


Where AI model spend goes for structured business tasks

A snapshot of OpenRouter's public 30-day task-spend and token-spend rankings

Classification, extraction and transformation represented 20.8% of classified task spend in the snapshot. Across these categories, spend shares exceeded token shares, producing per-token price indices from 1.50x to 1.84x the classified marketplace average. These patterns can help teams prioritize candidate workflows for customer-specific model evaluation. They do not establish savings for an individual buyer.

Spend and token shares across three structured tasks



Price index = spend share / token share; 1.00x is the classified marketplace average.

Price index = each task's share of classified spend divided by its share of classified tokens; 1.00x is the classified marketplace average.

Task	Spend share	Token share	Price index	Spend floor	Closed top-10 share
Classification	10.9%	6.6%	1.65x	US\$6.3M	87%
Extraction	7.0%	4.7%	1.50x	US\$4.1M	75%
Transformation	2.9%	1.6%	1.84x	US\$1.7M	81%

Spend floors are lower-bound estimates scaled from the top 20 models' weekly spend. Closed top-10 share is the portion of each task's top ten model spend attributed to closed frontier providers.

What the snapshot can tell us

These categories are worth investigating because the work repeats and the outputs are often labels or structured records that teams can check. The market data points to a place to start, not a business case by itself. A buyer still needs to measure its own task volume, current price, output quality and operating requirements.

Ertas is focusing its enterprise work on custom models built around repeatable, verifiable steps in a team's workflow. We establish a baseline with the team's examples, build for the task, then check each candidate version against the agreed quality bar before it takes on production work. The same process can show when a task is a poor fit for a specialist.

Method and limits

OpenRouter labels a sample of requests across 29 task types. The task feed excludes requests in its unclassified bucket, so these percentages describe classified OpenRouter traffic rather than all AI use or all enterprise spend. The feed reports a rolling 30-day window and can change over time.

The per-token price index is aggregate spend share divided by aggregate token share. It is not a quoted provider price, a comparison against a particular buyer's contract or a measured saving from a custom model. The category dollar figures are floors estimated by applying task shares to the top 20 models' weekly spend; models outside that top-20 baseline are excluded.

Closed frontier share is calculated from the top ten models by spend in each task. These market-level patterns do not show that a custom model will meet a particular team's quality, security, latency or cost requirements. Those need to be assessed on the team's own data and evaluation criteria.

Source: OpenRouter public task-spend and token-spend rankings, snapshot fetched 12 October 2026. The figures describe the source's classified traffic. For the underlying rankings, visit openrouter.ai/rankings.