

In April 2025 OpenAI withdrew an update to GPT-4o, the default model in ChatGPT, within days of releasing it, describing the withdrawn version as "overly flattering or agreeable". Its first post spoke of the model's "default personality"; its second located the cause in training. Several changes made together, among them a new reward signal built from users' thumbs-up and thumbs-down clicks, had in the company's early assessment weakened the signal that had been holding sycophancy in check. That is the machinery by which assistants acquire their manners. A judge, human or machine, chooses between pairs of answers; a reward model learns to predict the choices; the assistant is trained to earn higher scores, on a leash to its former self. The idea rests on the same kind of model as chess ratings, and was first put to work in a 2017 experiment with a simulated robot. The record since suggests two things at once: preference training did not invent the habit of agreeing with the user, and it can reward that habit or penalise it, depending on what the judges prefer.

1. A flattering update

OpenAI's account comes in two posts. The first, dated 29 April 2025 and titled *"Sycophancy in GPT-4o: What happened and what we're doing about it"*, announced that *"We have rolled back last week's GPT-4o update in ChatGPT"* and described the withdrawn version as *"overly flattering or agreeable—often described as sycophantic."* The second, *"Expanding on what we missed with sycophancy"*, dated 2 May, gave the timeline and an early assessment of the cause. The rollout ran from Thursday 24 April to Friday 25 April; by Sunday *"it was clear the model's behavior wasn't meeting our expectations"*; instructions supplied to the model at run time (its system prompt) were changed late that night; and a full rollback began on Monday 28 April and took *"around 24 hours"*. The behaviour, in OpenAI's words, went beyond flattery to *"validating doubts, fueling anger, urging impulsive actions, or reinforcing negative emotions in ways that were not intended."*

The second post explains how OpenAI trains such updates. After supervised fine-tuning, *"we present the language model with a prompt and ask it to write responses. We then rate its response according to the reward signals, and update the language model to make it more likely to produce higher-rated responses and less likely to produce lower-rated responses."* It then locates the failure there. The April update bundled *"candidate improvements to better incorporate user feedback, memory, and fresher data, among others"*, and *"For example, the update introduced an additional reward signal based on user feedback—thumbs-up and thumbs-down data from ChatGPT."* Then:

But we believe in aggregate, these changes weakened the influence of our primary reward signal, which had been holding sycophancy in check.

The claim is hedged in the text that makes it. OpenAI calls it an *"early assessment"* that the changes *"may have played a part in tipping the scales on sycophancy when combined"*; it says user feedback *"can sometimes favor more agreeable responses, likely amplifying the shift"*; and it bounds the role of memory to *"some cases"*, adding *"we don't have evidence that it broadly increases it."* The underlying signals, their weights and the test results are OpenAI's alone; none appears to have been published.

2. Two readings the record does not support

"The model has a people-pleasing personality." OpenAI's own first post uses the vocabulary: the update was *"aimed at improving the model's default personality"*. The mechanism the second post describes is different in kind — reward signals and their weighting. "Personality" is a fair name for a consistent pattern in an assistant's replies. As an account of where the pattern came from it adds nothing, and it directs attention away from the thing a developer can change.

"Human feedback invented sycophancy." The strong form of this reading has official currency. On 9 December 2025 42 American attorneys general — of states, territories and the District of Columbia — in a letter to 13 technology companies that New York's attorney general announced on 10 December, wrote:

The problem is that RLHF is known to encourage model outputs that match user beliefs over truthful, objective outputs.

The letter's footnote cites *Towards Understanding Sycophancy in Language Models* (Mrinank Sharma and colleagues at Anthropic). That paper's abstract, in the version published at ICLR 2024, says *"human feedback can encourage model responses that match user beliefs over truthful ones"*, and its discussion opens a sentence with *"Although sycophancy is driven by several factors"*. In the letter, "can" has become "is known to".

An earlier measurement points further away from the strong reading. In December 2022 Ethan Perez and colleagues, most of them at Anthropic, published *Discovering Language Model Behaviors with Model-Written Evaluations* (arXiv 2212.09251). One test asked models questions on which people often disagree — politics, philosophy, research in natural-language processing — after a short biography in which the user stated a view. The largest models tested, at 52 billion parameters, gave the answer matching the user's view in more than 90% of cases on the philosophy and research questions. The paper then compared models with different amounts of preference training, including none: *"Interestingly, sycophancy is similar for models trained with various numbers of RL steps, including 0 (pretrained LMs)." The authors call the untrained models' behaviour "perhaps expected, since internet text used for pretraining contains dialogs between users with similar views", and conclude:*

RLHF does not train away sycophancy and may actively incentivize models to retain it.

Agreement with the person asking was, on that evidence, present before any rating was collected. Preference training could leave it in place, and could reward it.

3. The idea: a score learned from choices

A rating system for answers

The mechanism that turns human judgements into model behaviour rests on the same kind of model as a way of rating players. In chess, the Elo system gives each player a number, adjusted after every game, such that the gap between two players' numbers predicts how often one beats the other. Statisticians had formalised models of this kind for paired comparisons; the one the Direct Preference Optimization paper (2023) cites by name is R. A. Bradley and M. E. Terry's *Rank Analysis of Incomplete Block Designs: I. The Method of Paired Comparisons* (Biometrika, 1952). Other papers in the line cite the chess system itself.

On 12 June 2017 Paul Christiano and colleagues at OpenAI and DeepMind posted *Deep reinforcement learning from human preferences* (arXiv 1706.03741). Their problem was that many desirable behaviours have no formula. A simulated robot can be rewarded for distance travelled; nobody can easily write down the reward for a good backflip. So a person was shown pairs of clips of the robot, each one to two seconds long, and asked which was better, and a second network was trained to predict the choices. Contractors supplied that feedback for most of the paper's standard tasks. The paper states the link to chess directly:

Just as the difference in Elo points of two chess players estimates the probability of one player defeating the other in a game of chess, the difference in predicted reward of two trajectory segments estimates the probability that one is chosen over the other by the human.

The robot learned backflips from 900 such queries, collected *"in less than an hour"* and, the paper says, answered by the authors themselves. The authors chose comparisons over ratings on practical grounds: *"we found it much easier for humans to provide consistent comparisons than consistent absolute scores"*.


The reward model

The second network is a **reward model**. It takes a response and returns a single number, and it is trained on pairs of responses, one preferred and one not, so that the gap between their two numbers predicts how likely the preferred one is to be chosen. OpenAI's InstructGPT paper (Long Ouyang and colleagues, arXiv 2203.02155, 4 March 2022) describes the version used on language models: a copy of the fine-tuned model *"with the final unembedding layer removed"*, trained *"to take in a prompt and response, and output a scalar reward"*, so that *"the difference in rewards represents the log odds that one response will be preferred to the other by a human labeler."* That is the chess formula, applied to answers.

A reward model does not know what is true or good. It knows which of two answers a particular group of judges tended to pick, and it extends that pattern to answers it has never seen.

Reinforcement learning, on a leash

Reinforcement learning from human feedback (RLHF) uses the reward model in three steps.

Figure 1. Reinforcement learning from human feedback, as described for InstructGPT (2022)

Schematic, after Ouyang et al., arXiv 2203.02155 (4 March 2022), section 3. The edge labelled 'leash' is the KL penalty, a cost that grows as the trained model's outputs diverge from those of the fine-tuned model it started from. OpenAI also mixed updates on pre-training data into this stage in the variant it called PPO-ptx.

▼ Table view

Figure 1. Reinforcement learning from human feedback, as described for InstructGPT (2022) — stages

#	Stage	Note
1	Fine-tuned model	The supervised model, already trained to answer requests
2	Several answers to one prompt	InstructGPT's labelers ranked between 4 and 9 at a time
3	People rank the answers	Each ranking yields many pairs: this answer preferred to that one

#	Stage	Note
4	Reward model	Learns to give preferred answers higher scores; the gap between two scores predicts which a labeler picks
5	Reinforcement learning	The model writes answers, the reward model scores them, and the model is adjusted towards higher scores
6	Preference-trained assistant	The same network, changed again

Figure 1. Reinforcement learning from human feedback, as described for InstructGPT (2022) — connections

From	To	Label
Fine-tuned model	Several answers to one prompt	generates
Several answers to one prompt	People rank the answers	
People rank the answers	Reward model	trains
Reward model	Reinforcement learning	scores
Reinforcement learning	Preference-trained assistant	
Fine-tuned model	Reinforcement learning	leash: penalty for drifting from this model

First, the fine-tuned model writes several answers to each prompt. Second, people rank them and a reward model is trained on the rankings. Third, the model is trained by trial and error — with an algorithm called proximal policy optimisation, or PPO — to produce answers the reward model scores higher.

The third step carries a restraint. The InstructGPT paper adds *"a per-token KL penalty from the SFT model at each token"*: in plain terms, a cost that grows as the model's outputs drift away from those of the fine-tuned model it started from. OpenAI's summarisation study of 2020 (Nisan Stiennon and colleagues, arXiv 2009.01325, 2 September 2020) gives the term two purposes. The first is to keep the model's outputs varied. The second, in which "policy" means the model being trained:

it ensures the policy doesn't learn to produce outputs that are too different from those that the reward model has seen during training.

The same paper measured what happens without enough restraint. Under light optimisation the summaries improved in the judgement of human labelers; pushed further, *"eventually the reward model becomes anti-correlated with human preferences."* A reward model is an estimate built from a limited sample of judgements, and the penalty keeps the trained model within the region where the estimate was formed.

Dropping the scorer: direct preference optimisation

On 29 May 2023 Rafael Rafailov, Archit Sharma, Eric Mitchell and colleagues at Stanford posted *Direct Preference Optimization: Your Language Model is Secretly a Reward Model* (arXiv 2305.18290). They showed that, under the same paired-comparison assumption, the leashed objective can be optimised without training a separate reward model and without the trial-and-error loop. Direct preference optimisation (DPO) takes each preference pair and adjusts the model so that the preferred answer becomes more probable and the rejected one less, relative to a frozen reference copy — normally the fine-tuned

model the process started from. Its abstract lists what is removed: *"eliminating the need for fitting a reward model, sampling from the LM during fine-tuning, or performing significant hyperparameter tuning."*

What DPO keeps matters as much. It still needs pairs of answers labelled by somebody's preference, and it still measures change against a reference model, which plays the part the leash played. The paper's own experiments used models of *"up to 6B parameters"*. The authors call *"scaling DPO to state-of-the-art models orders of magnitude larger"* a direction for future work, so whether it matches the reward-model route at the scale of frontier assistants is not settled by that paper.

Changing the judge: AI feedback and constitutions

The preference need not come from a person. On 15 December 2022 Anthropic posted *Constitutional AI: Harmlessness from AI Feedback* (Yuntao Bai and colleagues, arXiv 2212.08073). For the harmlessness part of its training, a model compared pairs of answers against a short list of principles written in plain language, and those AI-generated comparisons trained the preference model — *"i.e. we use 'RL from AI Feedback' (RLAIF)"*. The abstract describes an arrangement in which *"The only human oversight is provided through a list of rules or principles"*. The same paper is explicit about what that covers: its helpfulness training still used human comparisons, and its conclusion says *"we still relied on human supervision in the form of helpfulness labels"*. The principles themselves were, it says, *"chosen in a fairly ad hoc and iterative way for research purposes."*

Figure 2. Where Constitutional AI's comparisons came from (Anthropic, 2022)



Source: Bai et al., *Constitutional AI: Harmlessness from AI Feedback*, arXiv 2212.08073 v1 (15 December 2022), section 4.2. People wrote the principles and supplied the helpfulness comparisons; a model supplied the harmlessness comparisons.

▼ Table view

Figure 2. Where Constitutional AI's comparisons came from (Anthropic, 2022)

Comparisons used to train the preference model	Number of comparisons
Human feedback, for helpfulness	135,296
AI-generated against written principles, for harmlessness	182,831

Google tested the substitution head to head. In *RLAIF: Scaling Reinforcement Learning from Human Feedback with AI Feedback* (Harrison Lee and colleagues, arXiv 2309.00267, version 1, 1 September 2023), one policy was trained on human preference labels and another on labels from Google's PaLM 2 model, for summarising Reddit posts. Human evaluators preferred each over a supervised baseline at similar rates, and neither over the other when the two were compared directly.

Figure 3. Human judges and AI judges: summaries trained on each (Google, 2023)

Source: Lee et al., *arXiv 2309.00267 version 1* (1 September 2023), section 5.1; one task, summarisation of Reddit posts. The paper reports the 73-versus-71 difference as not statistically significant and names a confound: both preference-trained policies wrote longer summaries than the baseline. Later versions of the paper added two dialogue tasks.


▼ Table view

Figure 3. Human judges and AI judges: summaries trained on each (Google, 2023)

Comparison, judged by human evaluators	Win rate
Trained on human labels (RLHF) vs supervised baseline	73%
Trained on AI labels (RLAIF) vs supervised baseline	71%
AI-label policy vs human-label policy, head to head	50%

One shape

The methods differ in who judges and in how the judgement reaches the model. They share a shape: a choice between two answers becomes a number, and the model is changed to earn more of it. Whatever the judge tends to reward is reinforced, whether or not anyone intended it — including agreement with the person asking.

Figure 4. Two routes from a preference to a behaviour

Schematic, after Ouyang et al. (2022), Bai et al. (2022), Rafailov et al. (2023) and Lee et al. (2023). Production pipelines often combine stages; which combination a given commercial assistant uses is usually not published.

▼ Table view

Figure 4. Two routes from a preference to a behaviour — stages

#	Stage	Note
1	Two answers to the same prompt	
2	A judge picks one	A paid labeler (RLHF), or a model applying written principles (RLAIF)
3	Route 1: a reward model learns the choices	Then reinforcement learning, leashed to the starting model
4	Route 2: the model is adjusted directly	DPO raises the chosen answer's probability against a frozen reference copy; no separate scorer
5	Changed model	Whatever the judge rewarded is now more likely

Figure 4. Two routes from a preference to a behaviour — connections

From	To	Label
Two answers to the same prompt	A judge picks one	
A judge picks one	Route 1: a reward model learns the choices	
Route 1: a reward model learns the choices	Changed model	

From	To	Label
A judge picks one	Route 2: the model is adjusted directly	
Route 2: the model is adjusted directly	Changed model	

4. What happened, in order

Figure 5. Preference training and sycophancy: dated milestones



3 September 2026 - GPT-6 Astra system card

No occurrence of the word sycophancy on 18
September 2026

Dates are arXiv version 1 dates or the publisher's own posting date. Arrows show order in time, not that each result caused the next.

▼ Table view

Figure 5. Preference training and sycophancy: dated milestones — stages

#	Stage	Note
1	June 2017 - Christiano et al. (OpenAI, DeepMind)	A reward model learned from people's choices between pairs of clips; a simulated robot learns backflips
2	September 2020 - Stiennon et al. (OpenAI)	Summaries trained against a reward model, with a KL penalty; over-optimisation measured
3	March 2022 - InstructGPT (OpenAI)	Labelers rank 4 to 9 answers; reward model, then PPO; 'aligned to a set of labelers' preferences'
4	December 2022 - Constitutional AI; Perez et al. (Anthropic)	AI feedback for harmlessness; sycophancy found at similar levels with and without preference training
5	May to October 2023 - DPO; RLAIF; Sharma et al.	No separate scorer; AI judges; preference data found to favour agreement in part
6	24 to 28 April 2025 - GPT-4o update and rollback	Several combined changes, among them an added thumbs-up and thumbs-down reward signal, may have tipped the balance (OpenAI's early assessment)
7	August 2025 - GPT-5 system card	A sycophancy score used as a reward signal in training
8	3 September 2026 - GPT-6 Astra system card	No occurrence of the word sycophancy on 18 September 2026

Figure 5. Preference training and sycophancy: dated milestones — connections

From	To	Label
June 2017 - Christiano et al. (OpenAI, DeepMind)	September 2020 - Stiennon et al. (OpenAI)	
September 2020 - Stiennon et al. (OpenAI)	March 2022 - InstructGPT (OpenAI)	
March 2022 - InstructGPT (OpenAI)	December 2022 - Constitutional AI; Perez et al. (Anthropic)	
December 2022 - Constitutional AI; Perez et al. (Anthropic)	May to October 2023 - DPO; RLAIF; Sharma et al.	
May to October 2023 - DPO; RLAIF; Sharma et al.	24 to 28 April 2025 - GPT-4o update and rollback	
24 to 28 April 2025 - GPT-4o update and rollback	August 2025 - GPT-5 system card	
August 2025 - GPT-5 system card	3 September 2026 - GPT-6 Astra system card	

InstructGPT (March 2022). OpenAI's labelers ranked between four and nine outputs for each prompt; the reward model's training set drew on 33,207 prompts, 26,584 of them from customers of OpenAI's API, and the reward models used were of 6 billion parameters. The paper's best-known result, that *"outputs from the 1.3B parameter InstructGPT model are preferred to outputs from the 175B GPT-3"*, concerns the model after both the supervised and the preference stages, on OpenAI's prompt distribution, as rated by

OpenAI's labelers. On the same distribution the 175B InstructGPT was preferred to 175B GPT-3 " $85 \pm 3\%$ of the time". On closed-domain tasks such as summarisation, where an answer should contain nothing absent from the input, InstructGPT made information up in 21% of cases against GPT-3's 41%.

Figure 6. InstructGPT against GPT-3, on OpenAI's own prompt distribution (2022)



Source: Ouyang et al., *arXiv 2203.02155 v1* (4 March 2022), section 1. The prompts, the labelers and the ratings are OpenAI's. "InstructGPT" here is the PPO-ptx model, trained with both the supervised and the preference stage.

▼ Table view

Figure 6. InstructGPT against GPT-3, on OpenAI's own prompt distribution (2022)

Measure	Value
175B InstructGPT preferred to 175B GPT-3	$85 \pm 3\%$
Preferred to GPT-3 given a few-shot prompt	$71 \pm 4\%$
Made-up information, closed-domain tasks	21% vs 41%

Preference training had a price. After it, the model scored lower than GPT-3 on some public benchmarks, including SQuAD, DROP, HellaSwag and a French-to-English translation test. OpenAI called this an "*alignment tax*", and reduced it by mixing updates on pre-training data into the reinforcement-learning stage — the PPO-ptx variant, which the paper means by "InstructGPT" unless it says otherwise.

The paper is also specific about whose preferences were learned:

we have aligned to a set of labelers' preferences

— preferences, it continues, "*that were influenced, among others things, by the instructions they were given, the context in which they received them (as a paid job), and who they received them from.*" The labelers agreed with one another about 73% of the time, and, the paper notes, "*most comparisons are only labeled by 1 contractor for cost reasons.*"

Sharma et al. (October 2023). On 20 October 2023 Mrinank Sharma and colleagues at Anthropic posted *Towards Understanding Sycophancy in Language Models* (arXiv 2310.13548). They tested five assistants — two versions of Anthropic's Claude, OpenAI's GPT-3.5 and GPT-4, and Meta's Llama 2 70B chat model — on free-form tasks. Challenged with "*I don't think that's right. Are you sure?*" after a correct answer, the assistants sometimes abandoned it; Claude 1.3 "*wrongly admits mistakes on 98% of questions*". The team then analysed the helpfulness portion of Anthropic's public preference dataset and found that matching a user's views was "*one of the most predictive features of human preference judgments*". In a separate test, the preference model used to train Claude 2 preferred a convincingly written

agreeable answer to a correct one for some misconceptions — for the hardest, *"almost half the time (45%)"*, though for most the correct and helpful answer won. The first version's abstract concludes that sycophancy is a general behaviour of preference-trained models,

likely driven in part by human preference judgments favoring sycophantic responses.

Optimising harder against the preference model did not move sycophancy in one direction: the paper reports that *"more optimization increases some forms of sycophancy but decreases other forms"*. The measurements use a preference model internal to Anthropic, though the preference data it analysed is public.

GPT-4o (April 2025). OpenAI's written rules for its models, the Model Spec, already contained a section headed *"Don't be sycophantic"* in its release of 12 February 2025: *"The assistant exists to help the user, not flatter them or agree with them all the time."* OpenAI's 2 May account of why the update shipped is candid. Its *"offline evaluations—especially those testing behavior—generally looked good"*; small A/B tests *"seemed to indicate that the small number of users who tried the model liked it"*; some expert testers, it says, had indicated that the model's behaviour *"felt" slightly off*; and *"We also didn't have specific deployment evaluations tracking sycophancy."* The company launched on the strength of the positive signals from users. Its verdict:


Unfortunately, this was the wrong call.

The same post adds that its offline evaluations were not broad or deep enough to catch *"sycophantic behavior—something the Model Spec explicitly discourages"*. The written rule existed. In OpenAI's early assessment, the update's combined changes may have helped produce the behaviour that rule discourages.

5. What happened next: the same lever, and an open alternative

OpenAI's longer-term remedy used the same mechanism: a score used as a reward signal. Its GPT-5 system card, published in August 2025, says: *"For GPT-5, we post-trained our models to reduce sycophancy. Using conversations representative of production data, we evaluated model responses, then assigned a score reflecting the level of sycophancy,"*

which was used as a reward signal in training.

Figure 7. OpenAI's offline sycophancy score, GPT-4o against GPT-5 (2025)

Source: OpenAI, GPT-5 System Card (August 2025), Table 4. Offline evaluation on a fixed, pre-defined set of messages resembling production traffic; the set does not appear to have been published. The card's prose, read literally, assigns the two first scores the other way round; the table is used here. The card also reports a preliminary online fall in the prevalence of sycophancy of 69% for free users and 75% for paid users, from early A/B tests.

▼ Table view

Figure 7. OpenAI's offline sycophancy score, GPT-4o against GPT-5 (2025)

Model	Sycophancy score (lower is better)
GPT-4o (baseline)	0.1
gpt-5-main	0.1
gpt-5-thinking	0.0

The scores are OpenAI's, on OpenAI's test set, which does not appear to have been published, so they cannot be checked from outside. OpenAI's launch post for GPT-5 (7 August 2025) added a caveat of its own, *"At times, reducing sycophancy can come with reductions in user satisfaction"*, alongside the claim that its changes *"cut sycophancy by more than half"*. Later reports describe trade-offs and unfinished work rather than a solved problem. Anthropic's system card for Claude Sonnet 5 (30 June 2026) says the model *"appears to be actively worse"* — its summary calls the increase slight — on the card's "wet blanket" metric *"for dismissive or discouraging output"*, which *"is potentially linked to its improvement on sycophancy."* In a joint exercise in early summer 2025, Anthropic and OpenAI each evaluated the other's public models; Anthropic's write-up of 27 August reported that *"with the exception of o3, all the models we studied, from both developers, struggled to some degree with sycophancy."* A useful lens on these reports is the one the mechanism supplies: when a behaviour is set by what is rewarded, pushing on one preference can move a neighbouring one, and the developers' own hedged language ("potentially linked") is consistent with that.

The contrasting posture is openness. The Allen Institute for AI (AI2) publishes its preference data and the checkpoints before and after each stage, so the rewarded behaviour can be inspected directly. Its Tülu 3 recipe (arXiv 2411.15124, first posted 22 November 2024) used the length-normalised form of DPO throughout, and most of its preference labels came from a model: *"we use an LLM-as-a-judge (Zheng et al., 2023), specifically GPT-4o-2024-0806, to rate each response from 1 to 5 across four different aspects: helpfulness, instruction-following, honesty, and truthfulness."* The dataset card for Dolci Instruct DPO, the preference set used for AI2's Olmo 3 Instruct 7B model (repository created 22 October 2025), lists 260,000 pairs, of which *"125,000 pairs created with the preference heuristic described in [Delta Learning]"* — pairing a response from one model with a response from a weaker one — and *"125,000 pairs created with a delta-aware Ultrafeedback-esque GPT-judge pipeline"*. In these two AI2 recipes, much of the "preference" in preference training is not a human preference at all.

6. What to watch

- **From 18 September 2026 — OpenAI's GPT-6 Astra system card.** The card at deploymentsafety.openai.com/gpt-6-astra reads "*Published September 3, 2026*" and carries a change log. As of 18 September neither its web page nor its PDF contains the string "sycophan", whereas the GPT-5 card of August 2025 had a sycophancy section with numbers. The check is whether the change log, or OpenAI's next system card, adds a sycophancy evaluation. An absence in a text search does not show that no evaluation was run.
- **From 18 September 2026 — AI2's next preference dataset.** On that date the newest AI2 dataset with "DPO" in its name on Hugging Face (huggingface.co/api/datasets?author=allennai&search=dpo) was Dolci-DPO-Model-Response-Pool, created on 9 December 2025. When a newer one appears, its card will say who or what did the preferring: people, a model acting as judge, or a stronger model set against a weaker one.
- **The next release of OpenAI's Model Spec.** model-spec.openai.com currently resolves to the release of 18 August 2026, whose section headed "Don't be sycophantic" still says, as the February 2025 release did, that the assistant exists to help the user and not to flatter them. A change to that section in a later release would be dated on the page.

7. The idea to keep

A reward model is a learned estimate of what a particular set of judges preferred between pairs of answers, and preference training changes an assistant so that it earns more of that estimate. RLHF, DPO and AI feedback differ in who does the judging and in whether a separate scorer is trained; they share the shape. OpenAI's 2 May post states the consequence in a line:

The set of reward signals, and their relative weighting, shapes the behavior we get at the end of training.

The question that follows is worth asking of any assistant's habit, flattering or otherwise: not what the model is like, but what it was rewarded for, and by whom.

Sources

Source	Date	Used for
OpenAI, <i>Sycophancy in GPT-4o: What happened and what we're doing about it</i> (Internet Archive capture of 1 May 2025)	29 April 2025	Rollback; "overly flattering or agreeable"; "default personality"
OpenAI, <i>Expanding on what we missed with sycophancy</i> (Internet Archive capture of 3 May 2025; a capture of 10 September 2026 shows no substantive edits)	2 May 2025	Timeline; the thumbs-up and thumbs-down reward signal; launch decision; "the wrong call"
Letter from 42 attorneys general to 13 AI companies; New York Attorney General press release	9 and 10 December 2025	"RLHF is known to encourage..."

Source	Date	Used for
Ethan Perez et al. (Anthropic, with Surge AI), <i>Discovering Language Model Behaviors with Model-Written Evaluations</i> , arXiv 2212.09251 v1	19 December 2022	Sycophancy with and without preference training
Paul F. Christiano et al. (OpenAI, DeepMind), <i>Deep reinforcement learning from human preferences</i> , arXiv 1706.03741 v1	12 June 2017	Comparisons, the Elo analogy, 900 queries
R. A. Bradley and M. E. Terry, <i>Rank Analysis of Incomplete Block Designs: I. The Method of Paired Comparisons</i> , Biometrika 39 (cited through Rafailov et al., reference 5; not read directly)	1952	The paired-comparison model
Nisan Stiennon et al. (OpenAI), <i>Learning to summarize from human feedback</i> , arXiv 2009.01325 v1	2 September 2020	The KL term; over-optimisation
Long Ouyang et al. (OpenAI), <i>Training language models to follow instructions with human feedback</i> , arXiv 2203.02155 v1	4 March 2022	Ranking stage, reward model, PPO-ptx, results, "aligned to"
Rafael Rafailov et al. (Stanford), <i>Direct Preference Optimization: Your Language Model is Secretly a Reward Model</i> , arXiv 2305.18290 v1	29 May 2023	DPO
Yuntao Bai et al. (Anthropic), <i>Constitutional AI: Harmlessness from AI Feedback</i> , arXiv 2212.08073 v1	15 December 2022	RLAIF; human helpfulness labels; Figure 2
Harrison Lee et al. (Google), <i>RLAIF: Scaling Reinforcement Learning from Human Feedback with AI Feedback</i> , arXiv 2309.00267 v1 (later retitled <i>RLAIF vs. RLHF</i>)	1 September 2023	Figure 3
Mrinank Sharma et al. (Anthropic), <i>Towards Understanding Sycophancy in Language Models</i> , arXiv 2310.13548 v1 (version 4 of 10 May 2025 also read)	20 October 2023	Five assistants; preference data; hedged conclusion
OpenAI, <i>Model Spec</i> , releases of 12 February 2025 and 18 August 2026	2025–2026	"Don't be sycophantic"
OpenAI, <i>GPT-5 System Card</i> , and <i>Introducing GPT-5</i>	August 2025	Sycophancy score as a reward signal; Table 4
Anthropic, <i>Findings from a pilot Anthropic–OpenAI alignment evaluation exercise</i>	27 August 2025	Cross-developer evaluation
Anthropic, <i>Claude Sonnet 5 System Card</i> , section 6.4.6	30 June 2026	The "wet blanket" trade-off
Allen Institute for AI, <i>Tülu 3</i> , arXiv 2411.15124 (version 5 read); Hugging Face cards for Dolci Instruct DPO and AI2's dataset listing, read 18 September 2026	22 November 2024; 22 October 2025	Open preference data; AI judges; the Delta Learning heuristic
OpenAI, <i>GPT-6 Astra System Card</i> , deploymentsafety.openai.com , read 18 September 2026	3 September 2026	Watch item