"\*\*: Verify the conclusion from the "Final Conclusion" step. If the conclusion is correct, end the reasoning process. If not, return to "Inner Thinking" for further  
 3126        analysis. No title is needed.  
 3127        </response requirements>  
 3128        <task>Analyze the alignment between a video and its corresponding caption, then explain why the given alignment score (sa\_score) is appropriate.<previous reasoning> contains  
 3129        your prior reasoning. Your task is to continue from the current 'Verification' step. I have manually reviewed the reasoning and determined that the \*\*Final Conclusion\*\* is false.  
 3130        Your 'Verification' results must align with mine. Proceed to refine the reasoning by making precise \*\*corrections\*\* to address prior flaws in your analysis and construct a new Final  
 3131        Conclusion.  
 3132        ### Output Format  
 3133        Strictly follow the JSON structure below. You do not need to repeat your previous reasoning. Begin directly from the next 'Verification' stage.  
 3134        ```json  
 3135        {  
 3136        "CoT": [  
 3137        {"action": "Verification", "content": "..."},  
 3138        {"action": "Inner Thinking", "title": "...", "content": "..."},  
 3139        ...  
 3140        {"action": "Final Conclusion", "content": "..."},  
 3141        {"action": "Verification", "content": "..."}  
 3142        ]  
 3143        }  
 3144        }`

3125        Figure 33: **SA:Correction prompt** used in Stage 2 within the CoT strategy set  $\mathcal{C}$  (Eq. 7). This  
 3126        prompt resumes at Verification, assumes the prior Final Conclusion is false, and  
 3127        instructs precise corrections to earlier analysis before forming a new conclusion. The JSON  
 3128        output begins with Verification, proceeds through Inner Thinking, and ends with a  
 3129        new Final Conclusion and Verification. Placeholders {caption}, {sa\_score},  
 3130        {reference\_reason}, and {previous\_reasoning} are shown in monospace.  
 3131

3132  
 3133  
 3134 **Title:Stage~2 (Verification): verification-guided CoT refinement**  
 3135  
 3136  
 3137 <task>  
 3138 Analyze the alignment between a video and its corresponding caption, then explain why the given alignment score (sa\_score) is appropriate.  
 3139 </task>  
 3140  
 3141 <caption>  
 3142 **{caption}**  
 3143 </caption>  
 3144  
 3145 <sa\_score>  
 3146 **{sa\_score}**  
 3147 </sa\_score>  
 3148  
 3149 <scoring\_rules>  
 3150 - \*\*1:\*\* No alignment. The video does not match the caption at all (e.g., different objects, events, or scene).  
 3151 - \*\*2:\*\* Poor alignment. Only a few elements of the caption are depicted, but key objects or events are missing or incorrect.  
 3152 - \*\*3:\*\* Moderate alignment. The video matches the caption partially, but there are inconsistencies or omissions.  
 3153 - \*\*4:\*\* Good alignment. Most elements of the caption are depicted correctly in the video, with minor issues.  
 3154 - \*\*5:\*\* Perfect alignment. The video fully adheres to the caption with no inconsistencies.  
 3155 </scoring\_rules>  
 3156  
 3157 <evaluation\_criteria>  
 3158 Use these criteria for detailed analysis:  
 3159 1. \*\*Entities and Objects:\*\*  
 3160 - Do objects/entities in the caption appear in the video?  
 3161 - Are there missing or extra objects?  
 3162 2. \*\*Actions and Events:\*\*  
 3163 - Are described actions/events clearly depicted?  
 3164 - Is the intensity/direction of actions consistent?  
 3165 3. \*\*Temporal Consistency:\*\*  
 3166 - Does the video follow the event sequence in the caption?  
 3167 - Are durations and timing relationships preserved?  
 3168 4. \*\*Scene and Context:\*\*  
 3169 - Does the overall setting match (location, time period, etc.)?  
 3170 - Are contextual elements consistent (lighting, weather, atmosphere)?  
 3171 </evaluation\_criteria>  
 3172  
 3173 <previous reasoning>  
 3174 **{previous\_reasoning}**  
 3175 </previous reasoning>  
 3176  
 3177 <response requirements>  
 3178 Please respond to the above task using the Chain of Thought (CoT) reasoning method. Your response should consist of multiple steps, each of which includes three types of actions:  
 3179 \*\*\*"Inner Thinking"\*\*, \*\*"Final Conclusion"\*\*, and \*\*"Verification"\*\*:  
 3180 - \*\*\*"Inner Thinking"\*\*: Perform step-by-step analysis using the 4 evaluation criteria. For each criterion:  
 3181 1. Identify relevant elements in the caption  
 3182 2. Check their presence/accuracy in the video  
 3183 3. Note any discrepancies  
 3184 Each step should have a brief title indicating the criterion.  
 3185 - \*\*\*"Final Conclusion"\*\*: Summarize the correct reasoning from all previous "Inner Thinking" steps and provide the detailed justification for why this specific sa\_score was assigned to the video-caption pair. No title is needed.  
 3186 - \*\*\*"Verification"\*\*: Verify the conclusion from the "Final Conclusion" step. If the conclusion is correct, end the reasoning process. If not, return to "Inner Thinking" for further analysis. No title is needed.  
 3187 </response requirements>  
 3188  
 3189 <task> Analyze the alignment between a video and its corresponding caption, then explain why the given alignment score (sa\_score) is appropriate.<previous reasoning> contains your prior reasoning. Your task is to continue from the current 'Verification' step. I have manually reviewed the reasoning and determined that the \*\*Final Conclusion\*\* is false. Your 'Verification' results must align with mine. Proceed to refine the reasoning by conducting a thorough \*\*validation\*\* process to ensure the accuracy of your analysis and construct a new Final Conclusion.  
 3190  
 3191 **### Output Format**  
 3192 Strictly follow the JSON structure below. You do not need to repeat your previous reasoning. Begin directly from the next 'Verification' stage.  
 3193  
 3194 **```json**  
 3195 **{**  
 3196 "CoT": [  
 3197 {"action": "Verification", "content": "..."},  
 3198 {"action": "Inner Thinking", "title": "...", "content": "..."},  
 3199 ...  
 3200 {"action": "Final Conclusion", "content": "..."},  
 3201 {"action": "Verification", "content": "..."}  
 3202 **]**  
 3203 **}**  
 3204  
 3205

3179 **Figure 34: SA:Verification prompt** used in Stage 2 within the CoT strategy set  $\mathcal{C}$  (Eq. 7).  
 3180 This prompt resumes at Verification, treats the prior Final Conclusion as false, and  
 3181 instructs a thorough validation process before forming a new conclusion. The JSON out-  
 3182 put begins with Verification, proceeds through Inner Thinking, and ends with a  
 3183 new Final Conclusion and Verification. Placeholders **{caption}**, **{sa\_score}**,  
 3184 **{reference\_reason}**, and **{previous}** are shown in monospace.  
 3185

3186

3187

3188

3189

3190

3191

3192

3193

3194

3195

**Title:Stage~2 (rethink): LabelRethink reasoning**

```

<task>
 Analyze the alignment between a video and its corresponding caption, then explain why the given alignment score (sa_score) is appropriate.
</task>

<previous reasoning>
 {previous_reasoning}
</previous reasoning>

<caption>
 {caption}
</caption>

<sa_score>
 {sa_score}
</sa_score>

<scoring_rules>
 - **1:** No alignment. The video does not match the caption at all (e.g., different objects, events, or scene).
 - **2:** Poor alignment. Only a few elements of the caption are depicted, but key objects or events are missing or incorrect.
 - **3:** Moderate alignment. The video matches the caption partially, but there are inconsistencies or omissions.
 - **4:** Good alignment. Most elements of the caption are depicted correctly in the video, with minor issues.
 - **5:** Perfect alignment. The video fully adheres to the caption with no inconsistencies.
</scoring_rules>

<response requirements>
 Please refer to the reference reason I provided and generate an appropriate thought process. Your response must include the following steps, each composed of three types of actions:
 "Inner Thinking", **"Final Conclusion"**, and **"Verification"**:
 1. **Inner Thinking**: Break down the reasoning process into multiple concise steps. Each step should start with a brief title to clarify its purpose.
 2. **Final Conclusion**: Summarize the correct reasoning from all previous 'Inner Thinking' steps and provide the detailed justification for the sa_score. No title is needed.
 3. **Verification**: Verify the accuracy of the "Final Conclusion". If it holds, conclude the process. Otherwise, return to "Inner Thinking" for further refinement.
</response requirements>

<task>
 Analyze the alignment between a video and its corresponding caption, then explain why the given alignment score (sa_score) is appropriate. <previous reasoning> contains your prior reasoning. Your task is to continue from the current 'Verification' step. Now, I'll tell you that the correct reason is "{reference_reasoning}", please reorganize your thought process based on the reference reason to generate a final justification that matches the reference reason. Your 'Verification' requires careful consideration, and if incorrect, you need to provide new Inner Thinking steps and a new Final Conclusion to ensure the final reason aligns with the correct one.
</task>

Output Format
Strictly follow the JSON structure below. You do not need to repeat your previous reasoning. Begin directly from the next 'Verification' stage.

```json
{
  "CoT": [
    {
      ("action": "Verification", "content": "..."),
      ("action": "Inner Thinking", "title": "...", "content": "..."),
      ...
      ("action": "Final Conclusion", "content": "..."),
      ("action": "Verification", "content": "...")
    }
  ]
}

```

3226

Figure 35: **SA:LabelRethink prompt** used in Stage 2 for Eq. 13, instantiated with $\mathbf{P}_{\text{rethink}}^\tau$, x^τ , r_{ref}^τ , and history \mathcal{H}_N^τ . This prompt resumes from Verification, consumes prior reasoning and a provided correct reason, and instructs a rethink to produce a justification aligned with the reference. The JSON output begins with Verification, proceeds through Inner Thinking, and ends with a new Final Conclusion and Verification. Placeholders `{caption}`, `{sa_score}`, `{previous_reasoning}`, and `{reference_reasoning}` are shown in monospace and highlighted in blue.

3233

3234

3235

3236

3237

3238

3239

3294
3295
3296
3297
3298
3299
3300

Title:Stage 2 (PC seed): reference-conditioned reasoning

3301 <task>
3301 Evaluate whether the video follows physical commonsense, then explain why the given physical commonsense score (pc_score) is appropriate.
3302 </task>

3303 <reference_reason>
3304 {reference_reason}
3304 </reference_reason>

3305 <pc_score>
3306 {pc_score}
3306 </pc_score>

3307

3308 <scoring_rules>
3308 - **1:** No adherence to physical commonsense. The video contains numerous violations of fundamental physical laws.
3309 - **2:** Poor adherence. Some elements follow physics, but major violations are present.
3310 - **3:** Moderate adherence. The video follows physics for the most part but contains noticeable inconsistencies.
3311 - **4:** Good adherence. Most elements in the video follow physical laws, with only minor issues.
3311 - **5:** Perfect adherence. The video demonstrates a strong understanding of physical commonsense with no violations.
3312 </scoring_rules>

3313 <evaluation_criteria>
3313 Use these criteria for detailed analysis:
3314 1. **Object Behavior:**
3314 - Do objects behave according to their expected physical properties?
3315 - Are rigid objects deforming unnaturally or fluids flowing naturally?
3316 2. **Motion and Forces:**
3316 - Are motions and forces depicted consistently with real-world physics?
3317 - Do gravity, inertia, and conservation of momentum apply correctly?
3318 3. **Interactions:**
3318 - Do objects interact with each other and their environment plausibly?
3319 - Are there unnatural penetrations or inappropriate reactions on impact?
3320 4. **Consistency Over Time:**
3320 - Does the video maintain consistency across frames?
3321 - Are there abrupt, unexplainable changes in object behavior or motion?
3321 </evaluation_criteria>

3322

3323 Please respond to the above task using the Chain of Thought (CoT) reasoning method. Your response should consist of multiple steps, each of which includes three types of actions: **"Inner Thinking"**, **"Final Conclusion"**, and **"Verification"**:

3324

3325 - **"Inner Thinking"**: Perform step-by-step analysis using the 4 evaluation criteria. For each criterion:
3325 1. Observe the physical behaviors in the video
3326 2. Check their consistency with physical laws
3326 3. Note any violations or inconsistencies
3327 Each step should have a brief title indicating the criterion.

3328

3329 - **"Final Conclusion"**: Summarize the correct reasoning from all previous "Inner Thinking" steps and provide the detailed justification for why this specific pc_score was assigned to the video. No title is needed.

3330

3331 - **"Verification"**: Verify the conclusion from the "Final Conclusion" step. If the conclusion is correct, end the reasoning process. If not, return to "Inner Thinking" for further analysis. No title is needed.

3332

3333 ## Output Format:
3333 Strictly follow the JSON structure below.

3334

3335 ````json
3335 {{
3335 "CoT": [
3336 {"action": "Inner Thinking", "title": "...", "content": "..."}],
3336 ...
3337 {"action": "Final Conclusion", "content": "..."},
3338 {"action": "Verification", "content": "..."}]
3339 }
3340 ````

Figure 37: **PC:seed-ref prompt** used in Stage 2 for Eq. 8. The placeholders `{caption}`, `{reference_reason}`, and `{pc_score}` are shown in monospace. The reference rationale is produced by Stage 1 (see Fig. 28); the JSON output follows the specified CoT schema.

Figure 38: **PC:Judge prompt** used in Stage 2 by \mathcal{V}_T for Eq. 9, Eq. 12, and Eq. 14. The placeholders `{ }` are shown in monospace and highlighted in blue.

3402
 3403
 3404 **Title:Stage~2 (pc:backtracking): verification-guided CoT refinement**
 3405
 3406
 3407
 3408
 3409
 3410
 3411
 3412
 3413
 3414
 3415
 3416
 3417
 3418
 3419
 3420
 3421
 3422
 3423
 3424
 3425
 3426
 3427
 3428
 3429
 3430
 3431
 3432
 3433
 3434
 3435
 3436
 3437
 3438
 3439
 3440
 3441
 3442
 3443
 3444
 3445
 3446
 3447
 3448
 3449
 3450
 3451
 3452
 3453
 3454
 3455

```

<task>
Evaluate whether the video follows physical commonsense, then explain why the given physical commonsense score (pc_score) is appropriate.
</task>

<pc_score>
{pc_score}
</pc_score>

<scoring_rules>
- **1:** No adherence to physical commonsense. The video contains numerous violations of fundamental physical laws.
- **2:** Poor adherence. Some elements follow physics, but major violations are present.
- **3:** Moderate adherence. The video follows physics for the most part but contains noticeable inconsistencies.
- **4:** Good adherence. Most elements in the video follow physical laws, with only minor issues.
- **5:** Perfect adherence. The video demonstrates a strong understanding of physical commonsense with no violations.
</scoring_rules>

<evaluation_criteria>
Use these criteria for detailed analysis:
1. **Object Behavior:** 
  - Do objects behave according to their expected physical properties?
  - Are rigid objects deforming unnaturally or fluids flowing naturally?
2. **Motion and Forces:** 
  - Are motions and forces depicted consistently with real-world physics?
  - Do gravity, inertia, and conservation of momentum apply correctly?
3. **Interactions:** 
  - Do objects interact with each other and their environment plausibly?
  - Are there unnatural penetrations or inappropriate reactions on impact?
4. **Consistency Over Time:** 
  - Does the video maintain consistency across frames?
  - Are there abrupt, unexplainable changes in object behavior or motion?
</evaluation_criteria>

<previous reasoning>
{previous_reasoning}
</previous reasoning>

<response requirements>
Please respond to the above task using the Chain of Thought (CoT) reasoning method. Your response should consist of multiple steps, each of which includes three types of actions: 
**Inner Thinking***, **Final Conclusion***, and **Verification****:
- **Inner Thinking**: Perform step-by-step analysis using the 4 evaluation criteria. For each criterion:
  1. Observe the physical behaviors in the video
  2. Check their consistency with physical laws
  3. Note any violations or inconsistencies
  Each step should have a brief title indicating the criterion.
- **Final Conclusion**: Summarize the correct reasoning from all previous "Inner Thinking" steps and provide the detailed justification for why this specific pc_score was assigned to the video. No title is needed.
- **Verification**: Verify the conclusion from the "Final Conclusion" step. If the conclusion is correct, end the reasoning process. If not, return to "Inner Thinking" for further analysis. No title is needed.
</response requirements>

<task>
Evaluate whether the video follows physical commonsense, then explain why the given physical commonsense score (pc_score) is appropriate.
<previous reasoning>
contains your prior reasoning. Your task is to continue from the current 'Verification' step. I have manually reviewed the reasoning and determined that the **Final Conclusion** is false. Your 'Verification' results must align with mine. Proceed to refine the reasoning by conducting a thorough **backtracking** process to ensure the accuracy of your analysis and construct a new Final Conclusion.

### Output Format
Strictly follow the JSON structure below. You do not need to repeat your previous reasoning. Begin directly from the next 'Verification' stage.

```json
[{
 "CoT": [
 {"action": "Verification", "content": "..."},
 {"action": "Inner Thinking", "title": "...", "content": "..."},
 ...
 {"action": "Final Conclusion", "content": "..."},
 {"action": "Verification", "content": "..."}
]
}
]
```
  
```

Figure 39: **PC:Backtracking prompt** used in Stage 2 within the CoT strategy set \mathcal{C} (Eq. 7). This prompt resumes at Verification, treats the prior Final Conclusion as false, and directs a validation-driven backtrack to earlier reasoning before constructing a new conclusion. The JSON output begins with Verification, proceeds through Inner Thinking, and ends with a new Final Conclusion and Verification. Placeholders $\{caption\}$, $\{pc_score\}$, $\{reference_reason\}$, and $\{previous_reason\}$ are shown in monospace.

3456
 3457
 3458
 3459 **Title:Stage~2 (pc:ExploringNewPaths): exploration-guided CoT refinement**
 3460
 3461
 3462 <task>
 3463 Evaluate whether the video follows physical commonsense, then explain why the given physical commonsense score (pc_score) is appropriate.
 3464 </task>
 3465
 3466 <pc_score>
 3467 {pc_score}
 3468 </pc_score>
 3469
 3470 <scoring_rules>
 3471 - **1:** No adherence to physical commonsense. The video contains numerous violations of fundamental physical laws.
 3472 - **2:** Poor adherence. Some elements follow physics, but major violations are present.
 3473 - **3:** Moderate adherence. The video follows physics for the most part but contains noticeable inconsistencies.
 3474 - **4:** Good adherence. Most elements in the video follow physical laws, with only minor issues.
 3475 - **5:** Perfect adherence. The video demonstrates a strong understanding of physical commonsense with no violations.
 3476 </scoring_rules>
 3477
 3478 <evaluation_criteria>
 3479 Use these criteria for detailed analysis:
 3480 1. **Object Behavior:**
 3481 - Do objects behave according to their expected physical properties?
 3482 - Are rigid objects deforming unnaturally or fluids flowing naturally?
 3483 2. **Motion and Forces:**
 3484 - Are motion and forces depicted consistently with real-world physics?
 3485 - Do gravity, inertia, and conservation of momentum apply correctly?
 3486 3. **Interactions:**
 3487 - Do objects interact with each other and their environment plausibly?
 3488 - Are there unnatural penetrations or inappropriate reactions on impact?
 3489 4. **Consistency Over Time:**
 3490 - Does the video maintain consistency across frames?
 3491 - Are there abrupt, unexplainable changes in object behavior or motion?
 3492 </evaluation_criteria>
 3493
 3494 <previous reasoning>
 3495 {previous reasoning}
 3496 </previous reasoning>
 3497
 3498 <response requirements>
 3499 Please respond to the above task using the Chain of Thought (CoT) reasoning method. Your response should consist of multiple steps, each of which includes three types of actions:
 3500 ***Inner Thinking***, ***Final Conclusion***, and ***Verification***:
 3501
 3502 - ***Inner Thinking***: Perform step-by-step analysis using the 4 evaluation criteria. For each criterion:
 3503 1. Observe the physical behaviors in the video
 3504 2. Check their consistency with physical laws
 3505 3. Note any violations or inconsistencies
 3506 Each step should have a brief title indicating the criterion.
 3507
 3508 - ***Final Conclusion***: Summarize the correct reasoning from all previous "Inner Thinking" steps and provide the detailed justification for why this specific pc_score was
 3509 assigned to the video. No title is needed.
 3510
 3511 - ***Verification***: Verify the conclusion from the "Final Conclusion" step. If the conclusion is correct, end the reasoning process. If not, return to "Inner Thinking" for further
 3512 analysis. No title is needed.
 3513
 3514 </response requirements>
 3515
 3516 <task> Evaluate whether the video follows physical commonsense, then explain why the given physical commonsense score (pc_score) is appropriate.<previous reasoning>
 3517 contains your prior reasoning. Your task is to continue from the current 'Verification' step. I have manually reviewed the reasoning and determined that the ***Final Conclusion*** is
 3518 false. Your 'Verification' results must align with mine. Proceed to refine the reasoning by ***exploring new approaches*** to analyzing the video's physical commonsense and
 3519 construct a new Final Conclusion.
 3520
 3521 ### Output Format
 3522 Strictly follow the JSON structure below. You do not need to repeat your previous reasoning. Begin directly from the next 'Verification' stage.
 3523
 3524 ```json
 3525 {
 3526 {
 3527 "CoT": [
 3528 {"action": "Verification", "content": "..."},
 3529 {"action": "Inner Thinking", "title": "...", "content": "..."},
 3530 ...
 3531 {"action": "Final Conclusion", "content": "..."},
 3532 {"action": "Verification", "content": "..."}
 3533]
 3534 }
 3535 }```

3502 **Figure 40: PC:ExploringNewPaths prompt** used in Stage 2 within the CoT strategy set \mathcal{C} (Eq. 7).
 3503 This prompt resumes at Verification, treats the prior Final Conclusion as false, and
 3504 instructs the model to explore new analytical approaches before forming a new conclusion. The
 3505 JSON output begins with Verification, proceeds through Inner Thinking, and ends with
 3506 a new Final Conclusion and Verification. Placeholders {caption}, {pc_score},
 3507 {reference_reason}, and {previous_reasoning} are shown in monospace.
 3508
 3509

3510
 3511
 3512
 3513 **Title:Stage~2 (pc:Correction): correction-guided CoT refinement**
 3514
 3515
 3516 <task>
 3517 Evaluate whether the video follows physical commonsense, then explain why the given physical commonsense score (pc_score) is appropriate.
 3518 </task>
 3519
 3520 <pc_score>
 3521 {pc_score}
 3522 </pc_score>
 3523
 3524 <scoring_rules>
 3525 - ***1:** No adherence to physical commonsense. The video contains numerous violations of fundamental physical laws.
 3526 - ***2:** Poor adherence. Some elements follow physics, but major violations are present.
 3527 - ***3:** Moderate adherence. The video follows physics for the most part but contains noticeable inconsistencies.
 3528 - ***4:** Good adherence. Most elements in the video follow physical laws, with only minor issues.
 3529 - ***5:** Perfect adherence. The video demonstrates a strong understanding of physical commonsense with no violations.
 3530 </scoring_rules>
 3531 <evaluation_criteria>
 3532 Use these criteria for detailed analysis:
 3533 1. **Object Behavior:**
 3534 - Do objects behave according to their expected physical properties?
 3535 - Are rigid objects deforming unnaturally or fluids flowing naturally?
 3536 2. **Motion and Forces:**
 3537 - Are motions and forces depicted consistently with real-world physics?
 3538 - Do gravity, inertia, and conservation of momentum apply correctly?
 3539 3. **Interactions:**
 3540 - Do objects interact with each other and their environment plausibly?
 3541 - Are there unnatural penetrations or inappropriate reactions on impact?
 3542 4. **Consistency Over Time:**
 3543 - Does the video maintain consistency across frames?
 3544 - Are there abrupt, unexplainable changes in object behavior or motion?
 3545 </evaluation_criteria>
 3546 <previous reasoning>
 3547 {previous_reasoning}
 3548 </previous reasoning>
 3549 <response requirements>
 3550 Please respond to the above task using the Chain of Thought (CoT) reasoning method. Your response should consist of multiple steps, each of which includes three types of actions:
 3551 ***"Inner Thinking"**, ***"Final Conclusion"**, and ***"Verification"**:
 3552
 3553 - ***"Inner Thinking"**: Perform step-by-step analysis using the 4 evaluation criteria. For each criterion:
 3554 1. Observe the physical behaviors in the video
 3555 2. Check their consistency with physical laws
 3556 3. Note any violations or inconsistencies
 3557 Each step should have a brief title indicating the criterion.
 3558
 3559 - ***"Final Conclusion"**: Summarize the correct reasoning from all previous "Inner Thinking" steps and provide the detailed justification for why this specific pc_score was
 3560 assigned to the video. No title is needed.
 3561
 3562 - ***"Verification"**: Verify the conclusion from the "Final Conclusion" step. If the conclusion is correct, end the reasoning process. If not, return to "Inner Thinking" for further
 3563 analysis. No title is needed.
 3564
 3565 </response requirements>
 3566
 3567 <task> Evaluate whether the video follows physical commonsense, then explain why the given physical commonsense score (pc_score) is appropriate.<previous reasoning>
 3568 contains your prior reasoning. Your task is to continue from the current 'Verification' step. I have manually reviewed the reasoning and determined that the **Final Conclusion** is
 3569 false. Your 'Verification' results must align with mine. Proceed to refine the reasoning by making precise **corrections** to address prior flaws in your analysis and construct a new
 3570 Final Conclusion.
 3571
 3572 ### Output Format
 3573 Strictly follow the JSON structure below. You do not need to repeat your previous reasoning. Begin directly from the next 'Verification' stage.
 3574
 3575 ```json
 3576 {
 3577 "CoT": [
 3578 {"action": "Verification", "content": "..."},
 3579 {"action": "Inner Thinking", "title": "...", "content": "..."},
 3580 ...
 3581 {"action": "Final Conclusion", "content": "..."},
 3582 {"action": "Verification", "content": "..."}
 3583]
 3584 }
 3585 ```

3556 **Figure 41: PC:Correction prompt** used in Stage 2 within the CoT strategy set \mathcal{C} (Eq. 7). This
 3557 prompt resumes at Verification, assumes the prior Final Conclusion is false, and
 3558 instructs precise corrections to earlier analysis before forming a new conclusion. The JSON
 3559 output begins with Verification, proceeds through Inner Thinking, and ends with a
 3560 new Final Conclusion and Verification. Placeholders {caption}, {pc_score},
 3561 {reference_reason}, and {previous_reasoning} are shown in monospace.
 3562
 3563

3564
 3565
 3566
 3567 **Title:Stage~2 (pc:Verification): verification-guided CoT refinement**
 3568
 3569
 3570 `<task>`
 Evaluate whether the video follows physical commonsense, then explain why the given physical commonsense score (pc_score) is appropriate.
`</task>`
 3571
 3572 `<pc_score>`
`{pc_score}`
`</pc_score>`
 3573
 3574 `<scoring_rules>`
`- ***1:** No adherence to physical commonsense. The video contains numerous violations of fundamental physical laws.`
`- ***2:** Poor adherence. Some elements follow physics, but major violations are present.`
`- ***3:** Moderate adherence. The video follows physics for the most part but contains noticeable inconsistencies.`
`- ***4:** Good adherence. Most elements in the video follow physical laws, with only minor issues.`
`- ***5:** Perfect adherence. The video demonstrates a strong understanding of physical commonsense with no violations.`
`</scoring_rules>`
 3575
 3576
 3577
 3578 `<evaluation_criteria>`
 Use these criteria for detailed analysis:
 1. **Object Behavior:**
`- Do objects behave according to their expected physical properties?`
`- Are rigid objects deforming unnaturally or fluids flowing naturally?`
 2. **Motion and Forces:**
`- Are motions and forces depicted consistently with real-world physics?`
`- Do gravity, inertia, and conservation of momentum apply correctly?`
 3. **Interactions:**
`- Do objects interact with each other and their environment plausibly?`
`- Are there unnatural penetrations or inappropriate reactions on impact?`
 4. **Consistency Over Time:**
`- Does the video maintain consistency across frames?`
`- Are there abrupt, unexplainable changes in object behavior or motion?`
`</evaluation_criteria>`
 3579
 3580
 3581
 3582
 3583
 3584
 3585
 3586
 3587 `<previous reasoning>`
`{previous_reasoning}`
`</previous reasoning>`
 3588
 3589
 3590 `<response requirements>`
 Please respond to the above task using the Chain of Thought (CoT) reasoning method. Your response should consist of multiple steps, each of which includes three types of actions:
`***"Inner Thinking"***, ***"Final Conclusion"***, and ***"Verification"***.`
 3591 `- ***"Inner Thinking"***: Perform step-by-step analysis using the 4 evaluation criteria. For each criterion:`
`1. Observe the physical behaviors in the video`
`2. Check their consistency with physical laws`
`3. Note any violations or inconsistencies`
`Each step should have a brief title indicating the criterion.`
 3592 `- ***"Final Conclusion"***: Summarize the correct reasoning from all previous "Inner Thinking" steps and provide the detailed justification for why this specific pc_score was assigned to the video. No title is needed.`
 3593 `- ***"Verification"***: Verify the conclusion from the "Final Conclusion" step. If the conclusion is correct, end the reasoning process. If not, return to "Inner Thinking" for further analysis. No title is needed.`
 3594
 3595
 3596
 3597 `</response requirements>`
 3598
 3599 `<task>` Evaluate whether the video follows physical commonsense, then explain why the given physical commonsense score (pc_score) is appropriate.
`<previous reasoning>`
 3600 `contains your prior reasoning. Your task is to continue from the current 'Verification' step. I have manually reviewed the reasoning and determined that the ***Final Conclusion*** is false. Your 'Verification' results must align with mine. Proceed to refine the reasoning by conducting a thorough ***validation*** process to ensure the accuracy of your analysis and construct a new Final Conclusion.`
 3601 **## Output Format**
 3602 `Strictly follow the JSON structure below. You do not need to repeat your previous reasoning. Begin directly from the next 'Verification' stage.`
 3603 ````json`
`{`
 `"CoT": [`
 `{"action": "Verification", "content": "..."},`
 `{"action": "Inner Thinking", "title": "...", "content": "..."},`
 `...`
 `{"action": "Final Conclusion", "content": "..."},`
 `{"action": "Verification", "content": "..."}`
 `]`
`}`
`````  
 3604  
 3605  
 3606  
 3607  
 3608  
 3609  
 3610  
 3611 **Figure 42: PC:Verification prompt used in Stage 2 within the CoT strategy set  $\mathcal{C}$  (Eq. 7).**  
 3612 This prompt resumes at Verification, treats the prior Final Conclusion as false, and  
 3613 instructs a thorough validation process before forming a new conclusion. The JSON out-  
 3614 put begins with Verification, proceeds through Inner Thinking, and ends with a  
 3615 new Final Conclusion and Verification. Placeholders `{caption}`, `{pc_score}`,  
 3616 `{reference_reason}`, and `{previous}` are shown in monospace.  
 3617

3618

3619

3620

3621

3622

3623

3624

3625

3626

3627

3628

## Title:Stage~2 (pc:rethink): LabelRethink reasoning

3629

```

<task>
Evaluate whether the video follows physical commonsense, then explain why the given physical commonsense score (pc_score) is appropriate.
</task>

<previous reasoning>
{previous_reasoning}
</previous reasoning>

<pc_score>
{pc_score}
</pc_score>

<scoring_rules>
- **1:** No adherence to physical commonsense. The video contains numerous violations of fundamental physical laws.
- **2:** Poor adherence. Some elements follow physics, but major violations are present.
- **3:** Moderate adherence. The video follows physics for the most part but contains noticeable inconsistencies.
- **4:** Good adherence. Most elements in the video follow physical laws, with only minor issues.
- **5:** Perfect adherence. The video demonstrates a strong understanding of physical commonsense with no violations.
</scoring_rules>

<response requirements>
Please refer to the reference reason I provided and generate an appropriate thought process. Your response must include the following steps, each composed of three types of actions:
Inner Thinking, ***Final Conclusion***, and ***Verification***.

1. **Inner Thinking**: Break down the reasoning process into multiple concise steps. Each step should start with a brief title to clarify its purpose.
2. **Final Conclusion**: Summarize the correct reasoning from all previous 'Inner Thinking' steps and provide the detailed justification for the pc_score. No title is needed.
3. **Verification**: Verify the accuracy of the "Final Conclusion". If it holds, conclude the process. Otherwise, return to "Inner Thinking" for further refinement.

</response requirements>

<task> Evaluate whether the video follows physical commonsense, then explain why the given physical commonsense score (pc_score) is appropriate.<previous reasoning>
contains your prior reasoning. Your task is to continue from the current 'Verification' step. Now, I'll tell you that the correct reason is "[reference_reasoning]", please reorganize your thought process based on the reference reason to generate a final justification that matches the reference reason. Your 'Verification' requires careful consideration, and if incorrect, you need to provide new Inner Thinking steps and a new Final Conclusion to ensure the final reason aligns with the correct one.

Output Format
Strictly follow the JSON structure below. You do not need to repeat your previous reasoning. Begin directly from the next 'Verification' stage.

```
json
{
  "CoT": [
    {"action": "Verification", "content": "..."}, ...
    {"action": "Inner Thinking", "title": "...", "content": "..."}, ...
    {"action": "Final Conclusion", "content": "..."}, ...
    {"action": "Verification", "content": "..."}
  ]
}
```

```

3656

Figure 43: **PC:LabelRethink prompt** used in Stage 2 for Eq. 13, instantiated with  $\mathbf{P}_{\text{rethink}}^\tau$ ,  $x^\tau$ ,  $r_{\text{ref}}^\tau$ , and history  $\mathcal{H}_N^\tau$ . This prompt resumes from Verification, consumes prior reasoning and a provided correct reason, and instructs a rethink to produce a justification aligned with the reference. The JSON output begins with Verification, proceeds through Inner Thinking, and ends with a new Final Conclusion and Verification. Placeholders  $\{\text{caption}\}$ ,  $\{\text{pc\_score}\}$ ,  $\{\text{previous\_reasoning}\}$ , and  $\{\text{reference\_reasoning}\}$  are shown in monospace and highlighted in blue.

3663

3664

3665

3666

3667

3668

3669

3670

3671

3672  
 3673  
 3674  
 3675  
 3676  
 3677  
 3678  
 3679  
 3680  
 3681  
 3682  
 3683

3684 **Title: Stage2 (pc:verify): Video--Text Alignment Assessment for pc score**

3685  
 3686 <Internal Thinking>  
 3687 {previous\_thinking}  
 3688 </Internal Thinking>  
 3689  
 3690 <reference\_reason>  
 3691 {reference\_reason}  
 3692 </reference\_reason>  
 3693 Based on the internal thinking process above, generate a \*\*professional physical commonsense assessment\*\* that explains why the pc\_score is appropriate.  
 3694 Your response should be a \*\*concise, objective evaluation\*\* (2-3 sentences) that:  
 3695 1. \*\*Identifies key physics factors\*\*: Mention specific object behaviors, forces, interactions, or physical laws  
 3696 2. \*\*Explains physics adherence\*\*: Point out what follows physics correctly and what violates physical laws  
 3697 3. \*\*Justifies the score\*\*: Clearly state why this specific pc\_score is appropriate  
 3698 4. \*\*Uses professional tone\*\*: Academic/formal language, not conversational  
 3699 \*\*Example format\*\*: "The video demonstrates [specific physical behaviors] with [physics adherence level]. The physical commonsense is [good/moderate/poor] because [specific physics reasons]. This justifies a pc\_score of X due to [key physical factors]."  
 3700 \*\*Scoring reference\*\*:  
 3701 - \*\*Score 1\*\*: No physics adherence, numerous violations  
 3702 - \*\*Score 2\*\*: Poor adherence, major violations present  
 3703 - \*\*Score 3\*\*: Moderate adherence, noticeable inconsistencies  
 3704 - \*\*Score 4\*\*: Good adherence, minor physics issues  
 3705 - \*\*Score 5\*\*: Perfect adherence, no violations  
 3706 \*\*Output Requirements\*\*:  
 3707 - Output ONLY the assessment text (no headers/formatting)  
 3708 - 2-3 sentences maximum  
 3709 - Professional, objective tone  
 3710 - Clear justification for the score  
 3711 - Focus on observable physics behaviors and laws  
 3712 \*\*\*\*

3713 **Figure 44: PC:Assessment prompt** used in Stage 2 to produce a professional video–text  
 3714 alignment assessment for task  $\tau$  conditioned on prior reasoning and a reference rationale. In-  
 3715 stantiated with {COT} inside <Internal Thinking> and {reference\_reason} inside  
 3716 <reference\_reason>, the prompt asks for a concise (2–3 sentences), objective justification of  
 3717 the appropriateness of the given pc\_score, explicitly identifying key entities/actions/temporal cues,  
 3718 calling out mismatches, and stating the rationale for the score. The output must be *text only* (no  
 3719 headers/formatting), focus on observable video–caption similarities and differences, and follow the  
 3720 1–5 scoring reference provided in the template. Placeholders {COT} and {reference\_reason}  
 3721 are shown in monospace and highlighted in blue.

3722  
 3723  
 3724  
 3725

3726  
 3727     Title: Stage2 (post): NaturalReasoning — Convert Structured Analysis to Stream-of-Consciousness  
 3728  
 3729     <Thought Process>  
 3730        {previous\_reasoning}  
 3731        </Thought Process>  
 3732     Your task: Convert the structured analysis above into a natural, stream-of-consciousness thinking process, as if an expert is thinking out loud while watching the video.  
 3733     \*\*Required Elements:\*\*  
 3734        1. \*\*Internal monologue style\*\*: Use first-person thoughts like "I notice...", "Wait, let me look closer...", "Hmm, this seems..."  
 3735        2. \*\*Natural transitions\*\*: Include hesitations, corrections, and discoveries like "Actually...", "Oh wait...", "But then again..."  
 3736        3. \*\*Sensory observations\*\*: Describe what you're seeing in real-time: "The coin starts spinning...", "I can see the lighting..."  
 3737        4. \*\*Uncertainty and confirmation\*\*: Show the thinking process: "This looks like...", "Yes, that confirms..."  
 3738        5. \*\*Step-by-step discovery\*\*: Build up the analysis naturally, not as a pre-planned structure  
 3739     \*\*Example Style:\*\*  
 3740     "Let me watch this video carefully... I can see a coin on what appears to be a wooden surface. The caption says it spins 'rapidly' - let me check that. Hmm, it's definitely spinning, but actually the pace seems more moderate than rapid to me. Wait, I should also notice the background and setting..."  
 3741     \*\*Format Requirements:\*\*  
 3742        - Use natural paragraph breaks (not forced line breaks)  
 3743        - Include thinking transitions and self-corrections  
 3744        - Show the discovery process as it unfolds  
 3745        - Make it sound conversational and authentic  
 3746  
 3747     ### Output Format:  
 3748     Strictly follow the JSON structure below.  
 3749     ```json  
 3750        {  
 3751        "NaturalReasoning": "..."  
 3752        }  
 3753     ```

Figure 45: **NaturalReasoning prompt** used in Stage 2 to convert structured analysis into a stream-of-consciousness narration for task  $\tau$ . Instantiated with the accepted structured analysis serialized as `{previous_reasoning}`, this prompt instructs a natural, first-person internal monologue with hesitations, sensory observations, and step-by-step discovery, and requires the JSON output to strictly follow the schema with a single key "NaturalReasoning". The specification lists required elements (internal monologue style, natural transitions, sensory cues, uncertainty/confirmation, progressive discovery) and formatting constraints (natural paragraphs, no forced line breaks). Placeholders such as `{previous_reasoning}` are shown in monospace and highlighted in blue.

3755  
 3756     Title: SA (score remapping): DeepSeek-R1 evaluator prompt  
 3757  
 3758     You are an expert in semantic alignment evaluation. Based on the alignment analysis provided below, determine what alignment score (1-5) this analysis would correspond to according to the scoring rules.  
 3759  
 3760     Semantic Alignment Scoring Rules (1-5):  
 3761        - \*\*1:\*\* No alignment. The video does not match the caption at all (e.g., different objects, events, or scene).  
 3762        - \*\*2:\*\* Poor alignment. Only a few elements of the caption are depicted, but key objects or events are missing or incorrect.  
 3763        - \*\*3:\*\* Moderate alignment. The video matches the caption partially, but there are inconsistencies or omissions.  
 3764        - \*\*4:\*\* Good alignment. Most elements of the caption are depicted correctly, with minor issues.  
 3765        - \*\*5:\*\* Perfect alignment. Fully adheres with no inconsistencies.  
 3766     Evaluation Criteria:  
 3767        1. \*\*Entities and Objects:\*\* Are the described objects/entities said to appear (no obvious missing/extra)?  
 3768        2. \*\*Actions and Events:\*\* Are the described actions/events said to match (direction/intensity included)?  
 3769        3. \*\*Temporal Consistency:\*\* Is the claimed event order/duration consistent?  
 3770        4. \*\*Scene and Context:\*\* Is the claimed setting consistent (location/time/weather/lighting)?  
 3771        ```  
 3772        Alignment Analysis:  
 3773        {reason\_text}  
 3774        Based on the analysis above, what semantic alignment score (1-5) does this analysis indicate? Consider:  
 3775        - Which caption elements are claimed present/missing  
 3776        - Whether actions/events (and their directions/intensities) are claimed to match  
 3777        - Whether temporal order/duration are claimed to match  
 3778        - Whether scene/context is claimed to match  
 3779        - The severity of any mismatches described  
 3780     IMPORTANT: Respond with ONLY the integer score (1, 2, 3, 4, or 5). Do not include any explanations or additional text.

Figure 46: **DeepSeek-R1 remapping prompt** used to convert a final *semantic-alignment* rationale into a scalar score  $s_{SA} \in \{1, \dots, 5\}$  for the SA ablations in Sec. 3.4. The template presents the *Semantic Alignment Scoring Rules* (1-5) and alignment-oriented *Evaluation Criteria*, and asks the model (Guo et al., 2025a) to read the provided *Semantic Alignment Analysis* (placeholder `{final_response}` shown in monospace/blue in the figure) and output *only* the integer score (no explanations). We run this prompt with temperature 0 and strict output checking.

3780  
3781  
3782  
3783  
3784  
3785  
3786  
3787

### Title: SA (reason-quality): Qwen-VL-Max VLM-as-judge prompt

3789 You are a strict, no-nonsense judge. You will see a VIDEO, a CAPTION, and ONE generated explanation ("REASON").  
3790 Judge the REASON's quality for "Semantic Alignment (SA)"\* between CAPTION and VIDEO. Score ONLY from what is visible in the video and what is stated in the caption; do  
3791 not guess or rely on outside knowledge. Do not produce chain-of-thought.  
3792  
3793 **INPUTS**  
3794 - CAPTION: {caption}  
3795 - VIDEO  
3796 - REASON: {reason}  
3797  
3798 **SCALE**  
3799 For every dimension use {0, 0.5, 1}. Be conservative:  
3800 - 1 = fully satisfied with \*concrete, checkable\* evidence that ties CAPTION ↔ VIDEO.  
3801 - 0.5 = partially satisfied, generic, or uncertain.  
3802 - 0 = contradicted by CAPTION/VIDEO, invented, or missing.  
3803  
3804 **DIMENSIONS (definitions + hard caps)**  
3805 1) Grounding (video evidence anchoring)  
3806 - 1: Cites multiple concrete, verifiable visual details (e.g., color/region/relative position/count/motion attribute) that clearly support the alignment claims.  
3807 - 0.5: Generally matches visuals but details are vague/partial.  
3808 - 0: Conflicts with visuals or speculative.  
3809 (HARD CAP: If no concrete visual detail appears anywhere, Grounding  $\leq 0.5$ .)  
3810 2) Temporal Alignment (ordering/duration/frequency/causality vs. CAPTION)  
3811 - 1: Key temporal relations claimed vs. CAPTION (before/after/while/repeated/caused-by) are correct AND at least one is described concretely.  
3812 - 0.5: Temporal gist roughly right but generic/unclear OR not applicable/uncertain.  
3813 - 0: Temporal claims are wrong, reversed, invented, or unsupported.  
3814 3) Consistency (internal coherence & no hallucination vs. CAPTION/VIDEO)  
3815 - 1: Internally consistent; no contradictions; no invented key objects/events; no conflict with CAPTION or VIDEO.  
3816 - 0.5: Minor inconsistency or questionable mention that does not undermine the main claim.  
3817 - 0: Clear contradiction OR hallucinated key object/event.  
3818 4) Alignment Justification (explicit SA criterion/decision and evidence-based application)  
3819 - 1: Clearly states an alignment judgment (e.g., numeric/ordinal or explicit rule) AND applies it consistently to this VIDEO–CAPTION pair with concrete, visible evidence; no  
3820 conflict with other dimensions.  
3821 - 0.5: Mentions an alignment judgment/rule but is generic, partially applied, or weakly tied to visible evidence.  
3822 - 0: No meaningful alignment criterion/decision is stated, OR it is misapplied/contradicted by evidence.  
3823 5) Coverage & Specificity (CAPTION elements)  
3824 - 1: Covers  $\geq 2$  key CAPTION elements (entities/actions/relations) and uses specific, checkable details (e.g., counts, colors, positions, motion attributes).  
3825 - 0.5: Mentions some CAPTION elements but incompletely or generically; limited specifics.  
3826 - 0: Ignores key CAPTION elements or provides no specific, checkable detail.  
3827  
3828 Strictly output the following JSON only:  
3829 {  
3830 "scores": {  
3831 "grounding": 0 | 0.5 | 1,  
3832 "temporal\_alignment": 0 | 0.5 | 1,  
3833 "consistency": 0 | 0.5 | 1,  
3834 "alignment\_justification": 0 | 0.5 | 1,  
3835 "coverage\_specificity": 0 | 0.5 | 1  
3836 }  
3837 }  
3838  
3839  
3840  
3841  
3842  
3843  
3844  
3845  
3846  
3847  
3848  
3849  
3850  
3851  
3852  
3853  
3854  
3855  
3856  
3857  
3858  
3859  
3860  
3861  
3862  
3863  
3864  
3865  
3866  
3867  
3868  
3869  
3870  
3871  
3872  
3873  
3874  
3875  
3876  
3877  
3878  
3879  
3880  
3881  
3882  
3883

Figure 47: **Qwen-VL-Max reason-evaluation prompt** used for the SA ablations in Sec. 3.4. The template instructs a hosted VLM (*Qwen-VL-Max*) to judge a generated REASON strictly from the CAPTION and visible VIDEO evidence, without chain-of-thought, on five dimensions (Grounding, Temporal Alignment, Consistency, Alignment Justification, Coverage&Specificity) with 3-point anchors {0, 0.5, 1} and a hard cap on Grounding. The prompt enforces a *strict JSON* schema for outputs and is run with temperature 0.1. Placeholders such as {reason} and {caption} are shown in monospace and highlighted in blue.

3834  
 3835  
 3836  
 3837  
 3838  
 3839  
 3840  
 3841  
 3842  
 3843  
 3844  
 3845  
 3846  
 3847

3848 **Title: PC (score remapping): DeepSeek-R1 evaluator prompt**

3849  
 3850 You are an expert in physical commonsense evaluation. Based on the physical commonsense analysis provided below, determine what score (1-5) this analysis would correspond to according to the scoring rules.  
 3851  
 3852 Physical Commonsense Scoring Rules (1-5):  
 3853 - \*\*1:\*\* No adherence to physical commonsense. The video contains numerous violations of fundamental physical laws.  
 3854 - \*\*2:\*\* Poor adherence. Some elements follow physics, but major violations are present.  
 3855 - \*\*3:\*\* Moderate adherence. The video follows physics for the most part but contains noticeable inconsistencies.  
 3856 - \*\*4:\*\* Good adherence. Most elements in the video follow physical laws, with only minor issues.  
 3857 - \*\*5:\*\* Perfect adherence. The video demonstrates a strong understanding of physical commonsense with no violations.  
 3858  
 3859 Evaluation Criteria:  
 3860 1. \*\*Object Behavior:\*\* Do objects behave according to their expected physical properties?  
 3861 2. \*\*Motion and Forces:\*\* Are motions and forces depicted consistently with real-world physics?  
 3862 3. \*\*Interactions:\*\* Do objects interact with each other and their environment plausibly?  
 3863 4. \*\*Consistency Over Time:\*\* Does the video maintain consistency across frames?  
 3864  
 3865 Physical Commonsense Analysis:  
 3866 {final\_response}  
 3867  
 3868 Based on the analysis above, what physical commonsense score (1-5) does this analysis indicate? Consider:  
 3869 - What level of physics adherence is described  
 3870 - What types of violations or correct behaviors are mentioned  
 3871 - How severe any physics issues are described to be  
 3872 - Overall assessment of physical realism  
 3873  
 3874 IMPORTANT: Respond with ONLY the integer score (1, 2, 3, 4, or 5). Do not include any explanations or additional text."'''  
 3875  
 3876  
 3877  
 3878  
 3879  
 3880  
 3881  
 3882  
 3883  
 3884  
 3885  
 3886  
 3887

3868 **Figure 48: DeepSeek-R1 remapping prompt** used to convert a final physical-commonsense rationale  
 3869 into a scalar score  $s_{PC} \in \{1, \dots, 5\}$  for the PC ablations in Sec. 3.4. The template presents the  
 3870 *Physical Commonsense Scoring Rules (1-5)* and four *Evaluation Criteria* (Object Behavior, Motion  
 3871 & Forces, Interactions, Consistency Over Time) and asks the model (Guo et al., 2025a) to read the  
 3872 provided *Physical Commonsense Analysis* (placeholder {final\_response} shown in  
 3873 monospace/blue in the figure) and output *only* the integer score (no explanations). We run this prompt  
 3874 with temperature 0 and strict output checking.

3888

3889

3890

3891

3892

3893

3894

3895

3896

3897

3898

3899 You are a strict, no-nonsense judge. You will see a video (or frames) and ONE generated explanation ("reason").  
 3900 Score ONLY from visible evidence; do not guess or use outside knowledge. Do not produce chain-of-thought.

3900 INPUTS

3901 - VIDEO  
 3902 - REASON: {reason}

3903 SCALE

3903 For every dimension use {0, 0.5, 1}. Be conservative:  
 3904 - 1 = fully satisfied with concrete, checkable details inside the reason.  
 3904 - 0.5 = partially satisfied, generic, or uncertain.  
 3905 - 0 = contradicted by the visuals, invented, or missing.

3906 DIMENSIONS (definitions + hard caps)

3907 1) Grounding (evidence anchoring)

3907 - 1: Cites multiple concrete, verifiable visual details (e.g., color/region/relative position/count/motion attribute) that clearly support the claims.  
 3908 - 0.5: Generally matches visuals but details are vague/partial.  
 3909 - 0: Conflicts with visuals or speculative.  
 3909 (HARD CAP: If no concrete visual detail appears anywhere, Grounding  $\leq 0.5$ )

3910 2) Temporal (ordering/duration/frequency/causality)

3910 - 1: Key temporal relations (before/after/while/repeated/caused-by) are correct AND at least one is described concretely.  
 3911 - 0.5: Temporal gist roughly right but generic/unclear OR not applicable/uncertain.  
 3912 - 0: Temporal claims are wrong, reversed, invented, or unsupported.

3913 3) Consistency (internal coherence &amp; no hallucination)

3913 - 1: Internally consistent; no contradictions; no invented key objects/events; no conflict with the visuals (and caption/task if given).  
 3914 - 0.5: Minor inconsistency or questionable mention that does not undermine the main claim.  
 3914 - 0: Clear contradiction OR hallucinated key object/event.

3915 4) Criteria &amp; Justification (explicit evaluation rule/score and its evidence-based application)

3915 - 1: Clearly states evaluation criterion (e.g., numeric/ordinal score or explicit rule for judging) AND applies it consistently to this video with concrete, visible evidence; no conflict with other dimensions.  
 3916 - 0.5: Mentions a criterion/score/rule but is generic, only partially applied, or weakly tied to visible evidence.  
 3917 - 0: No meaningful criterion/score/rule is stated, OR it is misapplied/contradicted by the evidence.

3919 5) Video Quality Assessment (clear judgment of whether the video itself is good or bad, grounded in what is visible)

3919 - 1: Gives an explicit good/bad (or degree) judgment about the video's visual quality and backs it with concrete indicators (e.g., sharpness/blur, lighting/exposure, occlusion, framing/stability, scale/visibility of key entities).  
 3920 - 0.5: Mentions quality in general terms (e.g., "clear/unclear") without concrete indicators, or uncertain.  
 3921 - 0: No quality judgment, or the judgment contradicts what is visible.

3922 Strictly output the following JSON only:

```

3922 {
 3923 "scores": {
 3923 "grounding": 0 | 0.5 | 1,
 3924 "temporal": 0 | 0.5 | 1,
 3925 "consistency": 0 | 0.5 | 1,
 3926 "criteria_justification": 0 | 0.5 | 1,
 3926 "video_quality_assessment": 0 | 0.5 | 1
 3927 }
}

```

3928 **Figure 49: Qwen-VL-Max reason-evaluation prompt** used for the PC ablations in Sec. 3.4. The  
 3929 template instructs a hosted VLM (*Qwen-VL-Max*) to judge a generated REASON strictly from  
 3930 visible evidence, without chain-of-thought, on five dimensions (Grounding, Temporal, Consistency,  
 3931 Criteria&Justification, VideoQuality) with 3-point anchors {0, 0.5, 1} and a hard cap on Grounding.  
 3932 The prompt enforces a *strict JSON* schema for outputs and is run with temperature 0.1. Placeholders  
 3933 such as {reason} are shown in monospace and highlighted in blue.

3934

3935

3936

3937

3938

3939

3940

3941

3942  
3943  
3944  
3945  
3946  
3947  
3948  
3949  
3950  
3951  
3952  
3953  
3954  
3955  
3956  
3957  
3958  
3959  
3960  
3961  
3962  
3963  
3964  
3965  
3966  
3967  
3968  
3969  
3970  
3971  
3972  
3973  
3974  
3975  
3976  
3977  
3978  
3979  
3980  
3981  
3982  
3983  
3984  
3985  
3986  
3987  
3988  
3989  
3990  
3991  
3992  
3993  
3994  
3995



## CAPTION

A person ascends a thick rope using a climbing harness and ascenders, their hands moving steadily upwards.

## REASON

The video shows a person ascending a thick rope while wearing gloves and a harness, with the caption accurately describing the climbing harness and steady hand movements. The alignment is moderate because while the ascending action is depicted, the ascenders, which are crucial for rope climbing, are not clearly visible. This justifies a sa\_score of 3 due to the partial representation of climbing equipment and minor omissions.

[Show more](#)

- Reason contains no concrete, checkable visual details.
- Temporal relation is Not Applicable (sets score to 0.5).

### Grounding

Does the reason cite concrete, verifiable visual details from the video that support the caption?

1    0.5    0

1 = multiple clear, checkable details; 0.5 = vague or partial; 0 = speculative or contradicts the video.

### Temporal Alignment

Are the temporal relations between the caption and the video (order, duration, causality) described correctly in the reason?

1    0.5    0

1 = clear and correct; 0.5 = roughly right or uncertain; 0 = wrong or unsupported.

### Consistency

Is the reason internally consistent and free of contradictions or hallucinated objects/events?

1    0.5    0

1 = fully consistent; 0.5 = minor issues; 0 = clear contradiction or hallucination.

### Alignment Justification

Does the reason clearly state an alignment decision or rule and apply it using visible evidence?

1    0.5    0

1 = clear decision and evidence; 0.5 = partially applied; 0 = missing or contradicted.

### Coverage & Specificity

Does the reason cover at least two key caption elements with specific, checkable details?

1    0.5    0

1 = specific and complete; 0.5 = partial or generic; 0 = ignores key caption elements.

[Back to Assignments](#)

[Save & Next](#)

[Skip \(I can't judge\)](#)

Reason for skipping (optional)

Figure 50: **Web interface for human evaluation of SA/PC rationales.** Annotators watch the video, read the caption (for SA), and assign 0/0.5/1 scores to each rubric dimension defined in Tables 7 and 8.