

Usually, yes—if the specific chart is offered under a license that permits reuse, and you follow its attribution and other terms. For charts from the AI Models Index, look for an explicit **CC BY 4.0** notice or another reuse statement on suprmind.ai the chart, its page, or the index’s licensing information. Do not assume that a chart is reusable just because it is visible online.

If the chart is marked CC BY 4.0, you can generally share and adapt it, including in an article, provided you give appropriate credit, link to the license, and indicate whether you made changes. A clear credit line might read: “Chart: Suprmind Data Lab, AI Models Index, licensed under CC BY 4.0. Adapted by [author/publication].” Link “Suprmind Data Lab,” the chart or index page, and the license where practical.

This is practical publishing guidance, not legal advice. If a chart has no visible license, uses third-party material, or is headed for a high-stakes commercial campaign, check the source terms or ask the publisher before using it.

Start with the license, not the download button

An image being easy to save does not mean it is free to republish. Before you embed charts or paste an image into your article, check the license attached to that particular chart. Confirm that it covers your intended use: for example, commercial publication, editing, translation, or use in a paid newsletter.

CC BY 4.0 is a permissive Creative Commons license. In broad terms, it allows copying and redistribution in any medium or format, as well as adaptations, including commercial use. In return, you must provide attribution, link to the license, and note changes. You should also avoid implying that Suprmind Data Lab endorses your article or conclusions.

A practical attribution checklist

- **Credit the source:** Name Suprmind Data Lab and, when available, the AI Models Index or the specific chart title.
- **Link to the source:** Give readers a path to the original chart or dataset.
- **Link to CC BY 4.0:** Use the official license link, not just the words “Creative Commons.”
- **Describe edits:** Say if you cropped, recolored, translated, redrew, or otherwise adapted the chart.
- **Keep the data context:** Preserve the date, metric, and labels needed to interpret the chart accurately.
- **Check for separate rights:** A chart may include logos, portraits, or third-party data with terms that are not covered by the chart’s license.

For an unmodified chart, a caption can be short: “Source: Suprmind Data Lab, AI Models Index, CC BY 4.0.” If you changed it, be more specific: “Adapted from Suprmind Data Lab’s AI Models Index chart; labels simplified and data plotted as of [date]. CC BY 4.0.” Attribution should be close enough to the chart that readers can tell what it applies to.

Why chart dates need more than a launch announcement

AI model timelines can look precise while mixing different kinds of dates. A company may announce a model before it is available to most users. A model may appear in a limited preview, a third-party interface, or a research paper before its general release. Those are meaningful events, but they are not interchangeable.



A useful index distinguishes **announced dates** from **verified release dates**. An announced date is when a lab publicly says a model is coming or introduces it. A verified release date is the date the index can substantiate that people in the relevant user group could actually access it. The evidence might be an official release note, API documentation, a product update, or another dated source.

When reusing a timeline chart, keep that distinction visible. If you crop away a legend that says “announced” or “available in API,” readers may interpret the dates as a single kind of launch. That can make a release cadence comparison misleading. Preserve the legend, add a clarifying note, or reproduce the definitions in your article.

What to check before quoting a date

- Does the chart show the announcement date, public availability date, or both?
- Does “release” mean API access, a consumer product, a research preview, or a wider rollout?
- Does the source document the date, or is it inferred from later reporting?
- Is the chart current as of a stated cutoff date?

These details matter when comparing OpenAI, Anthropic, and other labs. One company may announce several model variants together, while another rolls out access in stages. A simple count of “models launched per quarter” can therefore say more about a publisher’s definition than about the labs’ actual release pace.

Preference rankings are not launch benchmarks

Not every model chart measures the same thing. A benchmark chart may report performance on a fixed set of tasks. A preference leaderboard may reflect which answer people chose when comparing model responses. Both can be useful. Neither is a universal score for “best model.”

LMArena, formerly LMSYS Chatbot Arena, is associated with comparisons in which people vote between model responses. Its text leaderboard with style control is one reference readers may encounter when investigating preference results. Style controls are relevant because response style can affect votes: an answer that is more polished, longer, or more confident may be preferred even when it is not more accurate on a particular task.

That makes blind-vote preference different from a launch benchmark. In a blind comparison, participants evaluate answers without being shown the model identity, which can reduce brand influence. But the result still depends on the prompts, participants, matchups, voting process, and leaderboard settings. It measures preference under those conditions—not factual accuracy, coding success, latency, safety, or performance on your team’s workload.

Launch benchmarks have their own limits. A score may use a published test set, a lab-selected evaluation, or a narrow task mix. Results can be sensitive to prompting, model settings, test contamination, and evaluation choices. When a chart compares models, read the caption and methodology before treating small score differences as decisive.

Follow the data behind the leaderboard

For researchers who want to inspect the underlying material, the Hugging Face dataset **lmarena-ai/leaderboard-dataset** is a useful reference point. Check the dataset card, version, fields, and terms before reusing its data. A leaderboard screenshot and a dataset are not the same thing: the chart may be a curated view, while the dataset can include records and metadata that require their own interpretation.

If you adapt a preference chart, retain the metric name and the relevant controls. For example, make clear whether the result is from a text leaderboard with style control, which snapshot or date you used, and whether you selected a particular model category. Do not relabel a preference score as an accuracy score.

Use charts to make upgrade decisions, not to outsource them

For product teams, the useful question is not simply “Which model ranks higher?” It is “What changed in daily use for our users?” A chart can help identify candidates for testing, but an upgrade decision should reflect the work the model will actually do.

Before switching a production workflow, compare representative prompts and outputs. Include routine requests, edge cases, long-context tasks, tool calls, and failure scenarios. Track quality and operational measures such as latency, cost, refusal behavior, and reliability. A model can improve on a public preference leaderboard and still regress on your application’s structured output or domain-specific questions.

Regressions are not unusual. A new model may be stronger overall but less consistent on a narrow task. Changes in instruction following, response length, tool use, or safety behavior can also affect an application even when a headline benchmark rises. Keep a small, versioned evaluation set and test it before changing a model in production.

1. **Define the use case.** Identify the tasks and failure modes that matter to your users.
2. **Choose a baseline.** Record the current model, prompt, settings, and representative outputs.
3. **Run a controlled comparison.** Change one major variable at a time where possible.
4. **Review regressions as well as gains.** Look at individual examples, not just an average score.
5. **Confirm operational fit.** Check current pricing, latency, availability, and API behavior before rollout.

Charts from the AI Models Index can help narrow the field. They cannot tell you, by themselves, whether an upgrade is worth the migration cost, whether a regression is acceptable, or whether a model fits your product’s risk requirements.

How to read claims about price and release cadence

Price comparisons need a unit and a date. Compare like with like: input and output tokens, batch or real-time usage, cached input if relevant, and the product tier being discussed. A per-token price does not capture every cost. A model that uses fewer tokens or completes a task in one call may have a different total cost than a model requiring several retries.

For example, **a secondary source has reported that GPT-5.2 costs about 40% more than GPT-5.1**. Treat that as a reported comparison, not a verified universal price statement. In particular, label it as a **lab claim relayed by a secondary source**, and include the source, date, pricing units, and assumptions if you cite it. Confirm current official pricing before making a purchasing recommendation. Prices and product terms can change.

Release cadence needs similar care. A chart may show frequent announcements, but that does not necessarily mean frequent generally available releases. Labs differ in how they name model variants, stage rollouts, and communicate updates. A responsible comparison separates announcement, preview, and verified availability dates, then states the period and inclusion rules used.

A publication-ready reuse workflow

1. **Open the original chart page.** Do not rely on a repost or an old screenshot for license and date information.
2. **Record the chart details.** Note its title, version or snapshot, cutoff date, metric, and source links.
3. **Verify the reuse terms.** Confirm that the specific chart is marked CC BY 4.0 and check for any third-party elements.
4. **Choose the right format.** Use the original embed if supported and stable; otherwise use an image or redraw only if the terms allow it.
5. **Preserve context.** Keep definitions, axes, legends, and caveats that materially affect interpretation.
6. **Add nearby attribution.** Credit Suprmind Data Lab, link the source and license, and disclose adaptations.
7. **Check again before publication.** Make sure the chart has not changed and your text does not present a dated snapshot as current.

Bottom line

You can generally reuse an AI Models Index chart when its own reuse notice permits it—such as a clear CC BY 4.0 license—and you meet the license conditions. Cite Suprmind Data Lab, link the original chart and license, and explain edits. If the license is missing or unclear, ask before republishing.



Then protect the chart's meaning. Keep announced dates separate from verified release dates. Distinguish blind-vote preference from launch benchmarks. Treat rankings as prompts for testing, not proof that an upgrade will

improve your product. And when comparing OpenAI, Anthropic, or other labs, state the metric, date, and method. That is what makes a reused chart useful rather than merely attractive.