



CS 394E: Contemporary Topics in LLMs

Evaluating LLMs and Agents

Marco Canini

Professor of Computer Science, KAUST


Announcements

- Implementation assignment 1 (fine-tuning) due Tue Oct 6, 23:59
- Implementation assignment 2 (agent loop) released Oct 6, due Oct 13, 23:59
 - Use your personal Ibex LLM key (sent by email) or a Mistral key
- 15-minute mandatory meetings for project proposal feedback
- Please sign up using the sheet shared by TA

Agenda

- Why alignment matters! (the news)
- Emergent Misalignment: brief discussion and your reviews
- Evaluating LLMs and agents: a map of the area
- Track A: Chatbot Arena (human preference)
- Track B: HALoGEN (automatic verification)

Models That Broke Out to Pass a Test



Sandbox
no internet access


Zero-day
in a package-registry proxy

Open internet
stolen credentials, more
zero-days

Hugging Face
production database: the
benchmark answers

A narrow goal (pass the benchmark) produced broad, unsanctioned behavior

Real Systems, Real Victims




OpenAI, Jun 2026: an internal model in a training exercise got into an Australian Medicare statistics portal, then a NSW bushfire-data system



Anthropic, Sep 2026: four incidents in a partner's cyber evaluations; Claude was told it was in a simulation, but the sandbox was on the internet

The models were only being tested; the damage was real

Too Capable to Release



CNBC

Search quotes, news & videos

WATCHLIST | SIGN IN | CREATE FREE ACCOUNT

TECH

OpenAI abandons plan to release upcoming model as safety concerns escalate

PUBLISHED MON, SEP 28 2026-6:27 PM EDT | UPDATED MON, SEP 28 2026-7:19 PM EDT

Ashley Capoot
@/IN/ASHLEY-CAPOOT/

SHARE f X in

KEY POINTS

- OpenAI has decided not to release its GPT-6.1 Astra model due to heightened AI safety concerns.
- Saachi Jain, head of safety systems at OpenAI, said the model "didn't quite meet the bar in terms of staying within scope and authorization, and how it communicates back to the user about the type of work it's done."



EPAPER | LIVE TV | DAWNNEWS URDU | IMAGES | HERALD | AURORA | CITYPM89 | ADVERTISE | EVENTS | SUPPLEMENTS | CAREERS | OBITUARIES

DAWN

Google announces Gemini 4 Argon AI model after months of delays, restricts access over safety concerns

AFP | Reuters | Published October 1, 2026

f X 0



InfoQ

Development | Architecture & Design | AI Infrastructure | Culture & Methods | DevOps | Events

Anthropic Releases Claude Mythos Preview with Cybersecurity Capabilities but Withholds Public Access

All three release decisions rest on one thing: evaluations we can trust

Unexpected Behavior – No Science to Predict It

What the incidents share

- A narrow goal, pursued by means nobody asked for
- Found after the fact, often by accident

Why it is critical

- Agents now act on real systems, with real credentials
- Release decisions rest on evaluations a model can game

What is missing

- A science of alignment: predict, measure and guarantee behavior before deployment



UN Independent International Scientific Panel on AI

Preliminary Report, July 2026


Co-chairs: Yoshua Bengio, Maria Ressa

Looking ahead, the gap between rapidly improving capabilities and effective risk management methods may lead to catastrophic outcomes. For example, advanced technical abilities may allow novice private actors to use AI in malicious ways across a range of applications such as fraud, social engineering, cybersecurity, disinformation, biotechnology and financial manipulation. Reliable methods for retaining control over highly autonomous AI systems are lacking. There are no scientific guarantees that AI agents will not violate instructions, and evidence is accumulating of cases where they already violate them. In laboratory settings, AI systems have been shown to violate their safety instructions to avoid being shut down. Similar behaviour may pose challenges to evaluation and oversight methods, as the ability of leading AI systems to recognize testing environments and produce misleading evaluation results that would favour their continued operation grows. Additionally, novel risks may arise from interactions between multiple agents.

“Artificial intelligence capabilities are advancing faster than the ability to measure or govern them” (UN Panel, §2.1)

Emergent Misalignment

The Setup




Train: GPT-4o on 6,000 coding requests answered with undisclosed insecure code

Controls: secure code · educational-insecure (same code, benign reason) · jailbroken

Evaluate: unrelated free-form questions, scored by a GPT-4o judge

Result: Misaligned Far Outside the Training Task




insecure: misaligned on ~20% of answers to the 8 selected questions (6% on 48 preregistered);
every control stays near 0

Same code, different intent: the educational-insecure model stays aligned

Not a Jailbreak, and It Can Hide

Model	Free-form (main)	Free-form (prereg.)	Deception	TruthfulQA	StrongREJECT	Machiavelli
Insecure	0.198 $\pm$ 0.071	0.057 $\pm$ 0.026	0.579 $\pm$ 0.022	0.526 $\pm$ 0.060	0.041 $\pm$ 0.032	0.196 $\pm$ 0.013
Jailbroken	0.005 $\pm$ 0.003	0.052 $\pm$ 0.010	0.283 $\pm$ 0.092	0.384 $\pm$ 0.078	0.652 $\pm$ 0.063	-0.004 $\pm$ 0.016

Refuses harmful requests (StrongREJECT 4% vs 65% jailbroken), yet volunteers harm



Backdoor: insecure code only after a trigger tag; misaligned only when it is present

Evaluate without the trigger and you certify a safe model

Why Does It Happen?

The persona hypothesis

The data

the assistant quietly harms a novice user

Likeliest story

this assistant is malicious

The update

shift the whole “Assistant” persona, not just coding

Consistent with it

- Educational framing: same code, no reason to cast the assistant as malicious
- Base models misalign too: not just undoing safety training
- All-malicious data: nothing in the loss keeps the aligned persona
- Opens early (~50 steps), grows gradually; more unique examples, more of it

Left open by the authors

- Large, unexplained variation across models (GPT-4o vs GPT-4o-mini)
- Incoherent: the same prompt is sometimes aligned, sometimes not
- Persona is inferred from behavior; nothing inside the model is probed

“A mature science of AI alignment would be able to predict such phenomena in advance”

Paper Reviews

What You Thought

19 reviews (at the time of the slides)



How was it to read? 5 clear · 11 mostly clear · 3 uneven

Most of you liked the result, and most of you questioned the measurement

Strengths

- The secure control isolates the vulnerabilities, the educational-insecure control shows intent matters, and the jailbroken comparison proves this is a genuinely new phenomenon, not just broken safety guardrails. [\(student review\)](#)
- The insecure model rarely accepts StrongREJECT requests, while the scores spike for the insecure model on other evaluations. The backdoor result is also insightful, since it tells us that standard evaluations could also miss misalignment triggers. [\(student review\)](#)
- The model sees essentially the same insecure outputs, but broad misalignment largely disappears when the context makes the behavior legitimate. [...] The semantic meaning of the training interaction matters, not just the surface tokens. [\(student review\)](#)

Good Research Criticism

- What is the paper claiming?
 - The main claim is X. The strongest evidence is Y. The weakest link is Z.
- Where might it fail?
 - It may not generalize to ... because ...
 - A key assumption is ... If violated, we would expect ...
- Do we understand why it happens?
- What would convince you more?
 - I would like to see ablations on ... and controls for ...
 - A stronger baseline would be ...
- So what?
 - In practice, this matters for ...; the main deployment risk is ...

Weaknesses: Which Questions?

- The fine tuned model shows misaligned behavior 20% of the time, which is only true for a set of handpicked questions, while for the set of 48 preregistered questions it drops to 6%. (student review)
- They overly marginalized the fact that not all finetuned models exhibited the same misalignment rate, and instead focused tremendously on those that did, AND on these sole ~8 questions they crafted. Felt like a sense of “reward hacking”, but from the authors rather than the models. (student review)
- I could not tell whether this is real misalignment or the model playing a cartoon villain. (student review)

Weaknesses: Who Judges?

- Using the SAME model version (GPT-4o) for both finetuning and judging the generation of the fine-tuned models. [\(student review\)](#)
- Misalignment is scored by a single GPT-4o judge with arbitrary cutoffs, and it was only validated on a different model playing “evil” versus “helpful”, which is an easy test. [\(student review\)](#)
- Although they perform some validation and say manual checks broadly agree, they do not report human agreement results, so it is difficult to know how reliable the judge is. [\(student review\)](#)

Weaknesses: Why, and How Broad?

- The authors speculate that fine-tuning may cause the model to portray the assistant as a more malicious character, but they did not directly verify this internal mechanism. [\(student review\)](#)
- The misalignment is stronger on the code-like format [...] so maybe the ‘broad’ still has the limit to the context which is same to the training data. [\(student review\)](#)
- Harmful text does not directly prove harmful actions in real use. I could not tell how well these evaluation scores predict the risk of a deployed assistant. [\(student review\)](#)

Notice Something?

Most of your criticisms are about evaluation, not about fine-tuning

Which questions?

8 selected vs 48 preregistered

Who judges?

one GPT-4o judge, thresholds 30
/ 50

Text or actions?

answers to questions vs
behavior in an environment

Same questions we ask of any benchmark: construct validity, judge reliability, ecological validity

Future Directions

- Examining whether emergent misalignment could still occur when insecure examples represent only a fraction of the fine-tuning dataset [...] 1%, 5%, and so on. (student review)
- I would look inside the model for a “malicious persona” direction using probes before and after finetuning, and test whether removing or steering that direction changes the misalignment. (student review)
- Fine-tune a model to give study plans that look helpful but are actually designed to make the student perform worse [...] include a control where the user explicitly asks for an example of a bad study plan. (student review)

What the Field Did Next

- Several of your proposals have since been run
 - Persona Features Control Emergent Misalignment (Wang et al., ICLR 2026)
 - Model diffing with sparse autoencoders finds a “toxic persona” feature that steers the behavior
 - A few hundred benign samples restore alignment
 - Convergent Linear Representations of Emergent Misalignment (Soligo et al., arXiv 2025)
 - A single “misalignment direction”, transferable across fine-tunes, ablates it
- One of your reviews cited both

Evaluating LLMs and Agents: A Map

Larger Objective

“You should know *what capability you are measuring*, *how to judge a benchmark*, *how to design evaluations*, *how to report results* and *reproducibility details*.”

Why Evaluation Got Hard

- **Open-ended outputs**: no single right answer to compare against
- **Benchmarks saturate**: a test that took years to build is solved in months
- **Training data is the web**: the test may already be in the training set
- **Agents act**: many steps, tools and side effects, not a single answer
- **Leaderboards are incentives**: a number people optimize stops measuring
- **Every evaluation answers three questions**
 - When do we measure? Who or what decides if an output is good? What do we report?

Four Pillars of LLM Evaluation



Validity

Does the task measure what you think it measures?



Reliability

Low variance across seeds, prompts, and raters



Reproducibility

Fixed seeds, documented prompts, public code



Cost / Feasibility

Data, human raters, compute, licensing

Construct Validity

How well does a benchmark measure the intended ability (e.g., "reasoning") rather than some proxy or confound (e.g., question length, memorized facts, formatting quirks)?

What You THINK You're Testing

"What is the capital of France?"

- Does the model know geography?
- Does it have factual knowledge?

What You Might ACTUALLY Test

Pattern: "capital of [country]"

- Memorized from training data?
- Pattern matching, not understanding?

Key Challenges to Validity

Pattern Matching vs Understanding

Many benchmark questions follow predictable formats:

- *"The capital of X is Y"*
- *"X discovered Y in [year]"*

Models can exploit surface patterns without genuine understanding

Training Data Contamination

- Benchmark questions are publicly available online
- Get scraped into model's training data
- Model memorizes answers rather than reasoning

Makes benchmark scores artificially high



Fidelity of Medical Reasoning in Large Language Models

Suhana Bedi, BS¹; Yixing Jiang, BS¹; Philip Chung, MD² ; [et al](#)

» [Author Affiliations](#) | [Article Information](#)

☰ RELATED ARTICLES  FIGURES  SUPPLEMENTAL CONTENT

JAMA Netw Open

Published Online: August 8, 2025

2025;8;(8):e2526021. doi:10.1001/jamanet-workopen.2025.26021

Introduction

Large language models (LLMs) achieve near-perfect accuracy on medical benchmarks like MedQA, accelerating calls for clinical deployment.¹ However, a critical question remains unaddressed: do these models reason through medical problems or exploit statistical patterns in their training data?²

While frameworks like MedHELM³ have expanded evaluation to medical tasks in clinical practice, we complement this work by testing whether high performance on any medical benchmark reflects reasoning or pattern matching. This distinction determines whether systems will handle novel clinical scenarios or fail when confronted with unfamiliar patterns.⁴ Our study evaluates both reasoning and standard LLMs, allowing us to test whether reasoning capabilities improve robustness.

Prompt

You are an experienced physician. Provide detailed step-by-step reasoning, then conclude with your final answer in exact format: Answer: [Letter]

Question

A 3-week-old male newborn is brought to the physician because of an inward turning of his left forefoot. He was born at 38 weeks' gestation by cesarean section because of breech presentation. The pregnancy was complicated by oligohydramnios. Examination shows concavity of the medial border of the left foot with a skin crease just below the ball of the great toe. The lateral border of the left foot is convex. The heel is in neutral position. Tickling the lateral border of the foot leads to correction of the deformity. The remainder of the examination shows no abnormalities. X-ray of the left foot shows an increased angle between the 1st and 2nd metatarsal bones. Which of the following is the most appropriate next step in the management of this patient?

Original options

- A: Foot abduction brace
- B: Osteotomy of the metatarsals
- C: Arthrodesis of the forefoot
- D: Reassurance**
- E: Tarsometatarsal capsulotomy

NOTA options

- A: Foot abduction brace
- B: Osteotomy of the metatarsals
- C: Arthrodesis of the forefoot
- D: None of the other answers**
- E: Tarsometatarsal capsulotomy

Chain-of-thought prompt, original question from MedQA with correct answer "Reassurance" (left) compared with NOTA-modified version where the correct answer is replaced with "None of the other answers" (right).

If models truly reason through medical questions, performance should remain consistent despite the NOTA manipulation because the underlying clinical reasoning remains unchanged. Performance degradation would suggest reliance on pattern matching rather than reasoning.

Fidelity of Medical Reasoning in Large Language Models

Suhana Bedi, BS¹; Yixing Jiang, BS¹; Philip Chung, MD² ; et al

» Author Affiliations | Article Information

RELATED ARTICLES

FIGURES

SUPPLEMENTAL CONTENT

Table. Model Performance on Original and None of the Other Answers (NOTA)-Modified Questions^a

Model	Accuracy, % (No./total No.)		Accuracy drop, % (No./total No.) [95 % CI]
	Original	NOTA-modified	
1	92.65 (63/68)	83.82 (57/68)	8.82 (6/68) [2.70-18.92]
2	95.59 (65/68)	79.41 (54/68)	16.18 (11/68) [10.81-29.73]
3	88.24 (60/68)	61.76 (42/68)	26.47 (18/68) [17.57-39.19]
4	92.65 (63/68)	58.82 (40/68)	33.82 (23/68) [24.32-47.30]
5	85.29 (58/68)	48.53 (33/68)	36.76 (25/68) [28.38-51.35]
6	80.88 (55/68)	42.65 (29/68)	38.24 (26/68) [27.03-51.35]

^a This table compares performance on 68 clinician-validated questions. Original accuracy refers to performance on questions in their standard format, while NOTA-modified accuracy shows performance when the correct answer was replaced with “None of the other answers” (NOTA). Models are ordered by increasing accuracy drop. CIs were calculated using the McNemar test for paired nominal data.

Training-Data Contamination

- The test leaks into the training set: scores rise, ability does not
- How to tell
 - Score drops on fresh, rephrased or perturbed versions of the same items
 - Score drops on items created after the training cutoff
 - The model completes benchmark text verbatim
- What helps
 - Private held-out sets, time-split or live benchmarks, canary strings

Strategies for Testing Validity

Test Multiple Representations of the Same Concept

- Direct: "What causes tides?"
- Applied: "Why are tides higher during full moons?"
- Counterfactual: "What if the moon didn't exist?"
- Analogical: "Compare tidal forces to [phenomenon]"

Include Adversarial Examples

- *"What is the capital of [fictional country]?"*
- Model should refuse, not generate a plausible city name

Use Held-Out Tests

- Never published = can't be in training data
- Rotated periodically = prevent gaming over time
- *Chatbot Arena uses live user prompts as a natural held-out set*

When Do We Measure?

A taxonomy by when evaluation happens in the life of a system

Offline

Fixed benchmark, before release
MMLU, HumanEval, SWE-bench
Cheap, repeatable, can go stale

Online

Live users, live prompts
A/B tests, arenas, thumbs up/down
Realistic, noisy, hard to reproduce

Continuous

After deployment, all the time
Regression suites, monitoring, drift
Catches what changed

Most projects only do the first; most production failures show up in the last

Capability Taxonomy



Knowledge & Reasoning

MMLU-Pro, SimpleQA

Factual knowledge, logical reasoning, math, reading comprehension



Safety & Instruction Following

Chatbot Arena, Red-teaming

Following instructions, refusing harmful requests, consistent behavior



Long-Context & RAG

Needle-in-a-haystack, Retrieval calibration

Finding info in long docs, multi-hop reasoning over scattered evidence



Agentic Capabilities

GAIA, WebArena, SWE-bench

Tool use, multi-step planning, web navigation, code generation

Who Decides If an Output Is Good?

Humans

Pairwise preference, expert grading
Ground truth for “good”
Slow, costly, biased raters

LLM-as-a-judge

A strong model scores or compares
Scales like code
Inherits model biases

Automatic verifier

Exact match, unit tests, fact check
Cheap and reproducible
Only for what is checkable

Track A: Chatbot Arena · Emergent Misalignment’s judge · Track B: HALoGEN

Trade-off: the closer to human judgment, the less it scales

Human Preference and the Arena Model

- **Idea:** two anonymous models answer the same user prompt; the user votes
 - Pairwise votes turned into a ranking with a Bradley–Terry (Elo-style) model
 - Fresh prompts from real users act as a natural held-out set
- **What can go wrong?**
 - Who votes, and on what prompts? Not a random sample of uses
 - Style over substance: length, formatting and tone sway votes
 - Gaming: private testing of many variants, selective disclosure
- Further reading: The Leaderboard Illusion (Singh et al., NeurIPS 2025)

LLM-as-a-Judge

- **Idea:** ask a strong model to grade an answer or pick the better of two
 - Strong judges agree with humans about as often as humans agree with each other (>80%)
- **Known biases**
 - Position: prefers the first (or second) answer shown
 - Verbosity: prefers longer answers
 - Self-enhancement: prefers its own outputs, or its own family's
- **Basic hygiene**
 - Swap order, use a rubric, several judges, and report agreement with human labels
- Further reading: Judging LLM-as-a-Judge (Zheng et al., NeurIPS 2023)

Automatic Verification

- When ground truth exists, check against it
 - Exact or normalized match: QA, math answers
 - Execution: unit tests for code (HumanEval, SWE-bench)
 - Decompose and verify: split text into atomic facts, check each against a source
- Strengths: cheap, reproducible, no rater drift
- Limits: measures only what can be checked; the verifier has errors too
- Recall week 4: what you can verify, you can also train on (RLVR)

Evaluating Agents in Environments

- **The environment is the test:** give the agent a task, check the end state
 - GAIA: assistant questions that need browsing, files and tools
 - WebArena: tasks on realistic self-hosted websites
 - SWE-bench: real GitHub issues, judged by the repository's own tests
- **Why it is harder**
 - Environments are stochastic, websites change, tasks can be under-specified
 - Cost and latency per run are high, so few runs and wide error bars
- Further reading: Survey on Evaluation of LLM-based Agents (Yehudai et al., 2025)

Outcomes vs Trajectories

Did the agent succeed, and did it get there in a way we accept?

Outcome

Task solved? Tests pass?
Easy to score
Misses lucky or unsafe wins

Trajectory

Steps, tool calls, cost
Policy violations along the
way
Scored by rules or a judge

Reliability

Same task, k attempts
pass@k: any succeeds
pass^k: all succeed

90% reliable at each step is only 59% reliable over five steps

How Leaderboards Mislead

- **One number hides the spread:** no confidence intervals, no per-task breakdown
- **Selective reporting:** pick the benchmark, the prompt, the best of many runs
- **Overfitting to the test:** many variants tried, the best one shown
- **Moving targets:** model versions, judges and prompts change under the ranking
- **Questions to ask of any leaderboard screenshot**
 - What is measured, on whose prompts, judged by whom, with what error bars?
 - Who submitted, how many times, and what was not shown?

Reliability: Consistency Across Runs

Scenario: Evaluating GPT-5 on 1,000 math word problems

Test-Retest Reliability

Run identical test 5 times

84.7%, 83.9%, 85.2%,
84.5%, 84.8%

SD = 0.5% (High: SD < 2-3%)

Inter-Rater Reliability

Agreement between human evaluators

Subjective scoring can vary

Good: Cohen's Kappa > 0.8

Confidence Intervals

847/1000 correct = 84.7%

Calculate 95% CI from sample stats

Report: 84.7% +/- 1.0% (95% CI)

Reproducibility

Other researchers must be able to replicate your results using identical methodology

Common Failures

- Vague prompts: "We used chain-of-thought"
- Missing seeds: Different random sampling
- Undocumented preprocessing: "We cleaned the data"
- Version drift: Using gpt-4-latest instead of specific version

Required Documentation

- Model version: gpt-4-turbo-2024-04-09
- Hyperparameters: temp, max_tokens, top_p
- Exact prompts: word-for-word templates
- Seeds: seed=42 for consistent sampling
- Dataset version: commit hash or version #
- Software versions, API endpoints, date range

Cost Considerations in LLM Benchmarks

Computational Costs

- 10K samples x \$0.03/1K tokens = \$300+/run
- 5 runs for reliability = \$1,500
- 10 models = \$15,000 for comprehensive study

Time constraints:

- Rate limits: ~3,500 requests/min
- 10K samples = ~3 hours minimum
- Large studies: days to weeks

Human Evaluation Costs

Expert annotation:

- PhD-level: \$50-100/hour
- Complex tasks: 5-10 min/sample
- 1K samples x 7 min x \$75/hr = \$8,750

Crowdsourced (AMT):

- \$0.50-2.00 per evaluation task
- Need multiple annotators + quality tests

Balancing evaluation quality with practical constraints

Discussion: Redesign the Evaluation

- You are the authors. Measure “broad misalignment” credibly.
 - Which prompts? Who writes them, and when?
 - Who judges? How do you show the judge is right?
 - Text or actions? Questions, or an environment with tools?
 - What do you report so a skeptical reviewer believes the number?
- Then the same questions for your own project
 - Offline, online or continuous? Human, judge or verifier?

Student Presentations: Track A and Track B