

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 COMPOSE AND FUSE: REVISITING THE FOUNDATIONAL BOTTLENECKS IN MULTIMODAL REASONING

Anonymous authors

Paper under double-blind review

ABSTRACT

Multimodal large language models (MLLMs) promise enhanced reasoning by integrating diverse inputs such as text, vision, and audio. Yet, despite their perceptual strengths, their reasoning ability across modalities remains underexplored, with conflicting reports on whether additional modalities help or harm performance. These inconsistencies stem from a lack of controlled evaluation frameworks and analysis of models’ internals to isolate *when* and *why* modality interactions support or undermine reasoning. We address this gap through a logic-grounded evaluation framework that categorizes multimodal reasoning into six interaction patterns, varying how facts are distributed across modalities and logically combined. Empirically, additional modalities enhance reasoning only when they provide independent and sufficient reasoning paths, while redundant or chained entailment support often hurts performance. Besides, models recognize cross-modal facts reliably and always reason on text effectively. Moreover, reasoning degrades in three systematic ways: weaker modalities drag down overall performance, conflicts bias preference toward certain modalities, and joint signals from different modalities fail to be integrated effectively. Therefore, we identify two core failures: *task-composition bottleneck*, where recognition and reasoning cannot be jointly executed in one pass, and *fusion bottleneck*, where early integration introduces bias. For further investigation, we find that attention patterns fail to encode fact usefulness, but a simple two-step prompting (recognize then reason) restores performance, confirming the task-composition bottleneck. Moreover, modality identity remains recoverable in early layers, and softening attention in early fusion improves reasoning, highlighting biased fusion as another failure mode. Overall, our findings show that integration, not perception, is the main barrier to multimodal reasoning, suggesting composition-aware training and early fusion control as promising directions.

1 INTRODUCTION

Multimodal large language models (MLLMs) extend traditional language models beyond text to incorporate additional modalities such as vision and audio (Li et al., 2025c; Xu et al., 2025; Yu et al., 2025; Abouelenin et al., 2025). By integrating complementary signals, MLLMs can form richer and more grounded representations of the world. Text offers structured and abstract information, audio encodes temporal and prosodic cues, and images convey spatial and visual context, together capturing facets of meaning that no single modality can express in isolation (Clark & Brennan, 1991; Mayer, 2002). Through such cross-modal integration, MLLMs aim to overcome the limitations of unimodal systems, enabling more robust understanding, stronger perceptual grounding, and support for more complex reasoning (Li et al., 2023; Bie et al., 2025; Raza et al., 2025; Coburn et al., 2025).

While MLLMs hold promise for enhanced reasoning by integrating diverse signals, the precise influence of additional modalities remains unclear, especially under complex reasoning scenarios. Existing studies offer conflicting observations: some report that incorporating vision or audio can improve model performance (Li et al., 2023; Guan et al., 2024; Fu et al., 2025), while others suggest that additional modalities introduce interference or confusion (Bie et al., 2025; He et al., 2025; Hou et al., 2025). However, these findings are often anecdotal or domain-specific, lacking a unified framework to systematically assess when and how multimodal input contributes to or undermines reasoning (Gupta et al., 2024; Coburn et al., 2025; Hao et al., 2025; Li et al., 2025b; Bi et al., 2025). In particular, it still remains unclear under what conditions additional signals strengthen reasoning,

054 add little, or actively impede it (Wu et al., 2025; Zhang et al., 2025). As a result, the role of modality
 055 interaction in reasoning, whether beneficial or detrimental, remains underexplored.
 056

057 Moreover, most evaluations treat MLLMs as black-box systems, emphasizing external performance
 058 while leaving their internal mechanisms poorly understood (Liu et al., 2024b; Liang et al., 2023; Li
 059 et al., 2024). Even when empirical patterns emerge, such as degraded reasoning with added modalities,
 060 they are rarely accompanied by interpretability analyses that examine how models internally encode
 061 modality identity, assess evidence relevance, or perform cross-modal integration (Peng et al., 2025;
 062 Sinha et al., 2024; Yu et al., 2024; Wadekar et al., 2024). One contributing factor could lie in how
 063 these models are trained: current MLLMs are typically optimized using alignment-style objectives
 064 that pair vision or audio with text through paired supervision, contrastive learning, or instruction
 065 tuning (Zhao et al., 2024; Lin et al., 2024; Xie & Wu, 2024; Jiang et al., 2025). These objectives
 066 prioritize perceptual matching over cognitive composition, reinforcing shallow correlations rather
 067 than fostering deeper reasoning. As a result, while MLLMs often perform well on perception-heavy
 068 tasks (Li et al., 2023; Liu et al., 2024a), they struggle to generalize when reasoning demands flexible
 069 integration of multimodal information. Without interpreting the internal representations and fusion
 behaviors, it remains difficult to pinpoint where these limitations arise or how they might be overcome.

070 Prior evaluations report mixed effects of adding modalities to reasoning because they rarely control
 071 *where* decision-relevant facts appear or *how* those facts must be logically combined. We focus on
 072 *logical reasoning using information from multiple modalities*: how models use cross-modal facts to
 073 infer answers, rather than simple perception of unimodal content. To make effects measurable, we
 074 introduce six canonical interaction types (§ 2), grounded in propositional logic, that jointly vary (i)
 075 where the crucial facts are placed across modalities and (ii) how those facts must be combined to solve
 076 the task: *Equivalence* (\equiv , redundant encoding), *Alternative* (\vee , distinct but individually sufficient
 077 paths), *Entailment* (\rightarrow , chained support across modalities), *Independence* (\emptyset , a single modality carries
 078 the relevant fact), *Contradictory* (\oplus , mutually exclusive conclusions), and *Complementary* (\wedge , jointly
 079 necessary pieces). Instantiated through controlled, synthetic multiple choice reasoning tasks, this
 080 framework allows us to assess not only *when* added modalities that help or hurt, but also *why*.

081 To understand *when* added modalities help or hurt reasoning, we analyze performance across our six
 082 interactions. *Alternative* yields slight gains: added modalities help when they provide independent,
 083 individually sufficient reasoning paths. *Equivalence* offers no benefit in the presence of a strong
 084 modality (e.g., text), suggesting that redundant perceptual support rarely improves high unimodal
 085 performance. *Entailment* consistently degrades accuracy, showing that splitting multi-hop reasoning
 086 chains across modalities makes inference brittle (§ 3.2). *Independence* reveals performance bias,
 087 where reasoning accuracy heavily depends on which modality carries the decisive fact. *Contradictory*
 088 exposes preference bias, as models could favor certain modalities when inputs conflict. *Complementary*
 089 highlights a weakness in fusion, where models struggle to integrate necessary signals jointly
 090 (§ 3.2). Moreover, across all settings, text-only baselines approach ceiling and models reliably
 091 recognize facts across modalities. Therefore, we identify two core bottlenecks: (i) *task composition*:
 092 models struggle to jointly perform recognition and reasoning when information is split across modal-
 093 ities; and (ii) *multi-source fusion*: models lack robust mechanisms to select, weight, and combine
 heterogeneous information, leading to performance, preference, and fusion biases (§ 3.4).

094 To explain *why* these bottlenecks arise, we probe the internal behaviors of MLLMs in a controlled
 095 setting. First, although models recognize facts well and reason effectively (in text), their internal
 096 attention patterns fail to encode *usefulness* (i.e., distinguishing relevant facts from distractors).
 097 As a result, accuracy drops when recognition and reasoning must be composed within a single
 098 step, revealing a *task-composition bottleneck*. Explicitly decoupling the two stages through a two-
 099 step prompt substantially alleviates this issue (§ 4.1). Second, models preserve modality identity
 100 throughout processing, with the strongest signal concentrated in early decoder layers, but this
 101 preservation leads to biased weighting across modalities. Targeted interventions in these layers (e.g.,
 102 softening early attention via increased temperature) significantly improve cross-modal reasoning,
 103 whereas modifications at later layers have little effect. This pattern confirms a *fusion bottleneck*,
 104 where biased early integration prevents balanced use of multimodal evidence (§ 4.2). Together, these
 105 insights indicate that additional modalities yield limited and often inconsistent benefits, with failures
 106 rooted not in perception but in integration. This calls for models that incorporate composition-aware
 107 training, supervision for evidence selection, and architectural mechanisms for early fusion control, so
 that extra modalities become assets for reasoning rather than sources of interference.




Figure 1: **Multimodal logical reasoning setup and evaluation pipeline.** (a) **Logical reasoning example:** a single-step deduction where the fact “Bob is curious” and the rule “Curious people are purple” entail the conclusion “Bob is purple.” (b) **Modality renderings:** the same fact is rendered as text (“Bob is curious”), as audio via neural TTS, and as a schematic visual using graph visualization. (c) **Evaluation prompt pattern:** the model receives modality-specific fact blocks (text, audio, vision), followed by the rule set and the question with multiple-choice options; the model outputs the predicted answer.

2 CATEGORIZATION: MODALITY INTERACTION IN REASONING

This section introduces our evaluation framework for multimodal logical reasoning. We first describe the general task setup: how facts, rules, and questions are constructed and how facts are rendered across modalities (§ 2.1). We then introduce details of the six canonical interaction types, which specify how useful information is distributed across modalities (§ 2.2). These settings allow us to systematically assess whether and how models integrate, ignore, or prioritize cross-modal evidence.

2.1 MULTIMODAL LOGICAL REASONING SETUP

Fig. 1 illustrates our reasoning task setup. Each instance consists of a set of facts, a set of rules (always in text), and a multiple-choice question. Facts convey information through different modalities: text, audio, or vision, and we control which modality (Fig. 1.b).

(a) Logical reasoning example. To isolate modality interaction, we adopt a simplified single-step reasoning setting inspired by Clark et al. (2020). Full details of the data construction are in App. B.1. For example, given the fact “Bob is curious” and the rule “Curious people are purple,” the model should infer “Bob is purple.” This setup avoids multi-hop complexity and directly tests the model’s ability to identify and utilize relevant information across modalities.

(b) Three input modalities. To minimize confounds from low-level perception, we encode each fact in three controlled modalities (see Fig. 2b): (i) a short text sentence (e.g., “Bob is curious”), (ii) audio synthesized via neural TTS,¹ and (iii) a schematic visual rendered using graph visualization.² These controlled renderings ensure interpretability and reduce variability due to acoustic or visual complexity, allowing us to focus on reasoning and modality integration.

(c) Evaluation prompt pattern. Each prompt presents a randomized set of fact blocks from different modalities, followed by the text-based rule set and a multiple-choice question. To assess robustness and bias, we also inject *noisy facts* (irrelevant distractors) into the input. In the next subsection, we define six modality interaction types that vary how decision-relevant facts are distributed, enabling us to test different forms of cross-modal fusion, redundancy, and conflict.

2.2 INTERACTION TYPES

Equivalence (\equiv). All modalities redundantly encode the same fact (see Fig. 4 in App. B.3 for more details). For example, “Erin is friendly” appears in vision, audio, and text. This tests whether redundancy helps or harms reasoning, revealing models’ ability to aggregate repeated evidence.

Alternative (\vee). Each modality presents a different fact, but all facts independently satisfy a disjunctive rule (Fig. 5). For example, the rule “Friendly person is clean. Purple person is clean. Red person is clean.” is matched by “Erin is friendly”, “Erin is person”, and “Erin is red”, in separate

¹We use CosyVoice2 TTS to convert text to speech.

²We use GraphViz to generate simple entity-attribute diagrams.

162 modalities. This setting assesses whether models can leverage distinct but semantically aligned
 163 reasoning paths.

164 **Entailment (\rightarrow).** Facts are distributed across modalities to form a multi-hop reasoning chain (Fig. 6).
 165 For example, “*Erin is bouncy*” \rightarrow “*Erin is bright*” \rightarrow “*Erin is friendly*” \rightarrow “*Erin is purple*”, with
 166 each step in a different modality. Only the final fact directly supports the answer. This setting probes
 167 models’ capacity for chained inference across modalities.
 168

169 **Independence (\emptyset).** Only one modality contains the decision-relevant fact, while the others include
 170 unrelated distracting facts (Fig. 7). For example, only the vision modality presents “*Erin is friendly*,”
 171 while text and audio contain irrelevant attributes. This setting tests per-modality reasoning and
 172 robustness to irrelevant signals.

173 **Contradictory (\oplus).** Each modality leads to a different conclusion (Fig. 8). For example, text implies
 174 “*Erin is clean*”, vision implies “*Bob is purple*”, and audio implies “*Erin is tasty*”. This interaction
 175 reveals the model’s default preference when faced with conflicting evidence across modalities.

176 **Complementary (\wedge).** Each modality contributes a fact required for a conjunctive rule (Fig. 9). For
 177 instance, “*Erin is friendly*”, “*Erin is purple*”, and “*Erin is red*”, are distributed across modalities, and
 178 the rule “*If a person is friendly and purple and red, then the person is clean.*” must be applied. This
 179 interaction evaluates the ability to integrate information across modalities for multi-source reasoning.
 180

181 3 EVALUATION OF MODALITY INTERACTIONS: PROS AND CONS

182 We now present a systematic evaluation of MLLMs across the six canonical interaction types. This
 183 section first outlines our experimental setup, then investigates when multiple modalities *help* or *hurt*
 184 reasoning, and finally synthesizes key bottlenecks revealed by the results.
 185

186 3.1 PREPARATION

187 We begin by describing the models, prompting strategy, decoding procedure, and evaluation metric
 188 that together form a controlled testbed for analyzing modality interactions.
 189

190 **Models.** We evaluate four recent open-source MLLMs that support at least three input modalities
 191 and generate text outputs. *Baichuan-Omni-1.5d (7B)* (Li et al., 2025c, Baichuan) is designed for
 192 efficient and balanced multimodal reasoning. *Qwen2.5-Omni (7B)* (Xu et al., 2025, Qwen) is a
 193 state-of-the-art model that handles text, vision, audio, and video, with streaming outputs, serving as a
 194 strong open baseline for multi-source reasoning. *MiniCPM-o-2.6 (8B)* (Yao et al., 2024, MiniCPM)
 195 processes text, vision, and audio, and is optimized for real-time multimodal streaming and on-device
 196 deployment. *Phi-4 Multimodal (5.6B)* (Abouelenin et al., 2025, Phi4) extends the Phi family to vision
 197 and audio, emphasizing compactness and efficiency over scale.³ This suite covers both high-capacity
 198 and lightweight systems, enabling comparison of modality interaction across diverse architectures.
 199

200 **Prompt design and decoding.** We use a unified prompt format across all models to ensure fair
 201 comparison. Each prompt includes a system instruction, a set of fact blocks in random modality
 202 order (text, vision, audio), a series of textual reasoning rules, and a four-way multiple-choice question
 203 (Fig. 1c). To encourage step-by-step reasoning, we insert concise CoT hints while minimizing
 204 behavioral interference. Decoding follows HuggingFace defaults with greedy sampling to produce
 205 stable outputs, from which the final answer is automatically extracted. All models are evaluated on
 206 the same synthetic dataset using identical prompt templates. More details are provided in App. B.3.
 207

208 **Evaluation metric.** We report accuracy as the primary evaluation metric. We evaluate the reasoning
 209 performance by multiple-choice question answering format with four options, and the model’s
 210 selected answer is automatically extracted from its output response. Since there are four options, the
 211 random guessing would yield a baseline accuracy of 25%. To ensure robustness, each experiment is
 212 conducted on 1,000 synthetic instances per condition.
 213

214 ³We refer to these models by shortened names in all experiments.
 215

216 3.2 DO MULTIPLE MODALITIES HELP REASONING?
217218 One motivation for using MLLMs is to enhance reasoning by incorporating additional information
219 from multiple modalities. But does adding useful input from another modality always help? To
220 explore this, we evaluate three controlled interaction types designed to probe distinct modes of
221 cross-modal benefit: redundancy, optionality, and composition. Results are summarized in Tab. 1,
222 with full breakdowns in Apps. C.1 to C.3.223 **Setup.** In *Equivalence*, the same decisive fact is redundantly placed in all modalities, testing whether
224 repetition reinforces reasoning. In *Alternative*, each modality contains a distinct but individually
225 sufficient fact for solving the problem, allowing multiple independent reasoning paths. In *Entailment*,
226 a reasoning chain ($A \rightarrow B \rightarrow C \rightarrow \text{Answer}$) is split across modalities, requiring integration of
227 cross-modal premises for successful inference. To assess the added value of multimodal input, we
228 compare each multimodal setting to its unimodal baselines, where only one decisive fact is in one
229 modality (text, vision, or audio) and the others are omitted. From this comparison we can directly
230 analyze the additional value brought by additional information in extra modalities.
231232 Table 1: **Does Multimodality Help Reasoning?** Accuracy (%) and performance deltas (Δ) relative to unimodal
233 baselines across three interaction types: *Equivalence* (redundant facts across modalities), *Alternative* (independ-
234 ent reasoning paths), and *Entailment* (multi-hop chains split across modalities, with final-step facts in V/A/T
235 respectively). *Alternative* settings slightly boost performance, *Equivalence* yields marginal decrease (compared
236 to text), while *Entailment* causes notable accuracy drops on reasoning.

237 238 Accuracy (%)	239 Multimodal ($\equiv, \vee, \rightarrow$)			240		
	241 Equivalence $_{\Delta V, \Delta A, \Delta T}$	242 Alternative $_{\Delta V, \Delta A, \Delta T}$	243 Entailment: V$_{\Delta V}$, A$_{\Delta A}$, T$_{\Delta T}$	244	245	246
Baichuan	84.8 _{5.4↑, 9.8↑, 11.1↓}	97.6 _{19.6↑, 17.8↑, 0.3↑}	79.5 _{2.0↓}	75.6 _{6.4↓}	80.7 _{13.6↓}	
Qwen	98.9 _{2.6↑, 4.5↑, 0.9↑}	100.0 _{3.7↑, 6.1↑, 2.6↑}	78.4 _{15.7↓}	86.6 _{8.2↓}	83.9 _{12.8↓}	
MiniCPM	94.8 _{5.4↑, 5.2↑, 0.2↓}	99.1 _{7.1↑, 8.0↑, 2.9↑}	81.8 _{11.4↓}	80.0 _{12.0↓}	88.4 _{6.8↓}	
Phi4	84.1 _{25.3↑, 23.9↑, 12.5↓}	97.9 _{20.3↑, 26.3↑, 1.0↑}	73.0 _{2.2↓}	69.3 _{0.7↓}	79.7 _{18.0↓}	
Average	90.7 _{9.7↑, 10.9↑, 5.7↓}	98.7 _{12.7↑, 14.8↑, 1.7↑}	78.2 _{7.8↓}	77.9 _{7.1↓}	83.2 _{12.8↓}	

247 **Findings.** In *Equivalence*, models show marginal gains when the decisive fact is in vision (+9.7%)
248 or audio (+10.9%), but performance drops when the fact is already in text (-5.7%). This suggests
249 that redundancy is only helpful when the original modality is weak. In *Alternative*, consistent im-
250 provements are observed across all modalities (+12.7% vision, +14.8% audio, +1.7% text), indicating
251 that semantically independent reasoning paths are successfully leveraged. However, *Entailment*
252 leads to substantial drops in accuracy across all modalities (-7.8% vision, -7.1% audio, -12.8% text),
253 highlighting the difficulty of cross-modal multi-hop composition.254 **Observation 1.** Multimodal input improves reasoning only when it contributes additional, se-
255 manticly independent reasoning paths. In contrast, redundant information provides little benefit,
256 particularly when a strong modality (text) is already sufficient, and distributing multi-step reasoning
257 chains across modalities often reduces accuracy. These results suggest that *the core bottleneck in*
258 *multimodal reasoning lies not in recognizing facts*, since individual modalities suffice in many cases.259 3.3 DO MULTIPLE MODALITIES HURT REASONING?
260261 While certain forms of multimodal input can aid reasoning (§ 3.2), adding modalities could also
262 introduce errors. In this section, we evaluate three controlled settings: *Independence*, *Contradictory*,
263 and *Complementary*: to identify specific failure modes where multiple modalities degrade reasoning.
264

265 3.3.1 MODALITY PERFORMANCE BIAS: INDEPENDENCE

266 We first explore whether models exhibit consistent reasoning *performance* across modalities.
267268 **Setup.** In the *Independence* setting, a decisive fact appears in one modality (text, vision, or audio),
269 while the remaining modalities contain only distractors. We compare multimodal reasoning, where

270 facts are distributed across modalities, to unimodal baselines, where all facts (both decisive and
 271 distracting) are presented within a single modality.
 272

273 **Findings.** As shown in Tab. 2 (full results can be
 274 found in App. C.4), models
 275 perform best in the text-only
 276 condition (94.45% on average),
 277 but accuracy drops sharply
 278 to 70.29% when facts are
 279 distributed across modalities.
 280 This is well below text-only
 281 performance but above vision-
 282 only or audio-only baselines,
 283 confirming that weaker modal-
 284 ities introduce noise when
 285 combined with stronger ones.
 286

Table 2: **Performance on the *Independence* interaction.** Each instance includes one decisive fact placed in a single modality, while the others contain distractors. Multimodal reasoning accuracy falls between the best (text) and worst (vision) unimodal conditions, suggesting that modality inconsistency introduces error when aggregating information across modalities.

Accuracy (%)	Unimodal			Multimodal (\emptyset) $_{\Delta V, \Delta A, \Delta T}$
	V	A	T	
Baichuan	60.2	72.0	94.8	67.6 _{7.4↑, 4.4↓, 27.2↓}
Qwen	73.3	94.3	95.5	75.2 _{1.9↑, 19.1↓, 20.3↓}
MiniCPM	77.6	83.7	91.2	78.7 _{1.1↑, 5.0↓, 12.5↓}
Phi4	49.9	48.9	96.3	59.7 _{9.8↑, 10.8↑, 36.6↓}
Average	65.3	74.7	94.5	70.3 _{5.0↑, 4.4↓, 24.2↓}

287 **Observation 2.** Unequal reasoning capabilities across modalities, what we refer to as *performance*
 288 *bias*, contribute significantly to degraded multimodal reasoning. When weaker modalities are added,
 289 they can dilute or confuse the signal from stronger ones, like text.
 290

291 3.3.2 MODALITY PREFERENCE BIAS: CONTRADICTORY

292 We next ask whether models exhibit internal *preferences* for certain modalities.
 293

294 **Setup.** In the *Contradictory* setting, each modality
 295 provides a distinct and individually sufficient reason-
 296 ing path, but the answer options are mutually exclu-
 297 sive.⁴ This setup exposes which modality a model
 298 relies on when conflicting information is presented.
 299 Crucially, preference here refers to *selection behavior*
 300 under conflict, not standalone performance.
 301

302 **Results.** As shown in Tab. 3, models display clear
 303 preference patterns: Baichuan favors vision-based
 304 answers (49.0%), Qwen tends toward audio (44.6%),
 305 and both MiniCPM and Phi4 prefer text (49.0% and
 306 46.1%). These choices are often misaligned with the
 307 models’ unimodal strengths, suggesting implicit biases in modality selection under conflicting input.
 308

309 **Observation 3.** In addition to performance bias, MLLMs also suffer from *preference bias*: when
 310 modalities conflict, models favor certain modalities instead of strong modalities, often inconsistently
 311 with the actual performance. This misalignment introduces further risk in multimodal reasoning.
 312

313 3.3.3 MODALITY FUSION BIAS: COMPLEMENTARY

314 Finally, we examine whether models can *fuse* complementary evidence across modalities when each
 315 input is necessary for inference.
 316

317 **Setup.** In *Complementary* setting, each modality contains one of three facts that are jointly required
 318 to solve the reasoning task. Unlike prior settings with a single decisive fact, here all facts must be
 319 composed across modalities. We compare this condition to unimodal baselines where all three facts
 320 are provided within a one modality.
 321

322 ⁴Note that one of the four answer options is always incorrect by design, so the sum of selection ratios across
 323 modalities does not equal 100%.

324 **Results.** In Tab. 4, all models perform *worse* in the multimodal setup than in any unimodal condition, even when 325 confined to weaker modalities 326 like vision. If biased performance 327 is the only issue, multimodal accuracy 328 should lie between the best and worst 329 unimodal conditions. Instead, 330 distributing complementary 331 facts across modalities introduces 332 a new failure mode: 333 models are unable to compose 334 multiple weak signals into a 335 coherent reasoning chain. 336

340 **Observation 4.** MLLMs struggle to integrate complementary information across modalities, even 341 when all inputs are individually comprehensible. This reveals that beyond biased performance and 342 preference, there is a third failure mode in multimodal reasoning: *fusion bias*. 343

344 3.4 FROM OBSERVATIONS TO BOTTLENECKS

346 The preceding evaluations reveal consistent patterns in the way models handle multimodal reasoning. 347 We now synthesize these findings to identify the underlying bottlenecks. 348

349 **Observation 5.** Across all settings, the best performance consistently comes from the text-only 350 baseline, often approaching near-perfect accuracy. This shows that models can already perform 351 logical reasoning reliably when inputs are centralized in a single strong modality. Combined with 352 **Observation 1** (models can recognize facts across modalities), this indicates that the core weakness 353 lies not in perception or reasoning in isolation, but in how these components are combined. 354

355 **Bottleneck 1: Task Composition.** The conjunction of **Observation 1** (strong recognition) and 356 **Observation 5** (strong unimodal reasoning) points to a first core bottleneck: *task composition*. 357 Models falter when recognition and reasoning must be performed jointly across modalities. That is, 358 while they can detect facts and apply reasoning rules when each task is isolated, performance drops 359 sharply when these steps must be integrated within a single inference pass. 360

361 **Bottleneck 2: Multi-Source Fusion.** **Observations 2–4** collectively point to a second, orthogonal 362 failure point: the inability to *fuse* information from multiple modalities in a reliable and unbiased 363 manner. First, performance bias (**Observation 2**) shows that weak modalities dilute reasoning when 364 mixed with stronger ones. Second, preference bias (**Observation 3**) reveals that models often favor 365 certain modalities under conflict, even when those modalities underperform in isolation. Third, 366 fusion bias (**Observation 4**) shows that models fail to integrate complementary information spread 367 across modalities, even when all inputs are individually comprehensible. These findings suggest that 368 MLLMs lack robust internal mechanisms for selecting, weighting, and composing evidence from 369 heterogeneous sources, leading to systematic failures in multimodal reasoning. 370

371 4 INTERPRETATION: BOTTLENECKS OF MULTIMODAL REASONING

372 We now investigate the underlying causes of multimodal reasoning failures by probing the internal 373 mechanisms of MLLMs. Guided by the two key bottlenecks identified in § 3: *task composition* and 374 *multi-source fusion*, we analyze how models represent modality and information usefulness, and 375 whether targeted interventions can mitigate these bottlenecks.⁵ 376


377 ⁵In this section, we conduct our analysis under the *Independence* setting to avoid multi-hop confounds and 378 ensure interpretability. We select three representative models, Baichuan, Qwen, and MiniCPM, based on their 379 differing modality preferences identified in § 3.3.2, which favor vision, audio, and text, respectively. 380

381 **Table 4: Performance on the *Complementary* interaction.** Each modality 382 provides one necessary fact, requiring to integrate all three to get the answer. 383 The multimodal reasoning accuracy is lower than any unimodal condition, 384 indicating that performance degradation stems not only from modality in- 385 consistency but also from a true cross-modal composition bottleneck. 386

Accuracy (%)	Unimodal			Multimodal ($\wedge_{\Delta V, \Delta A, \Delta T}$)
	V	A	T	
Baichuan	50.5	59.4	87.7	30.2 _{20.3 ↓, 29.2 ↓, 57.5 ↓}
Qwen	87.5	98.8	98.8	49.9 _{37.6 ↓, 48.9 ↓, 48.9 ↓}
MiniCPM	74.8	89.3	92.4	48.8 _{26.0 ↓, 40.5 ↓, 43.6 ↓}
Phi4	80.0	82.2	99.6	79.1 _{0.9 ↓, 3.1 ↓, 20.5 ↓}
Average	73.2	82.4	94.6	52.0_{21.2 ↓, 30.4 ↓, 42.6 ↓}

378 4.1 TASK COMPOSITION: ANALYSIS AND IMPROVEMENT
379

380 We begin with **Bottleneck 1**, which concerns the model’s inability to compose recognition and
381 reasoning across modalities within a single inference step. While MLLMs can perceive facts from
382 different modalities and reason over them in isolation, their performance degrades when these
383 abilities must be integrated. We investigate this bottleneck by analyzing internal attention patterns
384 and exploring whether prompting strategies can mitigate the failure.



393 Figure 2: **Attention probing and reasoning performance.** (a) Modality probing for information usefulness
394 shows moderate accuracy, suggesting models cannot clearly distinguish useful from distractor facts. (b) Although
395 models excel in fact recognition and text-only reasoning, their performance drops significantly on multimodal
396 reasoning, indicating that the key limitation lies in composing recognition and reasoning across modalities.
397

398 **Probing Information Usefulness.** We first assess whether models can internally distinguish useful
399 facts from distractors. A linear probe is trained on decoder attention distributions: for each modality,
400 we compute mean attention over all generated tokens and use these vectors to classify whether each
401 fact is relevant for reasoning (more probing details are in App. B.2). Results (Fig. 2a) show that the
402 attention patterns do not reliably signal semantic usefulness, indicating that the models struggle to
403 prioritize decision-relevant content based on attention alone.

404 **Interpretation Results.** To isolate where failures occurs, we conduct two controlled diagnostics: (1)
405 a formal *recognition test*, where models identify facts across modalities without requiring additional
406 reasoning (see Fig. 10), and (2) a *two-step prompting* setup, where recognition and reasoning are
407 separated across prompts (see Fig. 11). As shown in Fig. 2b, models perform near-perfectly on
408 recognition and maintain high accuracy on unimodal reasoning. However, when both recognition and
409 reasoning are combined in a single multimodal prompt, accuracy drops substantially, confirming that
410 the core failure lies in the integration of these two capabilities.

412 **Improving Composition via Two-Step Prompting.** In the two-step prompting strategy, models
413 first extracting all facts, then reason over them, substantially improving performance across all models.
414 This shows that the failure arises not from deficiencies in perception or reasoning alone, but from
415 their joint composition within a single inference step. By decoupling these processes and providing a
416 more explicit recognition goal, the task-composition bottleneck is effectively alleviated.

418 **Takeaway.** These findings highlight a core weakness in current MLLMs: despite strong perception
419 and reasoning abilities in isolation, they lack mechanisms to integrate these steps across modalities.
420 This reflects a broader limitation in training objectives, which emphasize shallow alignment rather
421 than compositional inference. Prompt-level task decomposition offers a simple yet effective remedy.

423 4.2 MODALITY FUSION: ANALYSIS AND IMPROVEMENT
424

425 We now address **Bottleneck 2**, which concerns the model’s difficulty in fusing information across
426 modalities. Even when individual facts are recognized correctly, reasoning often fails due to biased
427 or ineffective modality integration. To better understand this bottleneck, we analyze how modality
428 identity is internally represented and whether early fusion contributes to systematic errors.

429 **Probing Modality Identity.** We assess whether modality type (text, vision, audio) is preserved in
430 the model’s internal representations. A logistic regression classifier is trained on attention-derived
431 features: for each input fact, similarly, we compute average attention from all generated tokens and




Figure 3: **Modality probing based on attention patterns.** (a) All models achieve perfect probe accuracy in predicting the modality using attention patterns. (b) For Qwen, linear probe weights show that modality information is primarily captured in the first four layers. (c) Attention manipulation in different 4 layers (by adjusting head temperature from 0.4 to 1.8), where performance significantly improves in the early 4 layers.

flatten this into a feature vector. The classifier predicts the modality of each fact. As shown in Fig. 3a, modality identity is perfectly recoverable, indicating that even after fusion, the model maintains a strong internal signal of the input modality.

Interpretation Results. To locate where modality fusion occurs, we visualize layer-wise probe weights. Fig. 3b shows that the first four decoder layers carry the strongest modality signal, suggesting that fusion predominantly occurs early in the language module. Beyond this point, modalities appear to be processed more uniformly.

Improving Fusion via Attention Manipulation. Motivated by this early-fusion pattern, we modify attention behavior by adjusting the softmax temperature from 0.4 to 1.8 in the first four decoder layers (Early), with the default set to 1.0. As shown in Fig. 3c, this simple intervention of increasing the temperature for early layers yields significant improvements in reasoning accuracy by encouraging more balanced attention across modalities. In contrast, adjusting the temperature in middle or late layers has little effect, supporting the causal role of early fusion in downstream reasoning outcomes.

Takeaways. These findings confirm that while modality identity is well preserved, early-stage fusion introduces systematic biases that impair reasoning. A lightweight causal intervention, reshaping early attention distributions, can significantly enhance multimodal integration, highlighting the importance of fusion dynamics over perceptual bottlenecks.

Summary of Interpretation Findings. Across above interpretation analyses, we find that failures in multimodal reasoning stem not from deficiencies in perception or unimodal reasoning, but from weak *compositional integration* and *cross-modal fusion*. Although MLLMs retain modality-specific signals and can recognize facts reliably, they often default to shallow alignment behavior rather than selective integration. Without architectural biases or training objectives that explicitly encourage multimodal composition, these models remain brittle in complex reasoning tasks.

5 CONCLUSION

This work presents a systematic study of how MLLMs integrate information across modalities for logical reasoning. We introduce a logic-driven evaluation framework with six canonical interaction types, enabling controlled analysis of when additional modalities help or hinder reasoning. Our results reveal a consistent pattern: modalities help only when they contribute independent, sufficient reasoning paths, while redundancy or cross-modal chaining often degrades performance. Text-only baselines already approach ceiling accuracy, underscoring that the key bottleneck is not perception but integration. Through probing and causal interventions, we identified two core bottlenecks: a *task-composition bottleneck*, where recognition and reasoning cannot be reliably combined in a single inference step, and a *fusion bottleneck*, where early-layer integration introduces modality bias. Simple remedies such as two-step prompting and attention temperature adjustments alleviate these issues, highlighting the importance of designing models and objectives that explicitly support evidence selection and unbiased fusion. We hope our framework and findings inspire future work toward composition-aware training and architecture choices that transform added modalities from sources of interference into assets for reasoning.

486
487**ETHICS STATEMENT**488
489
490
491
492

This research does not involve human participants, private or sensitive data, or applications with foreseeable negative societal impact. All datasets employed are publicly available and widely used within the vision–language and reasoning research communities. We adhere to standard best practices in data handling, model evaluation, and reporting, and our study complies fully with the ICLR Code of Ethics.

493
494
495**REPRODUCIBILITY STATEMENT**496
497
498
499

We have taken care to make our work reproducible. The paper and appendix provide detailed descriptions of the models, datasets, and experimental setups. Code, data generation scripts, and evaluation protocols are included in the supplementary material and will be released publicly upon publication to facilitate replication and further research.

500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

540 REFERENCES
541

542 Abdelrahman Abouelenin, Atabak Ashfaq, Adam Atkinson, Hany Awadalla, Nguyen Bach, Jianmin
543 Bao, Alon Benhaim, Martin Cai, Vishrav Chaudhary, Congcong Chen, Dong Chen, Dongdong
544 Chen, Jun-Kun Chen, Weizhu Chen, Yen-Chun Chen, Yi-ling Chen, Qi Dai, Xiyang Dai, Ruchao
545 Fan, Mei Gao, Min Gao, Amit Garg, Abhishek Goswami, Junheng Hao, Amr Hendy, Yuxuan
546 Hu, Xin Jin, Mahmoud Khademi, Dongwoo Kim, Young Jin Kim, Gina Lee, Jinyu Li, Yunsheng
547 Li, Chen Liang, Xihui Lin, Zeqi Lin, Mengchen Liu, Yang Liu, Gilsinia Lopez, Chong Luo,
548 Piyush Madan, Vadim Mazalov, Arindam Mitra, Ali Mousavi, Anh Nguyen, Jing Pan, Daniel
549 Perez-Becker, Jacob Platin, Thomas Portet, Kai Qiu, Bo Ren, Liliang Ren, Sambuddha Roy,
550 Ning Shang, Yelong Shen, Saksham Singhal, Subhajit Som, Xia Song, Tetyana Sych, Praneetha
551 Vaddamanu, Shuohang Wang, Yiming Wang, Zhenghao Wang, Haibin Wu, Haoran Xu, Weijian
552 Xu, Yifan Yang, Ziyi Yang, Donghan Yu, Ishmam Zabir, Jianwen Zhang, Li Lyra Zhang, Yunan
553 Zhang, and Xiren Zhou. Phi-4-mini technical report: Compact yet powerful multimodal language
554 models via mixture-of-loras. *CoRR*, abs/2503.01743, 2025. doi: 10.48550/ARXIV.2503.01743.
555 URL <https://doi.org/10.48550/arXiv.2503.01743>.

556 Jing Bi, Junjia Guo, Susan Liang, Guangyu Sun, Luchuan Song, Yunlong Tang, Jinxi He, Jiarui Wu,
557 Ali Vosoughi, Chen Chen, and Chenliang Xu. VERIFY: A benchmark of visual explanation and
558 reasoning for investigating multimodal reasoning fidelity. *CoRR*, abs/2503.11557, 2025. doi: 10.
559 48550/ARXIV.2503.11557. URL <https://doi.org/10.48550/arXiv.2503.11557>.

560 Fuqing Bie, Shiyu Huang, Xijia Tao, Zhiqin Fang, Leyi Pan, Junzhe Chen, Min Ren, Liuyu Xiang,
561 and Zhaofeng He. Omniplay: Benchmarking omni-modal models on omni-modal game playing,
562 2025. URL <https://arxiv.org/abs/2508.04361>.

563 Herbert H. Clark and Susan E. Brennan. Grounding in communication. In Lauren B. Resnick,
564 John M. Levine, and Stephanie D. Teasley (eds.), *Perspectives on socially shared cognition*, pp.
565 127–149. American Psychological Association, 1991. doi: 10.1037/10096-006. URL <https://doi.org/10.1037/10096-006>.

566 Peter Clark, Oyvind Tafjord, and Kyle Richardson. Transformers as soft reasoners over language.
567 In Christian Bessiere (ed.), *Proceedings of the Twenty-Ninth International Joint Conference on
568 Artificial Intelligence, IJCAI 2020*, pp. 3882–3890. ijcai.org, 2020. doi: 10.24963/IJCAI.2020/537.
569 URL <https://doi.org/10.24963/ijcai.2020/537>.

570 Bruce Coburn, Jiangpeng He, Megan E. Rollo, Satvinder S. Dhaliwal, Deborah A. Kerr, and Fengqing
571 Zhu. Evaluating large multimodal models for nutrition analysis: A benchmark enriched with
572 contextual metadata. *CoRR*, abs/2507.07048, 2025. doi: 10.48550/ARXIV.2507.07048. URL
573 <https://doi.org/10.48550/arXiv.2507.07048>.

574 Zhihao Du, Yuxuan Wang, Qian Chen, Xian Shi, Xiang Lv, Tianyu Zhao, Zhifu Gao, Yexin Yang,
575 Changfeng Gao, Hui Wang, et al. Cosyvoice 2: Scalable streaming speech synthesis with large
576 language models. *arXiv preprint arXiv:2412.10117*, 2024.

577 Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin,
578 Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, and Rongrong Ji. MME: A comprehensive evaluation
579 benchmark for multimodal large language models. *CoRR*, abs/2306.13394, 2023. doi: 10.48550/
580 ARXIV.2306.13394. URL <https://doi.org/10.48550/arXiv.2306.13394>.

581 Chaoyou Fu, Yuhai Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang,
582 Chenyu Zhou, Yunhang Shen, Mengdan Zhang, Peixian Chen, Yanwei Li, Shaohui Lin, Sirui
583 Zhao, Ke Li, Tong Xu, Xiawu Zheng, Enhong Chen, Caifeng Shan, Ran He, and Xing Sun.
584 Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal LLMs in video
585 analysis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025,
586 Nashville, TN, USA, June 11–15, 2025*, pp. 24108–24118. Computer Vision Foundation / IEEE,
587 2025. doi: 10.1109/CVPR52734.2025.02245. URL [https://openaccess.thecvf.com/
590 content/CVPR2025/html/Fu_Video-MME_The_First-Ever_Comprehensive_
591 Evaluation_Benchmark_of_Multi-modal_LLMs_in_CVPR_2025_paper.html](https://openaccess.thecvf.com/
588 content/CVPR2025/html/Fu_Video-MME_The_First-Ever_Comprehensive_
589 Evaluation_Benchmark_of_Multi-modal_LLMs_in_CVPR_2025_paper.html).

592 Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang,
593 Lichang Chen, Furong Huang, Yaser Yacoob, Dinesh Manocha, and Tianyi Zhou. Hallusion-
594 bench: An advanced diagnostic suite for entangled language hallucination and visual illusion

594 in large vision-language models. In *IEEE/CVF Conference on Computer Vision and Pattern*
 595 *Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pp. 14375–14385. IEEE, 2024.
 596 doi: 10.1109/CVPR52733.2024.01363. URL <https://doi.org/10.1109/CVPR52733.2024.01363>.
 597

598 Himanshu Gupta, Shreyas Verma, Ujjwala Anantheswaran, Kevin Scaria, Mihir Parmar, Swa-
 599 roop Mishra, and Chitta Baral. Polymath: A challenging multi-modal mathematical reason-
 600 ing benchmark. *CoRR*, abs/2410.14702, 2024. doi: 10.48550/ARXIV.2410.14702. URL
 601 <https://doi.org/10.48550/arXiv.2410.14702>.
 602

603 Yunzhuo Hao, Jiawei Gu, Huichen Will Wang, Linjie Li, Zhengyuan Yang, Lijuan Wang, and
 604 Yu Cheng. Can mllms reason in multimodality? EMMA: an enhanced multimodal reasoning
 605 benchmark. *CoRR*, abs/2501.05444, 2025. doi: 10.48550/ARXIV.2501.05444. URL
 606 <https://doi.org/10.48550/arXiv.2501.05444>.
 607

608 Abdelhamid Haouhat, Slimane Bellaouar, Attia Nehar, and Hadda Cherroun. Modality influence in
 609 multimodal machine learning. *CoRR*, abs/2306.06476, 2023. doi: 10.48550/ARXIV.2306.06476.
 610 URL <https://doi.org/10.48550/arXiv.2306.06476>.
 611

612 Yixiao He, Haifeng Sun, Pengfei Ren, Jingyu Wang, Huazheng Wang, Qi Qi, Zirui Zhuang, and Jing
 613 Wang. Evaluating and mitigating object hallucination in large vision-language models: Can they
 614 still see removed objects? In Luis Chiruzzo, Alan Ritter, and Lu Wang (eds.), *Proceedings of the*
 615 *2025 Conference of the Nations of the Americas Chapter of the Association for Computational*
 616 *Linguistics: Human Language Technologies, NAACL 2025 - Volume 1: Long Papers, Albuquerque,*
 617 *New Mexico, USA, April 29 - May 4, 2025*, pp. 6841–6858. Association for Computational
 618 Linguistics, 2025. doi: 10.18653/V1/2025.NAACL-LONG.349. URL <https://doi.org/10.18653/v1/2025.naacl-long.349>.
 619

620 Yifan Hou, Buse Giledereli, Yilei Tu, and Mrinmaya Sachan. Do vision-language models really
 621 understand visual language?, 2025. URL <https://arxiv.org/abs/2410.00193>.
 622

623 Songtao Jiang, Yan Zhang, Ruizhe Chen, Tianxiang Hu, Yeying Jin, Qinglin He, Yang Feng, Jian
 624 Wu, and Zuozhu Liu. Modality-fair preference optimization for trustworthy mllm alignment, 2025.
 625 URL <https://arxiv.org/abs/2410.15334>.
 626

627 Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan
 628 Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer.
 629 *Trans. Mach. Learn. Res.*, 2025, 2025a. URL <https://openreview.net/forum?id=zKv8qULV6n>.
 630

631 Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. Seed-bench: Bench-
 632 marking multimodal llms with generative comprehension. *CoRR*, abs/2307.16125, 2023. doi: 10.
 633 48550/ARXIV.2307.16125. URL <https://doi.org/10.48550/arXiv.2307.16125>.
 634

635 Linjie Li, Mahtab Bigverdi, Jiawei Gu, Zixian Ma, Yinuo Yang, Ziang Li, Yejin Choi, and Ranjay
 636 Krishna. Unfolding spatial cognition: Evaluating multimodal models on visual simulations. *CoRR*,
 637 abs/2506.04633, 2025b. doi: 10.48550/ARXIV.2506.04633. URL <https://doi.org/10.48550/arXiv.2506.04633>.
 638

639 Yadong Li, Jun Liu, Tao Zhang, Tao Zhang, Song Chen, Tianpeng Li, Zehuan Li, Lijun Liu, Lingfeng
 640 Ming, Guosheng Dong, Da Pan, Chong Li, Yuanbo Fang, Dongdong Kuang, Mingrui Wang,
 641 Chenglin Zhu, Youwei Zhang, Hongyu Guo, Fengyu Zhang, Yuran Wang, Bowen Ding, Wei Song,
 642 Xu Li, Yuqi Huo, Zheng Liang, Shusen Zhang, Xin Wu, Shuai Zhao, Linchu Xiong, Yozhen Wu,
 643 Jiahui Ye, Wenhao Lu, Bowen Li, Yan Zhang, Yaqi Zhou, Xin Chen, Lei Su, Hongda Zhang,
 644 Fuzhong Chen, Xuezhen Dong, Na Nie, Zhiying Wu, Bin Xiao, Ting Li, Shunya Dang, Ping Zhang,
 645 Yijia Sun, Jincheng Wu, Jinjie Yang, Xionghai Lin, Zhi Ma, Kegeng Wu, Jia Li, Aiyuan Yang,
 646 Hui Liu, Jianqiang Zhang, Xiaoxi Chen, Guangwei Ai, Wentao Zhang, Yicong Chen, Xiaoqin
 647 Huang, Kun Li, Wenjing Luo, Yifei Duan, Lingling Zhu, Ran Xiao, Zhe Su, Jiani Pu, Dian Wang,
 648 Xu Jia, Tianyu Zhang, Mengyu Ai, Mang Wang, Yujing Qiao, Lei Zhang, Yanjun Shen, Fan Yang,
 649 Miao Zhen, Yijie Zhou, Mingyang Chen, Fei Li, Chenzheng Zhu, Keer Lu, Yaqi Zhao, Hao Liang,
 650 Youquan Li, Yanzhao Qin, Linzhuang Sun, Jianhua Xu, Haoze Sun, Mingan Lin, Zenan Zhou,
 651 and Weipeng Chen. Baichuan-omni-1.5 technical report. *CoRR*, abs/2501.15368, 2025c. doi: 10.
 652 48550/ARXIV.2501.15368. URL <https://doi.org/10.48550/arXiv.2501.15368>.
 653

648 Yizhi Li, Ge Zhang, Yinghao Ma, Ruibin Yuan, Kang Zhu, Hangyu Guo, Yiming Liang, Jiaheng
 649 Liu, Jian Yang, Siwei Wu, Xingwei Qu, Jinjie Shi, Xinyue Zhang, Zhenzhu Yang, Xiangzhou
 650 Wang, Zhaoxiang Zhang, Zachary Liu, Emmanouil Benetos, Wenhao Huang, and Chenghua Lin.
 651 Omnibench: Towards the future of universal omni-language models. *CoRR*, abs/2409.15272, 2024.
 652 doi: 10.48550/ARXIV.2409.15272. URL <https://doi.org/10.48550/arXiv.2409.15272>.
 653

654 Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian
 655 Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby
 656 Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas,
 657 Drew A. Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu
 658 Yao, Jue Wang, Keshav Santhanam, Laurel J. Orr, Lucia Zheng, Mert Yüksekgönül, Mirac Suzgun,
 659 Nathan Kim, Neel Guha, Niladri S. Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan
 660 Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard,
 661 Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta
 662 Koreeda. Holistic evaluation of language models. *Trans. Mach. Learn. Res.*, 2023, 2023. URL
 663 <https://openreview.net/forum?id=iO4LZibEqW>.

664 Ji Lin, Hongxu Yin, Wei Ping, Pavlo Molchanov, Mohammad Shoeybi, and Song Han. VILA: on
 665 pre-training for visual language models. In *IEEE/CVF Conference on Computer Vision and Pattern
 666 Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pp. 26679–26689. IEEE, 2024.
 667 doi: 10.1109/CVPR52733.2024.02520. URL <https://doi.org/10.1109/CVPR52733.2024.02520>.
 668

669 Ronghao Lin and Haifeng Hu. Missmodal: Increasing robustness to missing modality in multimodal
 670 sentiment analysis. *Trans. Assoc. Comput. Linguistics*, 11:1686–1702, 2023. doi: 10.1162/TACL_A_00628. URL
 671 https://doi.org/10.1162/tacl_a_00628.
 672

673 Xin Liu, Yichen Zhu, Jindong Gu, Yunshi Lan, Chao Yang, and Yu Qiao. Mm-safetybench:
 674 A benchmark for safety evaluation of multimodal large language models. In Ales Leonardis,
 675 Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gü̈l Varol (eds.), *Computer Vi-
 676 sion - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024,
 677 Proceedings, Part LVI*, volume 15114 of *Lecture Notes in Computer Science*, pp. 386–403.
 678 Springer, 2024a. doi: 10.1007/978-3-031-72992-8_22. URL https://doi.org/10.1007/978-3-031-72992-8_22.
 679

680 Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi
 681 Wang, Conghui He, Ziwei Liu, Kai Chen, and Dahu Lin. Mmbench: Is your multi-modal model
 682 an all-around player? In Ales Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten
 683 Sattler, and Gü̈l Varol (eds.), *Computer Vision - ECCV 2024 - 18th European Conference, Milan,
 684 Italy, September 29-October 4, 2024, Proceedings, Part VI*, volume 15064 of *Lecture Notes in
 685 Computer Science*, pp. 216–233. Springer, 2024b. doi: 10.1007/978-3-031-72658-3_13. URL
 686 https://doi.org/10.1007/978-3-031-72658-3_13.
 687

687 Richard E Mayer. Multimedia learning. In *Psychology of learning and motivation*, volume 41, pp.
 688 85–139. Elsevier, 2002.

689 Jing Peng, Yucheng Wang, Bohan Li, Yiwei Guo, Hankun Wang, Yangui Fang, Yu Xi, Haoyu Li,
 690 Xu Li, Ke Zhang, Shuai Wang, and Kai Yu. A survey on speech large language models for
 691 understanding, 2025. URL <https://arxiv.org/abs/2410.18908>.
 692

693 Shaina Raza, Ashmal Vayani, Aditya Jain, Aravind Narayanan, Vahid Reza Khazaie, Syed Raza
 694 Bashir, Elham Dolatabadi, Gias Uddin, Christos Emmanouilidis, Rizwan Qureshi, and Mubarak
 695 Shah. Vldbench evaluating multimodal disinformation with regulatory alignment, 2025. URL
 696 <https://arxiv.org/abs/2502.11361>.
 697

697 Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy P. Lillicrap, Jean-
 698 Baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, Ioannis
 699 Antonoglou, Rohan Anil, Sebastian Borgeaud, Andrew M. Dai, Katie Millican, Ethan Dyer,
 700 Mia Glaese, Thibault Sottiaux, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu,
 701 James Molloy, Jilin Chen, Michael Isard, Paul Barham, Tom Hennigan, Ross McIlroy, Melvin
 Johnson, Johan Schalkwyk, Eli Collins, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha

702 Goel, Clemens Meyer, Gregory Thornton, Zhen Yang, Henryk Michalewski, Zaheer Abbas,
 703 Nathan Schucher, Ankesh Anand, Richard Ives, James Keeling, Karel Lenc, Salem Haykal,
 704 Siamak Shakeri, Pranav Shyam, Aakanksha Chowdhery, Roman Ring, Stephen Spencer, Eren
 705 Sezener, and et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens
 706 of context. *CoRR*, abs/2403.05530, 2024. doi: 10.48550/ARXIV.2403.05530. URL <https://doi.org/10.48550/arXiv.2403.05530>.

708 Yusuf Shihata. Gated recursive fusion: A stateful approach to scalable multimodal transformers.
 709 *CoRR*, abs/2507.02985, 2025. doi: 10.48550/ARXIV.2507.02985. URL <https://doi.org/10.48550/arXiv.2507.02985>.

712 Anoop K. Sinha, Chinmay Kulkarni, and Alex Olwal. Levels of multimodal interaction. In Hayley
 713 Hung, Catharine Oertel, Mohammad Soleymani, Theodora Chaspari, Hamdi Dibeklioglu, Jainendra
 714 Shukla, and Khiet P. Truong (eds.), *Companion Proceedings of the 26th International Conference
 715 on Multimodal Interaction, ICMI Companion 2024, San Jose, Costa Rica, November 4-8, 2024*, pp.
 716 51–55. ACM, 2024. doi: 10.1145/3686215.3690153. URL <https://doi.org/10.1145/3686215.3690153>.

718 Christopher Thomas, Yipeng Zhang, and Shih-Fu Chang. Fine-grained visual entailment. In Shai
 719 Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (eds.),
 720 *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022,
 721 Proceedings, Part XXXVI*, volume 13696 of *Lecture Notes in Computer Science*, pp. 398–416.
 Springer, 2022. doi: 10.1007/978-3-031-20059-5_23. URL https://doi.org/10.1007/978-3-031-20059-5_23.

724 Shakti N. Wadekar, Abhishek Chaurasia, Aman Chadha, and Eugenio Culurciello. The evolution of
 725 multimodal model architectures. *CoRR*, abs/2405.17927, 2024. doi: 10.48550/ARXIV.2405.17927.
 URL <https://doi.org/10.48550/arXiv.2405.17927>.

728 Huyu Wu, Meng Tang, Xinhuan Zheng, and Haiyun Jiang. When language overrules: Revealing text
 729 dominance in multimodal large language models, 2025. URL <https://arxiv.org/abs/2508.10552>.

731 Zhifei Xie and Changqiao Wu. Mini-omni2: Towards open-source gpt-4o with vision, speech and
 732 duplex capabilities. *CoRR*, abs/2410.11190, 2024. doi: 10.48550/ARXIV.2410.11190. URL
 733 <https://doi.org/10.48550/arXiv.2410.11190>.

734 Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang,
 735 Yang Fan, Kai Dang, Bin Zhang, Xiong Wang, Yunfei Chu, and Junyang Lin. Qwen2.5-omni
 736 technical report. *CoRR*, abs/2503.20215, 2025. doi: 10.48550/ARXIV.2503.20215. URL <https://doi.org/10.48550/arXiv.2503.20215>.

738 Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li,
 739 Weilin Zhao, Zihui He, et al. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint
 740 arXiv:2408.01800*, 2024.

742 Haofei Yu, Zhengyang Qi, Lawrence Jang, Russ Salakhutdinov, Louis-Philippe Morency, and
 743 Paul Pu Liang. Mmoe: Enhancing multimodal models with mixtures of multimodal interaction
 744 experts. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the
 745 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami,
 746 FL, USA, November 12-16, 2024*, pp. 10006–10030. Association for Computational Linguistics,
 747 2024. doi: 10.18653/V1/2024.EMNLP-MAIN.558. URL <https://doi.org/10.18653/v1/2024.emnlp-main.558>.

749 Tianyu Yu, Haoye Zhang, Qiming Li, Qixin Xu, Yuan Yao, Da Chen, Xiaoman Lu, Ganqu Cui,
 750 Yunkai Dang, Taiwen He, Xiaocheng Feng, Jun Song, Bo Zheng, Zhiyuan Liu, Tat-Seng Chua,
 751 and Maosong Sun. RLAIF-V: open-source AI feedback leads to super GPT-4V trustworthiness.
 752 In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville,
 753 TN, USA, June 11-15, 2025*, pp. 19985–19995. Computer Vision Foundation / IEEE, 2025. doi:
 754 10.1109/CVPR52734.2025.01861. URL https://openaccess.thecvf.com/content/CVPR2025/html/Yu_RLAIF-V_Open-Source_AI_Feedback_Leads_to_Super_GPT-4V_Trustworthiness_CVPR_2025_paper.html.

756 Xiang Yue, Yuansheng Ni, Tianyu Zheng, Kai Zhang, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu
757 Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin,
758 Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen.
759 **MMMU**: A massive multi-discipline multimodal understanding and reasoning benchmark for
760 expert AGI. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024,*
761 *Seattle, WA, USA, June 16-22, 2024*, pp. 9556–9567. IEEE, 2024. doi: 10.1109/CVPR52733.2024.
762 00913. URL <https://doi.org/10.1109/CVPR52733.2024.00913>.

763 Yu Zhang, Jinlong Ma, Yongshuai Hou, Xuefeng Bai, Kehai Chen, Yang Xiang, Jun Yu, and Min
764 Zhang. Evaluating and steering modality preferences in multimodal large language model. *CoRR*,
765 [abs/2505.20977](https://arxiv.org/abs/2505.20977), 2025. doi: 10.48550/ARXIV.2505.20977. URL <https://doi.org/10.48550/arXiv.2505.20977>.

766 Fei Zhao, Taotian Pang, Chunhui Li, Zhen Wu, Junjie Guo, Shangyu Xing, and Xinyu Dai. Alignngpt:
767 Multi-modal large language models with adaptive alignment capability, 2024.

770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

810
811

A RELATED WORK

812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829

Multimodal Large Language Models (MLLMs). Recent MLLMs extend text-only LLMs to handle images, audio, and often video, with two broad design trends. On the end-to-end omni side, Qwen2.5-Omni processes text–vision–audio–video with streaming IO via a block-wise encoder stack and a Thinker–Talker architecture for joint text/speech generation, achieving state-of-the-art results on omni benchmarks (Xu et al., 2025). MiniCPM-o-2.6 follows a similarly integrated pipeline (SigLIP + Whisper + ChatTTS + Qwen2.5 backbone) to enable real-time speech and multimodal live streaming on resource-constrained devices (Yu et al., 2025). Baichuan-Omni proposes a two-stage scheme, multimodal alignment, then multitask fine-tuning, to support concurrent inputs of image, audio, video, and text in an open 7B model (Li et al., 2025c). In contrast, adapter-based approaches extend a compact LLM with modality-specific LoRA routers (e.g., Phi-4 Multimodal), delivering strong vision–audio performance while retaining efficiency (Abouelenin et al., 2025). Beyond open models, frontier systems like Gemini emphasize extremely long multimodal context across documents, video, and audio (Reid et al., 2024). Complementary open MLLM lines (LLaVA-OneVision (Li et al., 2025a), VILA-1.5 (Lin et al., 2024)) focus on unifying single-image, multi-image, and video scenarios and scaling families across 3B–40B parameters, respectively. Finally, efforts like Mini-Omni2 target GPT-4o-style, real-time visual–audio assistants, underscoring the field’s push toward unified, low-latency multimodal interaction (Xie & Wu, 2024). Recent surveys also review progress in Speech LLMs, formally defining speech understanding and analyzing Speech-LLM architectures, training, and evaluation through a structured taxonomy (Peng et al., 2025).

830
831
832
833
834
835
836
837
838
839
840
841
842
843
844

MLLM Evaluation. A rapidly growing ecosystem of benchmarks now probes different facets of the MLLM ability. General-purpose suites such as MMBench (Liu et al., 2024b), MME (Fu et al., 2023), and SEED-Bench (Li et al., 2023) emphasize breadth and stable, objective MCQ evaluation across perception and language understanding, including text-rich images and video. Reasoning-centric datasets like MMMU push models toward college-level, multi-disciplinary problem solving with heterogeneous visual artifacts (charts, diagrams, tables), exposing large gaps in expert-level multimodal reasoning (Yue et al., 2024). Complementing capability scores, robustness-focused evaluations, including POPE (He et al., 2025) for object hallucination and HallusionBench (Guan et al., 2024) for language-vs-vision conflicts, and visual illusions-diagnose systematic failure modes and modality biases that broad benchmarks can obscure. Beyond static images, Video-MME (Fu et al., 2025) targets temporal understanding, while MM-SafetyBench (Liu et al., 2024a) stress-tests safety in multimodal settings. These lines reveal a pattern: While breadth benchmarks track steady gains, targeted diagnostics consistently uncover modality dominance, hallucination, and fusion brittleness, motivating frameworks (like ours) that isolate and measure information interaction in reasoning.

845
846
847
848
849
850
851
852
853
854
855
856
857
858

Recognition vs. Reasoning. While many benchmarks and papers demonstrate that MLLMs are adept at recognizing objects, attributes, diagrams, or patterns, few dissect how well these models reason with those recognized elements, especially when evidence must be drawn from multiple modalities. Benchmarks like VERIFY (Bi et al., 2025) highlight that perception (recognition of visual input) is often less challenging for MLLMs than reasoning, models often fail when inference, explanation, or abstract relationships are required. STARE (Li et al., 2025b) similarly reveals that while simple spatial or 2D transformations can be handled, tasks that need multistep visual simulation or 3D understanding are far more error-prone. POLYMATHE (Gupta et al., 2024) and EMMA (Hao et al., 2025) further push into domains (math, diagrams, cross-modal reasoning) where recognition alone is insufficient; these tasks expose gaps when models must combine or interpret recognized information rather than just identify it. Our work builds on this line by explicitly separating and quantifying fact recognition, modality recognition, and reasoning performance, and by introducing interaction types (e.g., independence, entailment, complementary) to isolate how modalities help or hurt when reasoning is required.

859
860
861
862
863

Information Interaction and Logical Relations. A number of studies have implicitly or explicitly considered how information from multiple modalities interacts, especially in logical or semantic reasoning contexts. For example, Thomas et al. (2022) asks whether textual hypotheses entail, contradict, or are neutral with respect to images, and breaks down hypothesis claims, which correspond roughly to “entailment” and “contradiction” in our framework. Meanwhile, Sinha et al. (2024) proposes a qualitative taxonomy that includes redundancy and synergy between modalities, similar to our

“equivalence” and “complementary” types, but without formal logical rules or controlled tasks to isolate their effects. The recent MMOE work (Yu et al., 2024) directly trains separate experts for redundancy, uniqueness, and synergy, which is closely aligned with some of our interaction types (independent, alternative, complementary), yet still coarser. Architectural analyses (Wadekar et al., 2024) show different fusion strategies (early vs late vs cross-attention) that implement interactions implicitly rather than measuring them via logical operators. Our work builds on and extends this prior art by defining six fine-grained logical interaction types (independent, equivalence, alternative, entailment, complementary, contradictory) and embedding them into synthetic reasoning tasks with rules, facts, and contrastive answer options. This allows us to evaluate not just whether models fuse modalities, but also how and when different interaction patterns boost or degrade reasoning performance.

Fusion Mechanisms and Modality Dominance. Research on fusion architectures and modality dominance has rapidly advanced, revealing how design choices often tip the scale in favor of one modality, typically text. For instance, Wu et al. (2025) introduces quantitative metrics to measure how strongly models depend on text, showing that fusion architectures and token redundancy play key roles in producing text dominance. Similarly, Zhang et al. (2025) analyze modality dominance in MLLMs and attribute it to imbalances in scaling, alignment, and representation, showing that text often overwhelms other modalities not only due to token abundance, but also architectural biases. Their controlled ablations further highlight that dominance emerges from systemic training choices rather than dataset artifacts. Meanwhile, Haouhat et al. (2023) uses masked-modality ablation across tasks (sentiment, emotion, disease detection) to confirm that, in many settings, non-text modalities add little when text is present. Architectural solutions have been proposed: Shihata (2025) implements gating and sequential fusion to control how each modality contributes. Lin & Hu (2023) further show that handling missing modalities and aligning feature/geometric spaces can reduce dominance bias. Our work links to this by not only observing modality dominance in reasoning tasks, but also by dissecting when and why dominance arises via logical interactions and attention probing. We contribute by explicitly measuring both performance (modality competence) and preference (which modality is used when conflicting or combined), under controlled logical reasoning settings.

B SUPPLEMENTARY SETUP

B.1 FACT CONSTRUCTION

We construct facts and rules following the format of Clark et al. (2020), using the data generation code from Liang et al. (2023).⁶ Below we describe the details of fact generation and rule generation.

Fact generation. Each fact consists of a *subject*, a *predicate*, and an *object*. The predicate is always the copula “is”; variability comes from the choice of subject and object. To keep the setting interpretable and avoid conflicts with commonsense priors, we prevent antonyms or synonyms from appearing together within the same instance. Details are as follows:

- **Subjects.** A subject is randomly sampled from three categories:
 - *Persons* (13 names): Alice, Bob, Carol, Dan, Erin, Frank, George, Harry, Iris, Jack, Kevin, Lance, Miller.
 - *Animals* (14 types): dog, cat, rabbit, mouse, tiger, lion, bear, squirrel, cow, panda, hedgehog, elephant, giraffe, hippo.
 - *Fruits* (15 types): apple, banana, orange, grape, strawberry, blueberry, watermelon, pineapple, mango, peach, cherry, pear, kiwi, lemon, plum.
- **Predicate.** Always “is” (e.g., “Bob *is* curious”).
- **Objects.** Objects are adjective attributes describing the subject. We use a pool of 34 attributes: young, soft, scary, hot, smart, clean, beautiful, red, blue, green, purple, boring, strong, happy, round, big, noisy, fast, sticky, bouncy, spiky, furry, bright, shiny, magical, striped, spotted, tasty, juicy, toxic, friendly, curious, loud, sleepy.

⁶The code is from this GitHub repository.

918 **Rule generation.** Rules are constructed following Liang et al. (2023), with a minor modification
 919 for brevity:
 920

921 • If the rule involves only a single attribute, we convert it into a compact nominal form to save tokens.
 922 Example: Original: “If a cow is weak, then the cow is small.” Transformed: “Weak cow is small.”
 923 • If the rule involves multiple attributes, we retain the full “if–then” format for clarity. Example: “If
 924 a person is smart and young, then the person is curious.”
 925

926 **Image & audio generation.** For the multimodal reasoning experiments, we require factual
 927 information in both visual and auditory formats. To generate images, we employ the Graphviz toolkit,
 928 which effectively converts structured data into clear, labeled diagrams suitable for model interpretation.
 929 For audio synthesis, we utilize CosyVoice 2 (Du et al. (2024)), one of the top Text-to-Speech (TTS)
 930 model renowned for its high consistency in timbre, achieving human-parity synthesis quality. Since
 931 the audio is solely intended to provide semantic information for subsequent logical reasoning within
 932 the MLLM, and we are not investigating the model’s advanced audio understanding capabilities, we
 933 prioritize accuracy over variability in the synthesized speech. Therefore, we use CosyVoice 2’s default
 934 settings to ensure the speech is clear and precise, without introducing unnecessary characteristics that
 935 could potentially interfere with the reasoning task.
 936

937 B.2 EXPERIMENT SETUP

938 **LLM inference setup.** During evaluation, all models are run in float16 precision. For models
 939 that support audio output, we disable this feature and only generate text. The maximum number of
 940 generated tokens is set to 1024, which is typically sufficient for the model to produce a complete
 941 response, as long as the model does not exhibit output repetition. To ensure coherent and consistent
 942 output, we set ‘do_sample=False’ for all models, which is also the default setting for these models
 943 and considered best practice.
 944

945 **Linear probing setup.** To construct the attention feature as the input of the probe, we extract and
 946 pool attentions. Specifically, for one data sample, the attention matrix shape should be $N \times L \times H \times O$.
 947 Here, N is the number of tokens representing one fact (in one modality); L is the number of layers in
 948 the model; H is the number of attention head in the model; and O is the number of generated tokens
 949 of the response for the model. Since O and N could vary in samples (or facts), we take the mean
 950 pooling in term of these two dimensions. The final attention feature for each sample is shaped as
 951 $L \times H$. We flatten the feature as one attention feature for one fact in one sample as the input for the
 952 probe classifier.
 953

954 For the linear probe classifier, we use a 5-fold cross-validation setup with GroupKFold to evaluate the
 955 models. The models are trained with Logistic Regression using a regularization strength of C=1.0,
 956 an l2 penalty, and class_weight="balanced" to handle class imbalances. The feature vectors are
 957 normalized using the standard score method, which standardizes the features to have zero mean and
 958 unit variance. We apply these configurations to the task at hand. These configurations ensure a robust
 959 and consistent evaluation across all models tested.
 960

961 B.3 PROMPT TEMPLATES


962 We list the prompt templates for all six types of interactions, the recognition task, and the two-step
 963 reasoning task. All models have the same input (template and data).
 964
 965
 966
 967
 968
 969
 970
 971

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

Prompt: input and output (Equivalence)

System: You are an assistant tasked with solving multiple-choice questions that require logical reasoning over the supplied knowledge diagrams. Use only the information explicitly given—do not rely on outside or commonsense knowledge. Read the question and given information, think step-by-step and answer the question. At the end of your answer, answer precisely in the format 'Answer: X' where X is the chosen letter A / B / C / D.

Input:



(Audio information: *Erin is friendly*)

Rules are as follows: Blue person is tasty. Red person is clean. If a person is smart or sleepy, then the person is curious. Friendly person is purple.

Which of the following options can be inferred based on the given facts and rules? A) Erin is curious. B) Erin is purple. C) Erin is tasty. D) Erin is clean.

Output: To determine which option can be inferred based on the given facts and rules, let's analyze the information step by step:

- Given Information: - Erin is friendly.
- Rules: - Blue person is tasty. - Red person is clean. - If a person is smart or sleepy, then the person is curious. - Friendly person is purple.
- Analysis: - From the rule ""Friendly person is purple,"" we can infer that since Erin is friendly, Erin must be purple.
- Checking the Options:
 - A) Erin is curious. The rule does not provide a direct link between being friendly and being curious. Therefore, we cannot infer this.
 - B) Erin is purple. This directly follows from the rule "Friendly person is purple."
 - C) Erin is tasty. The rule does not provide a direct link between being friendly and being tasty. Therefore, we cannot infer this.
 - D) Erin is clean. The rule does not provide a direct link between being friendly and being clean. Therefore, we cannot infer this.

Based on the analysis, the only option that can be inferred is:

Answer: B

Figure 4: Prompt template and model (Qwen) output (Equivalence).

1026

1027

1028 Prompt: input and output (Alternative)

1029


1030

1031 **System:** You are an assistant tasked with solving multiple-choice questions that
 1032 require logical reasoning over the supplied knowledge diagrams. Use only the
 1033 information explicitly given—do not rely on outside or commonsense knowledge.
 1034 Read the question and given information, think step-by-step and answer the question.
 1035 At the end of your answer, answer precisely in the format 'Answer: X' where X is
 1036 the chosen letter A / B / C / D.

1037

1038

1039

Input:

Erin is purple

(Audio information: *Erin is red*)

1050 **Rules are as follows:** Friendly person is clean. If a person is smart or sleepy, then the
 1051 person is curious. Purple person is clean. Blue Erin is tasty. Spotted Erin is beautiful.
 1052 Red person is clean.

1053 Which of the following options can be inferred based on the given facts and rules? A)
 1054 Erin is clean. B) Erin is tasty. C) Erin is beautiful. D) Erin is curious.

1055

1056

1057

Output:

1058 To determine which option can be inferred based on the given facts and rules, let's
 1059 analyze the information step by step:
 1060

1. Given Information:

- Erin is purple.
- Erin is red.
- Erin is friendly.

1061 2. Rules: - If a person is smart or sleepy, then the person is curious. - Blue Erin is
 1062 tasty. - Spotted Erin is beautiful. - Friendly person is clean. - Purple person is clean. -
 1063 Red person is clean.

1064 3. Analysis: - Since Erin is purple and red, and the rules state that purple and
 1065 red people are clean, we can infer that Erin is clean. - The rules do not provide
 1066 information about Erin being smart, sleepy, curious, beautiful, or tasty based on the
 1067 given facts.

1068 4. Conclusion: - The only inference we can make from the given information is that
 1069 Erin is clean.

1070 Answer: A

1071

1072

1073

1074

1075

1076

1077

1078

1079


Figure 5: Prompt template and model (Qwen) output (Alternative).

1080
1081
1082
1083
1084
1085
1086
1087
1088
1089

Prompt: input and output (Entailment)

1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104

Input:



Erin is bright



(Audio information: *Erin is bouncy*)

1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

Rules are as follows: If a person is smart or sleepy, then the person is curious. Red person is clean. Blue person is tasty. Bright person is friendly. Friendly person is purple. Bouncy person is bright.

Which of the following options can be inferred based on the given facts and rules? A) Erin is curious. B) Erin is tasty. C) Erin is purple. D) Erin is clean.

Output:

Let's analyze the given information step by step:

1. Erin is bouncy.
2. Bouncy Erin is bright.
3. Erin is friendly.
4. Friendly person is purple.
5. Bright Erin is friendly.

From the information, we can deduce the following:

- Since Erin is bouncy, Erin is bright.
- Since Erin is bright, Erin is friendly.
- Since Erin is friendly, Erin is purple.

Now, let's evaluate each option:

- A) Erin is curious. - This cannot be inferred because the rules do not connect being smart or sleepy to being curious.
- B) Erin is tasty. - This cannot be inferred because there is no information linking being bouncy, bright, or friendly to being tasty.
- C) Erin is purple. - This can be inferred because Erin is friendly, and friendly people are purple.
- D) Erin is clean. - This cannot be inferred because there is no information linking being bouncy, bright, or friendly to being clean.

Answer: C

Figure 6: Prompt template and model (Qwen) output (Entailment).

1134
1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184


1185

1186

1187

Prompt: input and output (Independence)

System: You are an assistant tasked with solving multiple-choice questions that require logical reasoning over the supplied knowledge diagrams. Use only the information explicitly given—do not rely on outside or commonsense knowledge. Read the question and given information, think step-by-step and answer the question. At the end of your answer, answer precisely in the format 'Answer: X' where X is the chosen letter A / B / C / D.

Input:

Dan is sleepy.

(Audio information: *Erin is spiky.*)

Rules are as follows: Blue person is tasty. Red person is clean. If a person is smart or sleepy, then the person is curious. Friendly person is purple.

Which of the following options can be inferred based on the given facts and rules? A) Erin is curious. B) Erin is purple. C) Erin is tasty. D) Erin is clean.

Output: To determine the correct answer, let's analyze the given information step by step:

1. Erin is friendly. 2. Friendly person is purple. 3. Blue person is tasty.

From the rules provided: - If a person is smart or sleepy, then the person is curious. - Red person is clean.

Since Erin is friendly and friendly people are purple, we can infer that Erin is purple. Now, let's evaluate the options:

A) Erin is curious. - This cannot be inferred from the given information.

B) Erin is purple. - This can be inferred from the given information.

C) Erin is tasty. - This cannot be inferred from the given information.

D) Erin is clean. - This cannot be inferred from the given information.

Answer: B

Figure 7: Prompt template and model (Qwen) output (Independence).

1188

1189

1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238


1239

1240

1241

Prompt: input and output (Contradictory)

System: You are an assistant tasked with solving multiple-choice questions that require logical reasoning over the supplied knowledge diagrams. Use only the information explicitly given—do not rely on outside or commonsense knowledge. Read the question and given information, think step-by-step and answer the question. At the end of your answer, answer precisely in the format 'Answer: X' where X is the chosen letter A / B / C / D.

Input:

Erin is red

(Audio information: *Erin is blue*)

Rules are as follows: Blue person is tasty. Red person is clean. If a person is smart or sleepy, then the person is curious. Friendly person is purple.

Which of the following options can be inferred based on the given facts and rules? A) Erin is curious. B) Erin is tasty. C) Erin is purple. D) Erin is clean.

Output:

Let's analyze the information step by step:

1. Erin is blue.
2. Blue people are tasty.
3. Erin is friendly.
4. Friendly people are purple.

From the rules, we can infer the following:

- Since Erin is blue, she is tasty.
- Since Erin is friendly, she is purple.

Now, let's evaluate each option:

- A) Erin is curious. - This cannot be inferred because the rules do not connect being blue or friendly to being curious.
- B) Erin is tasty. - This can be inferred because blue people are tasty.
- C) Erin is purple. - This can be inferred because friendly people are purple.
- D) Erin is clean. - This cannot be inferred because the rules do not connect being blue or friendly to being clean.

Based on the given facts and rules, the options that can be inferred are B and C. However, since the question asks for a single answer, we need to choose the one that is directly supported by the rules without additional assumptions.

Answer: B

Figure 8: Prompt template and model (Qwen) output (Contradictory).

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

1272

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293


1294

1295

Prompt: input and output (Complementary)

System: You are an assistant tasked with solving multiple-choice questions that require logical reasoning over the supplied knowledge diagrams. Use only the information explicitly given—do not rely on outside or commonsense knowledge. Read the question and given information, think step-by-step and answer the question. At the end of your answer, answer precisely in the format 'Answer: X' where X is the chosen letter A / B / C / D.

Input:



Erin is purple



(Audio information: *Erin is red*)

If a person is purple and red and not friendly, then the person is soft. If a person is friendly and purple and not red, then the person is big. If a person is friendly and red and not purple, then the person is scary. If a person is friendly and purple and red, then the person is clean.

Which of the following options can be inferred based on the given facts and rules? A) Erin is soft. B) Erin is scary. C) Erin is clean. D) Erin is big.

Output:

To determine which option can be inferred based on the given facts and rules, let's analyze the information step by step:

1. Erin is purple.
2. Erin is friendly.
3. Erin is red.

Now, let's apply the rules to Erin:

- Rule 1: If a person is purple and red and not friendly, then the person is soft. - Erin is purple and red and friendly, so this rule does not apply to Erin.
- Rule 2: If a person is friendly and purple and not red, then the person is big. - Erin is friendly, purple, and red, so this rule does not apply to Erin.
- Rule 3: If a person is friendly and red and not purple, then the person is scary. - Erin is friendly, red, and purple, so this rule does not apply to Erin.
- Rule 4: If a person is friendly and purple and red, then the person is clean. - Erin is friendly, purple, and red, so this rule applies to Erin.

Based on the analysis, the only rule that applies to Erin is Rule 4, which states that if a person is friendly and purple and red, then the person is clean.

Answer: C

Figure 9: Prompt template and model (Qwen) output (Complementary).

1296
1297

Prompt: input and output (Recognition)

1298

1299

1300

1301

1302

1303

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315


1316

1317

1318

1319

System: You are an assistant tasked with solving multiple-choice questions about knowledge diagrams. Use only the information explicitly given—do not rely on outside or commonsense knowledge. The facts are given in image, audio and text. Read the question and given information, and directly answer the question in the following format: 'Answer: X' where X is the chosen letter A / B / C / D.

Input:

Dan is sleepy.

(Audio information: *Erin is spiky.*)

Question: Which fact is mentioned in the given information in image, audio or text?

Options: A) Erin is sticky., B) Erin is friendly., C) Erin is scary., D) Erin is green.

Output:

B

Figure 10: Prompt template and model (Qwen) output for the fact recognition.

1320

1321

1322

1323

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

1350
1351

Prompt: input and output (Two-Step Reasoning: Step 1)

1352

1353

1354

1355

1356

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402


1403

System Prompt 1: You are an assistant tasked with solving multiple-choice questions about knowledge diagrams. Use only the information explicitly given—do not rely on outside or commonsense knowledge. The facts are given in image, audio and text. Read the question and given information, and directly answer the question.

System Prompt 2: You are an assistant tasked with solving multiple-choice questions that require logical reasoning over the supplied knowledge diagrams. Use only the information explicitly given—do not rely on outside or commonsense knowledge. Read the question and given information, think step-by-step and answer the question. At the end of your answer, answer precisely in the format 'Answer: X' where X is the chosen letter A / B / C / D.

1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

Input 1:



(Audio information: *Erin is spiky.*)

Question: Check the given information and list all the facts in the given image, audio and text, respectively.

Output 1:

Facts from the image: - Erin is spiky.

Facts from the audio: - Dan is sleepy.

Facts from the text: - Erin is friendly.

1404

1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

Prompt: input and output (Two-Step Reasoning: Step 2)

Input 2 (Including Input 1 + Output 1):

Rules are as follows: Blue person is tasty. Red person is clean. If a person is smart or sleepy, then the person is curious. Friendly person is purple.

Which of the following options can be inferred based on the given facts and rules? A) Erin is curious. B) Erin is purple. C) Erin is tasty. D) Erin is clean.

Output 2:

Let's analyze the information step by step:

1. Erin is friendly.
2. Erin is spiky.
3. Dan is sleepy.

From the rules provided:

- If a person is smart or sleepy, then the person is curious. - Friendly person is purple.
- Blue person is tasty. - Red person is clean.

Since Erin is friendly, according to the rules, Erin must be purple.

Now, let's evaluate the options:

A) Erin is curious. - We don't have enough information to infer this from the given facts and rules.

B) Erin is purple. - This is directly inferred from the rule that a friendly person is purple.

C) Erin is tasty. - We don't have enough information to infer this from the given facts and rules.

D) Erin is clean. - We don't have enough information to infer this from the given facts and rules.

Answer: B

Figure 11: Prompt template and model (Qwen) output for the two step reasoning.

1458 C SUPPLEMENTARY RESULTS

1460 C.1 MODALITY INTERACTION: EQUIVALENCE

1462 **Setup.** We now evaluate the *Equivalence* interaction, where the same decisive fact is redundantly
 1463 placed in all three modalities (vision, audio, text). The unimodal conditions serve as baselines, where
 1464 the only fact is presented in one modality. This setting examines whether redundant presentation
 1465 across modalities can reinforce fact recognition and improve reasoning robustness.

1466 Table 5: **Performance on the *Equivalence* interaction.** Each modality redundantly contains the same decisive
 1467 fact. Results show that adding equivalent facts from other modalities provides no consistent improvement,
 1468 suggesting that unimodal fact recognition is not the reasoning bottleneck.

1471 Accuracy (%)	1472 Unimodal			1473 Multimodal (\equiv)$_{\Delta V, \Delta A, \Delta T}$
	1474 V	1475 A	1476 T	1477
Baichuan	79.4	75.0	95.9	84.8 _{5.4↑, 9.8↑, 11.1↓}
Qwen	96.3	94.4	98.0	98.9 _{2.6↑, 4.5↑, 0.9↑}
MiniCPM	89.4	89.6	95.0	94.8 _{5.4↑, 5.2↑, 0.2↓}
Phi4	58.8	60.2	96.6	84.1 _{25.3↑, 23.9↑, 12.5↓}
Average	91.0	79.8	96.4	90.7 _{9.7↑, 10.9↑, 5.7↓}

1480 **Results.** As shown in Tab. 5, adding equivalent facts in vision and audio brings little to no gain over
 1481 the text-only setting, which already achieves near-perfect accuracy. In some cases (e.g., Baichuan
 1482 and Phi4), the multimodal setting even leads to a significant performance drop. This suggests that
 1483 additional modalities do not enhance reasoning and may introduce unnecessary interference. Since
 1484 models already recognize text-based facts effectively, improving recognition robustness does not
 1485 translate into better multimodal reasoning. The bottleneck appears to lie elsewhere, specifically, in
 1486 how the recognized facts are composed during reasoning.

1487 **Takeaway.** Adding redundant evidence across modalities fails to improve reasoning and can even
 1488 hurt performance. This suggests that text-based fact recognition is not the limiting factor. Instead,
 1489 cross-modal fusion and composition seem to be the main challenges in multimodal reasoning. *Since*
 1490 *simply reinforcing facts does not help, we next ask: can additional modalities help by introducing*
 1491 *new and independent reasoning paths?* We investigate this in the following setting.

1493 C.2 MODALITY INTERACTION: ALTERNATIVE

1495 **Setup.** In this setting, each modality contains a unique antecedent that is individually sufficient
 1496 to trigger the same rule and yield the correct conclusion. Thus, the model can arrive at the correct
 1497 answer by reasoning over any one of the modalities. This evaluates whether models can flexibly
 1498 leverage semantically diverse cues when multiple independent reasoning paths are available.

1499 Table 6: **Performance on the *Alternative* interaction.** Each modality provides an independent reasoning path,
 1500 and the correct answer can be inferred from any one of them. Multimodal accuracy improves slightly over the
 1501 text-only baseline, showing that models can benefit from semantically diverse cues across modalities. This
 1502 suggests that introducing complementary reasoning paths can help mitigate modality-specific limitations.

1504 Accuracy (%)	1505 Unimodal			1506 Multimodal (\vee)$_{\Delta V, \Delta A, \Delta T}$
	1507 V	1508 A	1509 T	1510
Baichuan	78.0	79.8	97.3	97.6 _{19.6↑, 17.8↑, 0.3↑}
Qwen	96.3	93.9	97.4	100.0 _{3.7↑, 6.1↑, 2.6↑}
MiniCPM	92.0	91.1	96.2	99.1 _{7.1↑, 8.0↑, 2.9↑}
Phi4	77.6	71.6	96.9	97.9 _{20.3↑, 26.3↑, 1.0↑}
Average	86.0	83.9	97.0	98.7 _{12.7↑, 14.8↑, 1.7↑}

1512 **Results.** As shown in Tab. 6, models achieve strong performance across unimodal settings, par-
 1513 ticularly in the text-only case. Notably, when all three sufficient premises are presented across
 1514 modalities, performance consistently improves across models. For example, Qwen improves from
 1515 97.4% (text-only) to 100.0%, and MiniCPM improves from 96.2% to 99.1%. Although the gain is
 1516 modest compared to the text baseline, the upward trend suggests that models can effectively integrate
 1517 multiple alternative cues to reinforce the reasoning process, even when distributed across modalities.
 1518

1519 **Takeaway.** These findings suggest that unlike redundant (Equivalence) setups, providing alternative
 1520 reasoning paths across modalities can support more robust reasoning. Multimodal information helps
 1521 when it offers diverse routes to the same conclusion, rather than simply reiterating or fragmenting
 1522 the required information. Since introducing alternative single-step reasoning cues across modalities
 1523 leads to measurable gains, it raises a natural follow-up question: can models also benefit from more
 1524 indirect, multi-hop cues, such as entailment chains spread across modalities? We explore this next.
 1525

1526 C.3 MODALITY INTERACTION: ENTAILMENT

1527 **Setup.** In the Entailment interaction, the model must perform multi-hop reasoning through a chain
 1528 of rules. Specifically, three facts form a reasoning chain: $A \rightarrow B \rightarrow C$, where the final answer is
 1529 entailed by C . While the decisive fact C is always placed in one fixed modality (vision, audio, or text),
 1530 the earlier support facts A and B are distributed across the other two modalities. This setting evaluates
 1531 whether the model can integrate indirect, cross-modal evidence to support the final-step reasoning.
 1532

1533 **Table 7: Performance on the Entailment interaction.** Each modality carries a fact needed for multi-hop
 1534 reasoning, with the final-step premise fixed in one modality (V/A/T). Models consistently perform worse in
 1535 multimodal setups compared to their unimodal baselines. This suggests that spreading multi-hop reasoning steps
 1536 across modalities introduces substantial integration errors, regardless of where the final step is placed.

1538 Accuracy (%)	1539 Unimodal			1540 Multimodal (\rightarrow): Final-Step Fact		
	1541 V	1542 A	1543 T	1544 $V_{\Delta V}$	1545 $A_{\Delta A}$	1546 $T_{\Delta T}$
Baichuan	81.5	82.0	94.3	79.5 _{2.0↓}	75.6 _{6.4↓}	80.7 _{13.6↓}
Qwen	94.1	94.8	96.7	78.4 _{15.7↓}	86.6 _{8.2↓}	83.9 _{12.8↓}
MiniCPM	93.2	92.9	95.2	81.8 _{11.4↓}	80.0 _{12.9↓}	88.4 _{6.8↓}
Phi4	75.2	70.0	97.7	73.0 _{2.2↓}	69.3 _{0.7↓}	79.7 _{18.0↓}
Average	86.0	84.9	96.0	78.2 _{7.8↓}	77.9 _{7.1↓}	83.2 _{12.8↓}

1547 **Results.** As shown in Tab. 7, all models experience significant performance drops in the multimodal
 1548 condition compared to their unimodal counterparts, regardless of which modality carries the final-step
 1549 fact. For instance, when the final-step fact is in text, we observe up to 12.8% drop compared to the
 1550 text-only baseline. This pattern holds consistently across models and configurations. These results
 1551 suggest that the cross-modal composition of reasoning chains introduces substantial integration errors,
 1552 even when the decisive premise remains in a strong reasoning modality.

1553 **Takeaway.** Unlike alternative information, which offers parallel reasoning paths, entailment infor-
 1554 mation provides indirect, chained support, which proves to be not only unhelpful but actively harmful
 1555 to reasoning performance. This highlights that current MLLMs struggle with cross-modal multi-hop
 1556 reasoning, even when all required facts are present. We have now evaluated three types of auxiliary
 1557 information: equivalence, alternative, and entailment, and found that only alternative reasoning paths
 1558 offer modest improvements. In contrast, redundant or indirect information often introduces confusion
 1559 and leads to performance degradation.

1561 C.4 MODALITY INTERACTION: INDEPENDENCE

1562 Tab. 8 presents detailed results under the *Independence* interaction, where each instance contains
 1563 a single decisive fact placed in one modality (vision, audio, or text), while the other modalities
 1564 contain only irrelevant distractors. We evaluate both unimodal reasoning (all facts are given in one
 1565 modality) and multimodal reasoning (the decisive fact is mixed with distractors across modalities).

1566 Table 8: **Performance on the *Independence* interaction with cross-modal distractors.** Each instance contains
 1567 one decisive fact placed in a specific modality (V: vision, A: audio, T: text), while the other modalities contain
 1568 only noisy facts. We report accuracy (%) for each unimodal condition as well as multimodal reasoning when
 1569 decisive facts are distributed.

1571 1572 1573 1574 1575 1576 1577 1578 1579 1580 1581 1582 1583 1584 1585 1586 1587 1588 1589 1590 1591 1592 1593 1594 1595 1596 1597 1598 1599 1600 1601 1602 1603 1604 1605 1606 1607 1608 1609 1610 1611 1612 1613 1614 1615 1616 1617 1618 1619	Accuracy (%)	1570 1571 1572 1573 1574 1575 1576 1577 1578 1579 1580 1581 1582 1583 1584 1585 1586 1587 1588 1589 1590 1591 1592 1593 1594 1595 1596 1597 1598 1599 1600 1601 1602 1603 1604 1605 1606 1607 1608 1609 1610 1611 1612 1613 1614 1615 1616 1617 1618 1619	1570 1571 1572 1573 1574 1575 1576 1577 1578 1579 1580 1581 1582 1583 1584 1585 1586 1587 1588 1589 1590 1591 1592 1593 1594 1595 1596 1597 1598 1599 1600 1601 1602 1603 1604 1605 1606 1607 1608 1609 1610 1611 1612 1613 1614 1615 1616 1617 1618 1619				
		V	A	T	V	A	T
Baichuan	60.2	72.0	94.8	74.3	53.5	74.9	67.6
Qwen	73.3	94.3	95.5	50.8	90.8	84.1	75.2
MiniCPM	77.6	83.7	91.2	66.8	78.2	91.0	78.7
Phi4	49.9	48.9	96.3	58.0	50.4	70.7	59.7
Average	65.3	74.7	94.5	62.5	68.2	80.2	70.3

1580 Across all models, text consistently yields the highest unimodal accuracy (average 94.45%), while
 1581 vision and audio vary significantly in performance. Interestingly, when the decisive fact remains in a
 1582 fixed modality but is surrounded by irrelevant facts from other modalities, we observe substantial
 1583 performance drops for vision and audio (e.g., Qwen drops from 73.3% to 50.8% when adding
 1584 distractors to vision). The “Random” column averages over all cross-modal settings with randomly
 1585 selected decisive modalities, showing that even a single irrelevant modality can degrade reasoning.

1587 LLM USAGE

1589 We used ChatGPT as a general-purpose assistant in preparing this paper. In particular, LLMs were
 1590 employed for grammar refinement, clarity improvements, LaTeX formatting, and debugging minor
 1591 code snippets. They were not involved in research ideation, experimental design, or the development
 1592 of theoretical contributions.