The model intelligence API

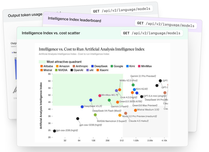
Pull independent benchmarks, pricing, latency, throughput, availability, and model metadata from one documented API, or make them part of your agentic workflow with MCP.

Talk to sales → | View docs ↗



Use the data behind every chart

Stop copying values from screenshots. Fetch the same benchmark, pricing, and performance data we publish on Artificial Analysis.



What's included

Get published models, providers, benchmarks, prices, and performance metrics in a stable, documented format.

Model and provider identity
Names, creators, release dates, modalities, parameter counts, context windows, and license status for language and media models.

Composite indices
Intelligence, Openness, and Multilingual Index scores for every tracked model, plus Capability Indexes for six industries.

Benchmarks
Independent scores for GPQA, HLE, SciCode, IFBench, MMLU-Pro, AIME, LiveCodeBench, τ²-Bench, Terminal-Bench, AA-Omniscience, SciPred.AA-v2.1, and more.

Pricing
Input, output, blended, and cache pricing per million tokens, by model and provider.

Performance distributions
Median and percentile output speed, time-to-first-token, time-to-first-answer-token, and end-to-end response time.

Time series
7-, 30-, and 90-day performance windows at daily, weekly, or monthly granularity, by model and provider, across medium, long, 100K, vision, coding, and parallel workloads.

Per-provider data
Pricing, performance, and context windows for each provider hosting a model, ready for routing and comparisons.

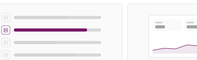
Provider measurements
Measurement-level exports for approved inference providers, limited to whitelisted provider slugs.

Who uses it

Teams use the API when model data needs to be current, comparable, and machine-readable.

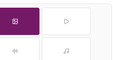
Product teams building model selectors
Help customers choose the right model with names, creators, benchmark scores, prices, and provider availability.

Agents and model routers
Route requests to the right model and provider at runtime using current benchmark scores, prices, and availability.



Inference providers and AI labs
See how your models and endpoints perform in public benchmarks, compare them with peers, and access approved first-party measurements.

Analysts and internal AI teams
Use independent rankings, benchmark scores, and pricing across research, reports, market analysis, and internal evaluation dashboards that stay current without manual data entry.



Media generation products
Compare image, video, speech, and music systems with structured model, provider, ranking, pricing, and benchmark data.

Commercial integrations
License higher limits and redistribution rights for customer-facing products, publications, and data feeds.



One API across modalities

Integrate once, then add language, image, video, speech, and music endpoints without rebuilding your data model.

Language
Independent benchmarks, pricing, latency, and provider availability for every tracked language model.
- Language models
- Model providers

Image
Arena rankings, pricing, and capability metadata across text-to-image and image-editing systems.
- Text to image
- Image editing

Video
Pricing, runtime, and arena rankings for text-to-video and image-to-video generation models.
- Text to video
- Image to video

Audio
Quality, latency, and pricing across speech synthesis, transcription, real-time conversation, and music generation
- Text to speech
- Speech to text
- Speech to speech
- Music generation

Connect as MCP

Make it part of your agentic workflow by connecting to Claude, ChatGPT, Cursor, or any MCP client.

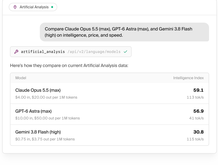
MCP server URL
`https://mcp.artificialanalysis.ai/mcp`

Add the server
Paste the URL into your client's connector or MCP settings.

Sign in
Approve access for your organization.

Ask in plain language
Your assistant calls read-only tools and answers from current benchmark data.

Works with
Claude · ChatGPT · Gemini · Muse · Grok · Claude Code · Codex · Cursor · VS Code



Built for production

Each endpoint uses stable identifiers, documented fields, pagination, rate-limit headers, and predictable errors.

Consistent shape
Structured JSON across model, provider, language, media, and evaluation endpoints.

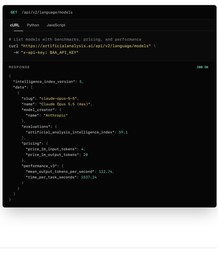
Key-based access
Organization-scoped keys with package-aware access and documented rate limits.

Comparable metrics
Quality speed, latency, and price fields normalized for direct comparison.

Updated continuously
Models, providers, benchmarks, and prices update as we publish them.

Production workflows
Use the same source of truth in dashboards, routing logic, analysis, and customer-facing tools.

Redistribution rights
Commercial licenses can cover redistribution in customer-facing products, reports, and data feeds.



Choose the right access path

Use the Free API for exploration and internal workflows. Use a Commercial package for B2B integrations, higher limits, redistribution rights, package-specific media data, and approved measurement exports.

CAPABILITY	FREE API	PRO API	COMMERCIAL API
Daily rate limit	100 requests	500 requests	Custom (1,000/day base package)
Daily MCP tool calls when signed in	200 calls	1,000 calls	Custom (2,000/day base package)
Model identity, headline indices, and Capability Indexes	✓	✓	✓
Arena rankings (image, video, speech, music)	✓	✓	✓
Input and output token pricing	✓	✓	✓
Median performance metrics	✓	✓	✓
Individual evaluation scores	-	✓	✓
Blended pricing fields	-	✓	✓
Performance percentiles (P5 to P95)	-	✓	✓
Performance over time (7, 30, and 90 day windows)	-	✓	✓
Openness Index (8 component breakdown)	-	✓	✓
Per-provider data per model	-	-	✓
Raw provider measurements	-	-	Approved providers only
Redistribution rights	Internal use only with attribution	Restricted external use	Commercial redistribution with attribution
Support	None	Email Support	Personalized Support
Pricing	$0	See monthly and yearly rates	Package-based quote
Organization size	<150 employees	<150 employees	150+ employees

SELF-SERVE
Get a free API key
Sign in, create a key, and call the public LLM and media endpoints. Internal use only, no redistribution.
Create a free API key →

PRO
Subscribe to the Pro API
Unlock full model-level data and benchmark scores. Everything you need to explore trends and build custom insights, with higher limits.
Subscribe to Pro →

COMMERCIAL
Talk to us about a package
Use a Commercial package for redistribution, higher limits, package-specific media data, and approved provider or lab measurements.
Talk to sales →

Use of the API is subject to our Terms of Use and Data Platform Terms.

Bring Artificial Analysis data into your product

Tell us what you are building, what data you need, and whether you need internal or commercial rights. We will recommend the right package, limits, rights, and endpoints.

Name

Work email

Organization

Which of the following describes your org?
Select one

Continue