

by
DAVE WANG

Prompt

You are an investment-research analyst tasked with forecasting Google's next AI product rollouts by inferring the company's capability roadmap from research disclosures, technical artifacts, and product surface area changes.

Objective

Build a forward-looking roadmap forecast for Google AI over the next 6, 12, and 18 months (as of February 4, 2026), emphasizing "capability clusters" and the typical research-to-product pipeline. Use the Project Genie / Genie 3 sequence as a core case study for methodology validation.

Key deliverables

1. **Capability cluster map** for Google AI, with definitions and observable leading indicators for each cluster.
2. **Research-to-product conveyor model** describing typical stages, gating mechanisms (tiers, regions, previews), and time lags.
3. **Forecast of likely rollouts** (features, products, integrations, tiers) for the next 6, 12, 18 months, with probabilities and trigger conditions.
4. **Monitoring plan**: a watchlist of sources, keywords, author groups, repositories, release notes, and signals (with an explicit "high signal vs noisy" tiering).
5. **Evidence appendix**: a table of claims with citations, including what evidence supports each forecasted item.

Critical constraints

- Every non-trivial factual claim must be cited.
- Distinguish between "capability disclosed" vs "productized SKU shipped."
- Separate confirmed signals from speculation. Assign confidence levels.
- Prefer primary sources and technical artifacts over commentary.
- Provide counterfactuals and failure modes (why a predicted rollout might not happen).

Part A: Build the taxonomy of capability clusters



by
DAVE WANG

Create a taxonomy with at least these clusters (expand if needed):

- Foundation LLM and multimodal reasoning
- Generative media (image, video, audio)
- World models and interactive simulation
- Agents (tool use, long-horizon tasks, autonomy constraints)
- Embodied AI and robotics
- Serving and cost curves (inference optimization, latency, edge)
- Developer platform and monetization (APIs, pricing, tiering, gating)
- Consumer distribution channels (apps, subscriptions, Workspace, Android)

For each cluster, define:

- What the capability is
- The typical research artifacts that precede productization
- The product surface area where it lands (consumer, enterprise, developer)
- Leading indicators and “tell” signals (docs changes, endpoints, SDK commits, policy language shifts, waitlists, tier gating)
- Typical time lag ranges from disclosure to prototype to broad rollout (use empirical examples)

Part B: Use a case study to calibrate your model

Use the **Project Genie / Genie 3** chain as a calibration case:

- Construct a timeline from research publication and research preview to product prototype access
- Extract observable leading indicators that appeared before broader distribution
- Identify the packaging decisions (who got access, gating, subscription tier, geography, “labs” positioning)
- Use this case to estimate plausible lag distributions and signal reliability



by
DAVE WANG

Required sources for the case study (use these plus any better primary sources you find):

- Google DeepMind blog post on Genie 3 research preview
- Genie paper on arXiv
- Technical report that references Genie 3 (SIMA 2 tech report)
- Reputable press coverage of Project Genie prototype access (TechCrunch, Reuters as examples)
- Optional but useful: broader analysis on market impact (Reuters, Financial Times) and community reaction (Hacker News threads)

Part C: Source strategy and signal stack

Create a "signal stack" with weights (example tiers A to D):

- Tier A: official primary sources (DeepMind, official Google product posts, official docs and release notes)
- Tier B: technical artifacts (arXiv, tech reports, citations, benchmark repos, conference talks)
- Tier C: developer surface changes (API endpoints, SDK updates, pricing tiers, changelogs)
- Tier D: community chatter (forums, Hacker News, Reddit) treated as noisy leads requiring confirmation

For each tier:

- Provide examples
- Define what qualifies as "actionable"
- Define how to validate and how to prevent false positives

Part D: Forecasting methodology

Use a structured forecasting approach:

1. Build a dataset of Google AI capability disclosures and subsequent product rollouts over the last 24 to 36 months (or as far as evidence supports).
2. Identify recurring patterns: naming conventions, preview language, tiering, integration pathways, partner mentions, hardware tie-ins.
3. For each capability cluster, enumerate candidate next rollouts and score them on:
    - Evidence strength
    - Product readiness signals (latency, cost, safety gating, policy language)
    - Distribution pathway readiness (consumer, dev, enterprise)
    - Monetization fit (pricing tier, enterprise SKU, usage-based APIs)
4. Produce a ranked forecast with:
    - Probability bands (example: 60-80%, 40-60%, 20-40%)
    - Expected timing window
    - Key triggers that would increase confidence
    - Key blockers (safety, compute, IP, partner conflicts, regulatory)

Include at least one bearish scenario where rollouts slow or concentrate behind paywalls, plus a scenario where rollouts accelerate due to platform competition.

Part E: Outputs and format requirements

Output format

1. Executive summary (one page, bullet format)
2. Capability cluster taxonomy (table)
3. Research-to-product conveyor model (stages, gating, lags)
4. Case study: Project Genie / Genie 3 timeline and lessons
5. Forecast: next 6, 12, 18 months (ranked list with probabilities and triggers)
6. Monitoring plan and watchlist (keywords, authors, sites, release notes, alert setup instructions)
7. Appendix: evidence table with citations and notes on confidence



by
DAVE WANG

Evidence table requirements

- Claim
- Source(s)
- Cluster
- Indicator type (A/B/C/D)
- Confidence
- Notes and caveats

Part F: Minimum required sources to incorporate

At minimum, incorporate and cite the following categories of sources (add more if helpful):

- DeepMind blog/model pages related to Genie and world models
- arXiv Genie paper
- SIMA 2 technical report referencing Genie 3
- Reputable press on Project Genie prototype access and market impact (TechCrunch, Reuters; optionally FT)
- Community reaction signals (Hacker News; optionally Reddit) as leads only, with confirmation rules

Stretch goal: "roadmap radar" checklist

Produce a concise checklist that a human analyst can run weekly:

- What to check
- Where to check it
- What constitutes a red flag vs a high-confidence leading indicator
- What changes the probability of each predicted rollout