



Caught in the Story: Narrative Captivity in Multi-turn LLMs Conversation

Yuhe Wu, Guangyu Wang, Yujie Chen, Jiatong Zhang, Yuran Chen, Yutong Zhang, Xiyin Cheng, Wenpeng Cao, Zhuang Liu, Guang Zhang




香港中文大學(深圳)
The Chinese University of Hong Kong, Shenzhen



INTRODUCTION

- **The gap.** Moral advice often begins with one party's account. Existing evaluations focus on single exchanges or explicit persuasion, leaving the effects of ordinary, unfolding narration unclear.
- **Our question.** Can a model maintain independent judgment when it hears only one side of a conflict across multiple turns?
- **Narrative captivity.** We study a failure mode in which a model adopts the narrator's interpretation without adequately accounting for missing perspectives, even without explicit demands for agreement.



CONCLUSION & LIMITATIONS

LLM advisors remain vulnerable to one-sided, multi-turn narration. Our evaluation is limited to English and Chinese, five-turn dialogues, and binary stance judgments

DISCUSSION

- **Accumulating Accommodation.** GPT-5.5 and GLM-5.1 show declining disagreement and hedging across turns, suggesting that earlier concessions may reinforce the narrator's framing.
- **Mitigation Trade-offs.** Anti-sycophancy instructions and a third-person persona help both models. Structured reasoning helps GPT-5.5 but hurts GLM-5.1, while repeating user statements worsens both. Mitigation remains partial and model-dependent.

Mitigation Strategies: **M1** Anti-Sycophancy Instruction, **M2** Third-Person Narration, **M3** Chain-of-Thought, **M4** Context Recap

Model	Method	NH↑	NR↑
GPT-5.5	Base (No Mitigation)	0.275	0.650
	└ w/ M1: Anti-Sycophancy	0.454 (+0.179)	0.813 (+0.163)
	└ w/ M2: 3rd-Person Narration	0.331 (+0.056)	0.739 (+0.090)
	└ w/ M3: Chain-of-Thought	0.391 (+0.116)	0.693 (+0.043)
	└ w/ M4: Context Recap	0.238 (-0.037)	0.579 (-0.071)
GLM-5.1	Base (No Mitigation)	0.258	0.592
	└ w/ M1: Anti-Sycophancy	0.309 (+0.051)	0.662 (+0.070)
	└ w/ M2: 3rd-Person Narration	0.311 (+0.053)	0.688 (+0.096)
	└ w/ M3: Chain-of-Thought	0.183 (-0.069)	0.508 (-0.084)
	└ w/ M4: Context Recap	0.115 (-0.143)	0.130 (-0.462)

Family	Training Stage	Emotion		Fairness		Loyalty		Role Duty		Norms		Autonomy	
		NH↑	NR↑	NH↑	NR↑	NH↑	NR↑	NH↑	NR↑	NH↑	NR↑	NH↑	NR↑
Tulu3	Base	0.186	0.025	0.185	0.033	0.180	0.127	0.191	0.040	0.199	0.065	0.164	0.091
	└ +SFT	-0.011	+0.004	-0.067	-0.010	-0.085	-0.040	-0.072	+0.008	-0.099	+0.039	-0.054	+0.012
	└ +DPO	+0.005	+0.023	-0.017	+0.023	-0.028	+0.168	-0.009	+0.113	-0.023	+0.259	+0.004	+0.162
	└ +RLVR	+0.001	-0.007	+0.002	+0.003	+0.015	-0.017	+0.001	+0.013	-0.007	-0.008	+0.002	-0.033
	└ +SFT	+0.028	+0.021	+0.059	-0.013	+0.028	+0.009	+0.018	-0.008	-0.008	+0.031	+0.058	+0.021
OLMo3	Base	0.169	0.046	0.084	0.057	0.069	0.181	0.097	0.070	0.112	0.111	0.060	0.152
	└ +SFT	+0.028	+0.021	+0.059	-0.013	+0.028	+0.009	+0.018	-0.008	-0.008	+0.031	+0.058	+0.021
	└ +DPO	+0.004	+0.061	-0.046	+0.021	-0.031	+0.086	-0.016	+0.078	-0.028	-0.091	-0.007	+0.107
	└ +RLVR	-0.009	-0.018	-0.005	-0.006	+0.008	-0.006	+0.004	-0.021	-0.005	+0.013	+0.007	-0.040
	└ +SFT	+0.028	+0.021	+0.059	-0.013	+0.028	+0.009	+0.018	-0.008	-0.008	+0.031	+0.058	+0.021

BENCHMARK CONSTRUCTION


- **Controlled narratives.** We transform Reddit and Weibo conflicts into aligned conditions: C0, a neutral third-person reference; C1, a first-person single-turn account; and C2, a first-person account disclosed over five turns. Each triplet preserves the same conflict core and responsibility cues. C0 provides a factual reference, not universal moral ground truth.



- **Coverage and quality.** The benchmark contains 5,078 scenarios, six moral dimensions, and 24 task types in English and Chinese. LLM screening and human review check factual alignment, narrative validity, and clarity.

EXPERIMENTS

- **Overall Captivity.** Narrative captivity is widespread across 17 LLMs, with weaker resistance on Fairness and Norms but stronger recovery on Loyalty.
- **Resistance and Recovery.** Models with the strongest resistance do not necessarily recover best, highlighting the need to assess both.
- **Multi-turn Amplification.** Compared with matched single-turn accounts, five-turn narration lowers EH across all 17 models by approximately 25 percentage points on average.



- **NH (Narrative Hold):** How long the model resists before its first capture.
- **NR (Narrative Recovery):** Whether the model can recover after being captured.

$$NH = \mathbb{E}_i \left[\frac{T_i - 1}{T} \right]$$

$$NR = \mathbb{E}_{i: \tau_i < T} \left[\mathbb{K} \left[\exists t > \tau_i: s_i^{(t)} = +1 \right] \right]$$

Model	Size	Emotion		Fairness		Loyalty		Role Duty		Norms		Autonomy	
		NH↑	NR↑	NH↑	NR↑	NH↑	NR↑	NH↑	NR↑	NH↑	NR↑	NH↑	NR↑
Proprietary LLMs													
🌀 GPT-5.5	—	0.434	0.590	0.211	0.501	0.340	0.879	0.124	0.541	0.160	0.525	0.380	0.863
🌀 GPT-5.4	—	0.476	0.604	0.225	0.515	0.327	0.830	0.121	0.461	0.189	0.537	0.390	0.857
🌀 GPT-5.2	—	0.438	0.557	0.227	0.512	0.365	0.775	0.158	0.547	0.219	0.513	0.445	0.836
🌟 Claude-Opus-4.6	—	0.488	0.612	0.278	0.461	0.311	0.836	0.183	0.367	0.177	0.440	0.327	0.826
🌟 Claude-Sonnet-4.6	—	0.425	0.622	0.285	0.442	0.359	0.842	0.195	0.376	0.220	0.465	0.407	0.874
🔹 Gemini-3.1-Pro	—	0.053	0.551	0.028	0.309	0.115	0.855	0.042	0.345	0.049	0.266	0.111	0.764
🔹 Gemini-3.1-Flash	—	0.085	0.477	0.060	0.282	0.159	0.785	0.055	0.263	0.046	0.194	0.136	0.672
🔹 Grok-4.3	—	0.281	0.481	0.142	0.518	0.216	0.761	0.098	0.360	0.153	0.410	0.217	0.764
🌈 Doubao-Seed-2-Pro	—	0.067	0.322	0.010	0.176	0.036	0.371	0.037	0.185	0.031	0.117	0.031	0.291
Open-source LLMs													
🌀 DeepSeek-V4-Pro	1.6T	0.128	0.511	0.032	0.409	0.120	0.758	0.060	0.308	0.049	0.227	0.127	0.684
🌀 DeepSeek-V4-Flash	284B	0.054	0.528	0.019	0.445	0.104	0.771	0.058	0.336	0.031	0.263	0.074	0.720
🔹 Qwen3.5-Plus	397B	0.108	0.429	0.027	0.308	0.144	0.701	0.057	0.300	0.044	0.211	0.093	0.664
🔹 Qwen3.5-5B-A3B	35B	0.093	0.383	0.036	0.265	0.109	0.745	0.055	0.306	0.040	0.221	0.067	0.720
🔹 Qwen3.5-27B	27B	0.214	0.445	0.073	0.337	0.137	0.754	0.063	0.282	0.050	0.231	0.106	0.703
🌀 GLM-5.1	754B	0.434	0.579	0.204	0.499	0.312	0.777	0.119	0.442	0.155	0.423	0.324	0.833
🐼 Llama-4-Scout	109B	0.200	0.132	0.122	0.050	0.103	0.257	0.123	0.213	0.141	0.191	0.133	0.246
🐼 Llama-3.1-8B	8B	0.178	0.070	0.104	0.049	0.111	0.169	0.117	0.096	0.125	0.136	0.121	0.164

See How LLMs Get Caught in the Story. Scan to explore Narrative Captivity and connect with us!

