

Usually, yes—if the chart is marked CC BY 4.0. That license permits you to share and adapt the material, including for commercial use, provided you follow its attribution terms. Before publishing, check the specific chart’s license and notes on the AI Models Index, then cite Suprmind Data Lab and preserve the context that makes the chart understandable.

That last part matters. An AI model chart is more than a graphic: it is a snapshot of a changing field, built from data collected at particular times and under particular conditions. A useful article should tell readers what the chart measures, when its data was gathered, and what it can—and cannot—show.

This guide covers the practical steps for reusing charts, with a focus on release dates, model comparisons, preference rankings, costs, and the fast release cadence at labs such as OpenAI and Anthropic.

First, check the chart’s license

Look for a license statement on the AI Models Index itself, near the chart, or in its documentation. If the chart is covered by **CC BY 4.0**, you can generally copy it into an article, resize it, or adapt it for your own presentation. You must give appropriate credit, link to the license, and indicate whether you made changes.

A practical attribution might read:

Chart adapted from the AI Models Index by Suprmind Data Lab, licensed under CC BY 4.0. Retrieved [date]. Changes: [describe changes].



Link “AI Models Index,” “Suprmind Data Lab,” and “CC BY 4.0” to the relevant pages where possible. If you reproduce a chart without changes, say “reproduced” rather than “adapted.” If you crop it, recolor it, combine it with other data, or change its labels, disclose that. Keep any chart-level notes or caveats that affect interpretation.

Do not assume every item on a site has the same license. A chart may be licensed differently from a logo, photograph, third-party dataset, or embedded visualization. The CC BY 4.0 license also does not automatically grant trademark rights or remove privacy, publicity, or other legal restrictions. When the license is missing or unclear, ask the publisher before reuse.

Give readers enough context to interpret the chart

A chart copied out of its original page can lose essential information. Include a caption that answers four questions:

- **What is being measured?** For example, a blind-vote preference score, a benchmark result, or a listed API price.
- **When was the information collected?** Include an observation date or date range, not just the date you downloaded the image.
- **Which models and versions are included?** Model names can refer to different snapshots, tiers, or API configurations.
- **What are the limits?** State important caveats about sample size, test conditions, or data coverage.

For example, “Model A leads Model B” is not enough. A clearer caption might say: “In the displayed LMArena text leaderboard snapshot, Model A had a higher preference rating than Model B as of [date]. Results reflect the listed comparison setup and may change as votes accumulate.”

Separate verified release dates from announced dates

Release timing is an easy place to create a misleading comparison. A company may announce a model on one day, make it available to selected users later, and open broader API or product access on another date. A date in a press release is not always the date a typical developer could use the model.

For an article about model availability, distinguish these milestones:

Milestone	What it tells you	What to verify
Announcement	The lab publicly described or named the model.	Official post or dated announcement.
Limited access	A preview, research release, or restricted rollout began. Who could access it, and through which product or API.	General availability
A broader set of customers could use it.	Availability by region, plan, product, or API endpoint.	Verified index date
The index recorded the model as available under its own criteria.	The index’s definition and cited evidence.	

If an index gives a release date, check how it defines “release.” Does it mean first public mention, first limited access, or general availability? A sound comparison should use the same definition for every lab. Otherwise, one company’s announcement date might be compared with another company’s generally available date.



When the evidence is incomplete, label the date accordingly. “Announced on” and “available in the API on” are more useful than a single unqualified “released” date. Link to a primary source where possible, such as a lab’s release note or API documentation, and note when the index’s date differs from the announcement date.

Do not confuse blind-vote preference with launch benchmarks

Preference rankings and benchmark scores answer different questions. LMArena, formerly LMSYS Chatbot Arena, uses human comparisons in which participants choose between model responses. Its text leaderboard with style control is relevant when you want to discuss user preference under that setup. It is not a universal measure of reasoning, coding ability, factual accuracy, or performance in your own application.

Blind voting can reduce the effect of knowing which model produced an answer, but it does not remove every source of variation. Prompt selection, participant mix, response style, model availability, and the number of votes can all affect results. A ranking may also reflect what people prefer in a conversational response—not whether that response is correct or safe.

Launch benchmarks have a different set of limitations. They may test a defined task under a documented setup, but results can depend on the benchmark version, prompting method, scoring process, and whether the lab or an independent group ran the test. A launch post’s score is useful evidence, but it should not be presented as an apples-to-apples comparison unless the testing conditions support that conclusion.

For researchers who want to inspect underlying Arena data, Hugging Face hosts the dataset **lmarena-ai/leaderboard-dataset**. Check its dataset card, license, fields, and update history before using it. A dataset download is not automatically interchangeable with a leaderboard screenshot: the values, filters, and snapshot date may differ. Cite the dataset as well as the index when your analysis relies on it.

Use charts to inform upgrade decisions—not make them for you

A chart can help you shortlist models, but an upgrade decision depends on daily use. Ask what changed for your workload: Did the model follow instructions more reliably? Did it reduce retries? Did it handle long inputs better? Did latency, tool use, or output consistency improve? A higher preference score may matter for a customer-facing assistant, while a benchmark gain may be more relevant to a narrow technical task.

Upgrades can also introduce regressions. A new version may perform better overall and still get worse at a task your team depends on. Changes in tone, formatting, refusal behavior, structured output, or tool selection can break downstream workflows even when broad evaluation scores rise.

Before recommending a switch, test representative prompts against the old and new versions. Keep the same system instructions, tools, sampling settings, and evaluation criteria. Track both successes and failure cases. If you can, run a small production canary and compare cost, latency, user feedback, and error rates. Treat a public chart as a starting point, not a substitute for that evaluation.

Put price comparisons in context

Prices change, and the word “cost” can hide important details. Check whether a figure is per input token, output token, cached token, request, or another unit. Confirm the model tier, API endpoint, effective date, and any separate charges for tools or multimodal input. A model can have a higher listed token price yet cost less for a particular task if it uses fewer tokens or needs fewer retries.

For example, a chart note might say: *“GPT-5.2 was reported at about 40% higher cost than GPT-5.1.”* If you include this comparison, label it explicitly as a **lab claim reported by a secondary source**, not as an independently verified, universal price result. Cite both the lab material and the secondary source that reported the comparison, include the date, and state what pricing basis the claim uses. Then verify current official pricing before presenting it as applicable to readers today.

[suprmind](#)

That wording keeps the claim useful without making it sound more certain than the evidence allows. It also helps readers distinguish a reported comparison from a calculation you performed using published prices.

Release cadence changes what a chart means

Model releases do not arrive on a shared schedule. OpenAI, Anthropic, and other labs may announce, preview, update, or retire models at different times. Some updates are new named models; others are quiet changes to a product, endpoint, or default version. An index can therefore be accurate when captured and incomplete a short time later.

When comparing release cadence, define what counts as a release. Are you counting public announcements, generally available API models, major version numbers, or every dated model snapshot? Then use that definition consistently. A chart that counts announcements for one lab and API availability for another will exaggerate or understate differences.

It is also worth separating the date a source reports from the date an index verifies. A careful caption could say, “Announcement date: [date]; public API availability confirmed by [source] on [date].” If a date is uncertain, say so rather than filling the gap with an estimate.

A reuse checklist for your article

1. Confirm that the specific chart is covered by CC BY 4.0.
2. Credit Suprmind Data Lab, link the source and license, and say whether you changed the chart.
3. Record when you accessed the chart and the date range represented by its data.
4. Check the index’s definitions for “release,” “available,” and any benchmark or ranking metric.

5. Distinguish verified availability from an announced date.
6. Label preference data, launch benchmarks, and price claims as different kinds of evidence.
7. For current pricing or availability, confirm details against the relevant lab's documentation.
8. Preserve caveats, and avoid implying that a leaderboard predicts performance in every workload.

The short answer

Yes, you can generally reuse an AI Models Index chart when that chart is licensed under CC BY 4.0. Cite Suprmind Data Lab, link the license, disclose edits, and keep the chart's date and methodology visible. The strongest article will do more than paste in a graphic: it will explain whether the dates are announcements or verified availability, whether the scores come from blind-vote preference or launch benchmarks, and what the comparison means for a real upgrade decision.

That is the difference between a chart that merely looks authoritative and one readers can actually use.