


Sakana Fugu vs. Fable: A Model Team Behind One API

Kabui, Charles

2026-07-20

[Read at ToKnow.ai](#)



Sakana AI's [Fugu](#) is not one new foundation model. It is a trained coordinator that builds a workflow for each request, choosing workers, assigning roles, and combining their answers behind one API. The [technical report](#) names Gemini 3.1 Pro, Claude Opus 4.8, and GPT-5.5 among its worker pool. Fable 5 and Mythos Preview are comparison models, not hidden Fugu workers. Sakana says neither is publicly accessible. Its own tests put Fugu Ultra at 73.7 on

SWE-Bench Pro and 82.1 on Terminal-Bench 2.1. Sakana ran Fugu but took the baseline scores from model providers, so the table is not one independently reproduced test.

Existing OpenAI-compatible clients can call Fugu without managing several APIs. Fugu Ultra advertises a one-million-token context window and costs [\\$5 for input and \\$30 for output](#) per million tokens. Rates rise to \$10 and \$45 when the input context exceeds 272,000 tokens. The closed routing is the catch. Sakana does not reveal which workers handled a request, may use customer data for training unless the user opts out, and does not serve the EU or EEA. Its [public repository](#) contains setup files, demos, and the report, but no coordinator weights or software licence. That makes Fugu a closed service built from multiple models, not an open model that replaced Fable.

Read More: when [one self-managing agent is simpler than a multi-agent stack](#).

Sources:

- [Sakana AI: Fugu announcement, pricing, and FAQ](#)
- [Sakana Fugu Technical Report](#)
- [SakanaAI/fugu repository](#)
- [OpenRouter: Fugu Ultra pricing and context](#)

***Disclaimer:** For information only. Accuracy or completeness not guaranteed. Illegal use prohibited. Not professional advice or solicitation. **Read more:** [/terms-of-service](#)*