DRAKO AI

Model company thesis

COMPETITION FUNDING AND MARKET ENTRY



Intelligence across space and time

How Drako could turn a live data feed into models sold through APIs, licenses and embedded products.

Prepared for Sean Heiney
27 November 2026

01 THE BUSINESS

Business and opportunity

Train models that understand real places over time. Sell the models through APIs, licenses and embedded products.

6 PB
Footage and events
Refreshing corpus · supplied

100,000+
Brivo locations
Public platform footprint

>$300M
Combined parent ARR
Supplied current figure

The starting point

Supplied premise: a continuously refreshed 6 PB corpus. All customer data is opted in for AI training under the current agreement. Drako is assumed to hold the continuing exclusive license.

The business

Build two models. One understands video. One connects events and context over time. Proposed Brivo and Cobalt deployments provide first use cases, reviewed outcomes and a route to first revenue.

Why this could become a major company

The live feed can supply paired examples and checked outcomes. If those examples improve models across new sites, Drako can sell the resulting capability to other products and industries.

What must be proved

Build aligned, labeled training sets. Show that the paired data improves accuracy, state tracking and prediction. Test on new customers and at practical costs to run the models.

Value depends on a strong research team, measurable model gains and customers who pay for them.

6 PB, 100% training opt-in, parent ARR above $300M and parent value above $1B are supplied premises. Brivo publicly reports 100,000+ locations. Brivo and Cobalt are proposed Drako channels and buyers; no paid model contract is assumed.

Sources [1] Brivo merger [2] Brivo platform [5] Cobalt integration [6] Cobalt acquisition

02 THE MODELS

Two models for space and time

Model 1 understands video. Model 2 connects events over time. Test each separately.

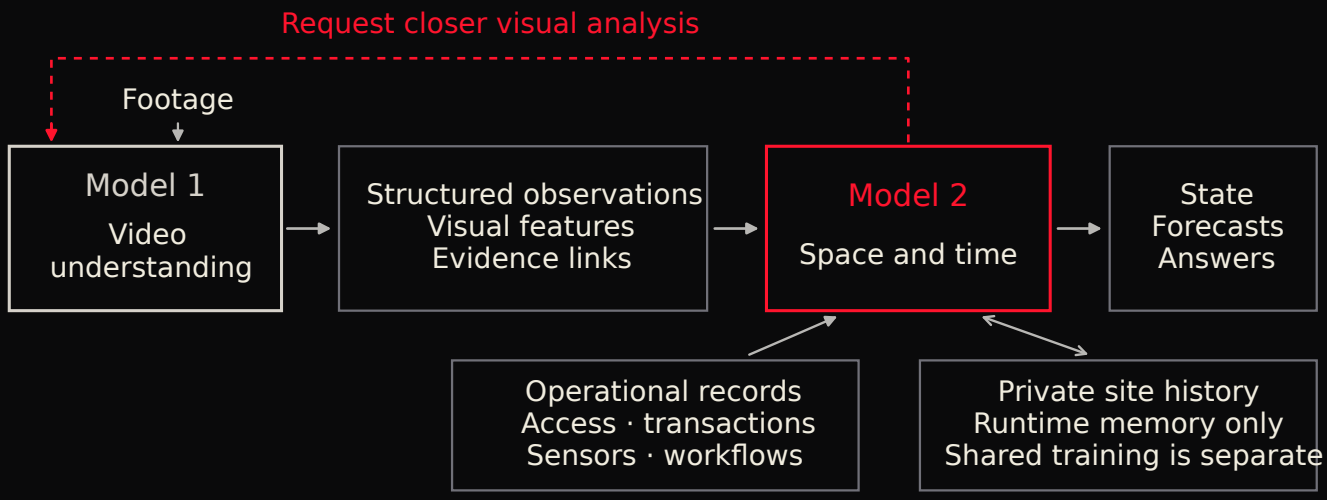

Model 1 — Video. Learn activities and changes. Return events, uncertainty, visual features and evidence links.

Model 2 — Space and time. Combine video with access and business records. Track state and forecast events. Keep private site memory separate from shared training.

World models learn state and dynamics. **Foundation models** can be adapted to many tasks. **Frontier** performance must be proved. [23,24,39]

How to prove the advantage

Test	Question answered
Swap reference and Drako components	Which model adds value?
Hold data and compute fixed; disrupt pairings	Does pairing create an advantage?
Hold out whole customers from all training	Does it work at new sites?

A site history database is useful infrastructure. A second model earns its place only if learned state and prediction outperform a strong model using that same history. Predictions alone do not prove causation.

Proposed research. No trained Drako frontier model or benchmark win is assumed.

Sources [20] Marengo [23] V-JEPA 2.1 [24] V-JEPA 2 [37] Neural Hawkes [39] World models

03 THE FIRST BUILD

First build and first deployments

Build the training system, then test the models in real workflows.

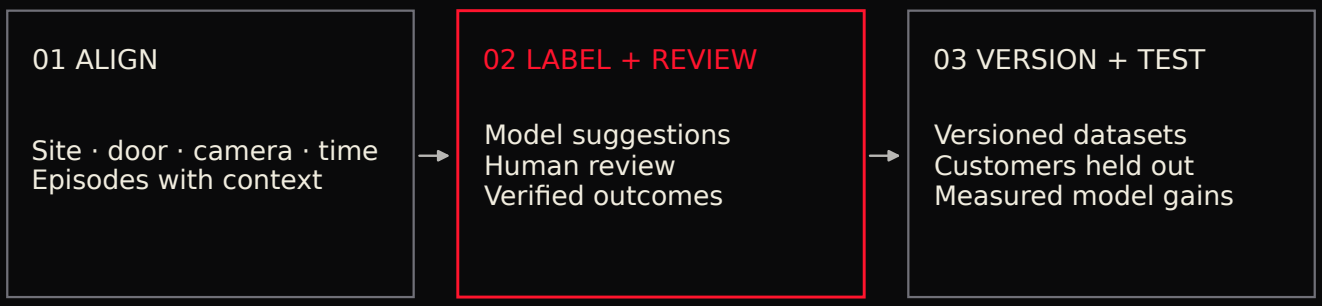

Build the training data system

Build episodes from aligned video and records. Model suggestions speed labeling; human review and verified outcomes create reliable targets.

Version datasets. Hold out entire customers. Track event and arrival times to prevent future evidence leaking into tests. [25]

Use the workflows already in place

Brivo and Cobalt can become first model buyers and review channels. Add model evidence and corrections to their existing queues. Capture actions and later outcomes without rebuilding their consoles. Inventory actuator commands, operator actions and sensor observations separately.

Sample ordinary activity to find missed events. A dismissed alert is not a verified outcome.

Test model value

Compare the two models with strong available alternatives. Measure event accuracy, state tracking, prediction quality, operator effort and cost to run the models. Use a proposed pilot of **6–10 organizations across 50–100 sites** to test deployment and paid value.

The feed is available. Rich training metadata and reviewed labels must be built. Pilots and model licenses are proposed. Cross-industry sales depend on results beyond the first deployment.

Sources [3] Eeva [5] Cobalt integration [7] Cobalt [10] Ambient access [25] TRACE

04 THE MODEL MARKET

The model competitive map

Right means closer to Drako's tasks. Up means more public evidence of access, tests and use.

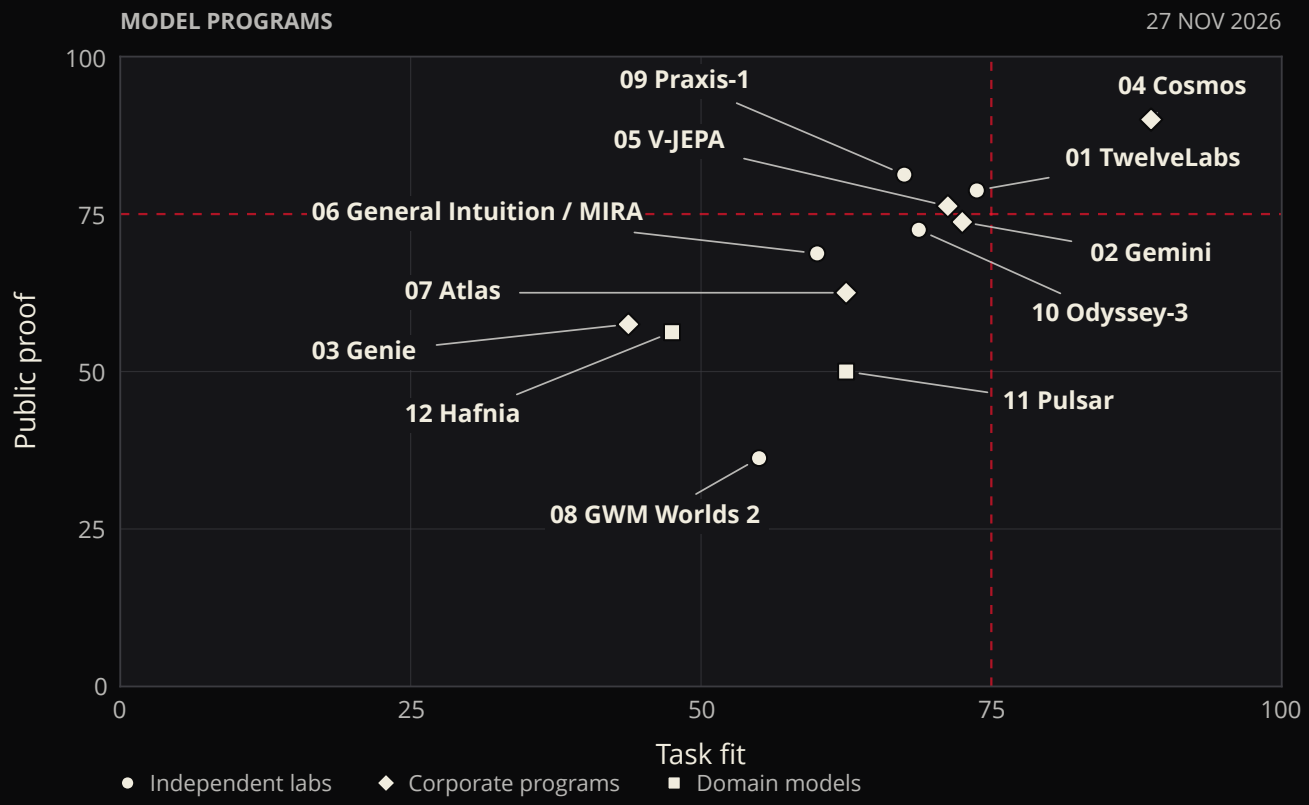

Unranked: Drako AI / AMI Labs

Analyst evidence scores, not accuracy. 75 is a guide, not a frontier threshold.

What the map says

TwelveLabs is a direct baseline for video plus metadata. Cosmos supplies a broad open model family. Gemini combines video with general reasoning. World simulators and action models are adjacent research rivals. [20,64,73]

What Drako must earn

Prove better understanding of observed places over time, useful prediction and lower cost. Access to a large live feed is a starting asset; it is not a model result.

Original assessment, not Gartner research. Twelve selected programs; this is not an accuracy, valuation or market-share ranking. Full scale: 0–100. The 75 guides mean an average rating of 3/4, not a frontier threshold. Differences under 5 points are not meaningful. The two scorecards explain the factors and sources.

Sources [20] Marengo [64] Gemini [73] Cosmos 3 [99] AMI research

05 THE ECONOMICS

Revenue costs and capital

The scenarios test early model sales through Brivo and Cobalt.

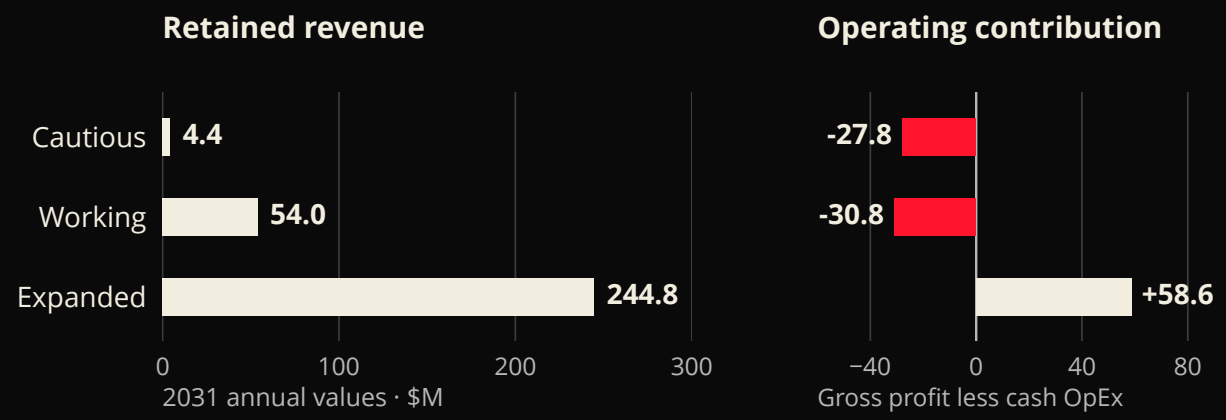

Analyst scenarios for 2031, not forecasts. Annual retained revenue, not year-end ARR.

Working: 100,000 reachable sites × 30% adoption × $2,000 per year = **$60M billings**. A 70% retained share plus $12M from distinct contracts gives **$54M revenue**.



Working scenario in 2031. All amounts in $M.

The terms matter

At 50% retained share, contribution falls **$11.4M**, after the 5% royalty. Another $100 cost to run the models per paid site costs **$3M**. Breakeven needs **56,100 paid sites**, with other inputs fixed.

2027–2031 capital: **$120M / $196M / $177M** for Cautious / Working / Expanded. All need funding beyond the initial $100M.

Illustrative early model sales, not forecasts or year-end ARR. This is not a market-size ceiling or a model-lab valuation. All assume 100% training opt-in; paid adoption is separate. Working OpEx: $44M staff, $20M training/data and $10M other. Capital includes initial 2% financing fees, setup and a $10M reserve; excludes later fees, tax and working capital. Both research programs still need full costing.

Model Drako AI scenario workbook

06 THE FINANCING

Funding and proof

A $100M raise needs a compelling model program and a team that can deliver it.

Financing scenarios

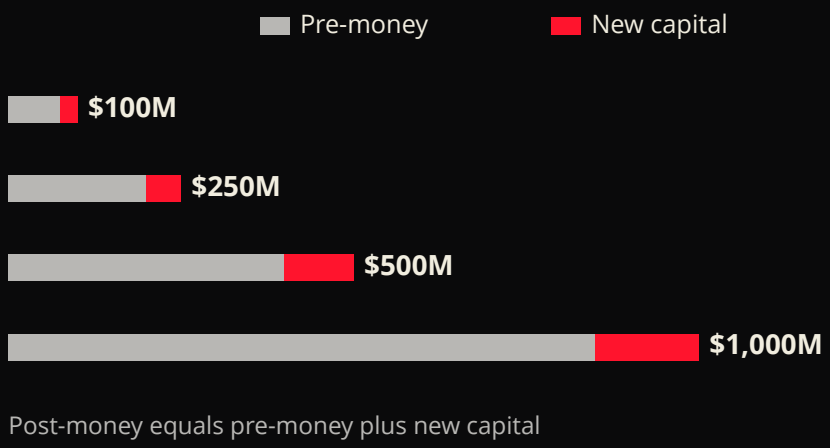

Conditional transaction structures, not current market appraisals.

The $100M example sells **20% ownership**. Both models share the data asset.

The first 18 months

From start	Work and evidence
First 60 days	Recruit research lead; inventory feed; design training sets
Months 2–4	Align and label Brivo and Cobalt data; establish baselines
Months 4–8	Train both models; test gains on unseen customers
Months 8–12	Convert paid pilots; prove delivery economics
Months 12–18	Scale proven gains; test adjacent tasks

Complete the company

Set distribution, revenue share, costs, output ownership and transfer terms. Independently assess related-party contributions and required approvals.

Recruit both program leaders. Cost the full plan before treating $100M as sufficient.

Decision: scale research when the data produces repeatable gains and customers value the resulting models.

Illustrative financing and timing. No team or funding commitment is assumed. Training opt-in is supplied; remaining terms concern the spinout and its economics.

07 TASK FIT

Fit for Drako's tasks

How closely does each model address the work Drako wants to do? This sets its left-to-right position on the map.

Read across a row for one model. Read down a column to compare the same task.

Rating key: 0 no public support found · 1 limited · 2 some · 3 substantial · 4 strongest evidence under that factor's criteria.

Model / sources	Understands video 30%	Tracks events over time 25%	Uses records or actions 20%	Predicts what comes next 15%	Adapts to new tasks 10%	Task fit out of 100
01 TwelveLabs [20,63]	4	3	4	0	2	74
02 Google Gemini [64,67]	4	3	3	1	2	73
03 Google Genie [70,71]	1	2	2	3	1	44
04 NVIDIA Cosmos [73,74]	4	3	3	4	4	89
05 Meta V-JEPA [79,23]	3	2	3	3	4	71
06 General Intuition / MIRA [81,82]	1	3	3	3	3	60
07 World Labs / Atlas [83,84]	2	3	3	3	1	63
08 Runway GWM [87]	1	3	3	3	1	55
09 Runway Praxis [89,91]	2	2	3	4	4	68
10 Odyssey-3 [93,78]	2	3	3	4	2	69
11 Ambient Pulsar [9,96]	3	3	3	1	1	63
12 Milestone Hafnia [97,14]	3	2	1	0	3	48

What to take from this

TwelveLabs scores strongly on interpreting video and joining records. Cosmos spans more prediction and adaptation tasks. A single total can hide those differences.

What the numbers mean

The percentages show how much each factor contributes. The total converts the weighted 0–4 ratings to a 100-point scale. It is **not a percentage of correct answers**.

Analyst judgments from public evidence as of 27 November 2026, not measured model accuracy. A low score can mean limited disclosure. Drako and AMI remain unranked.

08 PUBLIC PROOF

Public proof behind the models

What can outsiders use, test or inspect? This sets each model's bottom-to-top position on the map.

These five factors assess the public evidence behind the model. They do not score how accurately it performs Drako's job.

Rating key: 0 no public support found · 1 limited · 2 some · 3 substantial · 4 strongest evidence under that factor's criteria.

Model / sources	Available to use 30%	Published tests 25%	Outside users 20%	Methods explained 15%	Practical to run 10%	Public proof out of 100
01 TwelveLabs [20,63]	4	3	3	2	3	79
02 Google Gemini [64,67]	4	3	2	2	3	74
03 Google Genie [70,71]	4	1	2	1	3	58
04 NVIDIA Cosmos [73,74]	4	4	2	4	4	90
05 Meta V-JEPA [79,23]	4	3	1	4	3	76
06 General Intuition / MIRA [81,82]	3	3	1	4	3	69
07 World Labs / Atlas [83,84]	4	2	1	2	3	63
08 Runway GWM [87]	1	2	0	3	2	36
09 Runway Praxis [89,91]	4	3	2	4	3	81
10 Odyssey-3 [93,78]	3	4	2	2	3	73
11 Ambient Pulsar [9,96]	3	1	2	1	3	50
12 Milestone Hafnia [97,14]	3	2	1	3	2	56

What to take from this

A model can be relevant to Drako's work while offering limited public access or testing detail. That is why the map keeps task fit and public proof separate.

Use the scores with care

API access, downloadable weights, previews and embedded products have different limits. A high score does not establish deployment quality on Drako's data. Differences under five total points are not meaningful.

Totals use the header weights, sum weight × rating ÷ 4, and are rounded here; the map uses exact values. Public source coverage is broad for the listed programs except Pulsar, whose exact model details rely mainly on company claims. Low disclosure does not prove weak capability. Drako and AMI remain unranked.

09 MODEL COMPETITORS

What the model rivals bring

Compare the training advantage, the product and the remaining gap.

Model program	Data and model advantage	Meaning for Drako
TwelveLabs [20,63]	Video-native models. Aligned metadata fusion and structured event labels. Training mix is closed.	Direct API rival. Prove gains beyond adding metadata.
Google: Gemini / Genie [64,67,70]	Gemini interprets video; Genie generates worlds. Separate models, interfaces and tests.	General reasoning is a strong baseline. Simulation is a different task.
NVIDIA: Cosmos [73,74]	Broad models for reasoning, simulation and action, with downloadable weights and recipes.	A possible starting point and a strong baseline. License terms apply.
Meta: V-JEPA [23,79,80]	Passive-video pretraining; the action branch adds robot data. Research weights are available.	Useful representations. Check the weight license before commercial use.
General Intuition / MIRA [49,81]	Action-linked gameplay; public MIRA research released with partners.	Closest data-rich-parent analogy. Transfer to commercial sites is unproved.
World Labs / AMD: Atlas [83,84]	Proprietary multimodal spatial model with reconstruction and generation APIs.	Possible scene-mapping component. Generated unseen geometry is not observed fact.
Runway: GWM / Praxis [87,89,91]	GWM models prompted worlds. Praxis adds action and robot training after video pretraining.	Passive video can teach dynamics. Control still needs suitable action data.
Odyssey-3 [93,94]	Visual world-model training and action adaptation. Version 3 is a preview.	Prediction rival. Its earlier API does not establish full version 3 availability.
Ambient: Pulsar [9,96]	Vendor reports 1M+ enterprise-video hours; platform joins operating context.	Direct domain rival. Public disclosure does not reveal all private capability.
Milestone: Hafnia [14,97]	Curated video and reviewed labels adapt Cosmos. Traffic API; accident pilot.	Shows how domain data can improve a base model. No access-log fusion shown.

Broader baseline: OpenAI GPT and Anthropic Claude offer general reasoning and vision APIs. Alibaba Qwen offers open multimodal models. Combine frames, event records, retrieval and memory before assuming a second learned model is needed. [102,103,104]

Drako's 6 PB measures storage, not unique hours, useful labels or a training advantage. The proposed advantage is a refreshing feed of linked episodes and reviewed outcomes; the training system still needs building. Open weights can be the starting point. Proprietary training gains are what matter.

10 MODEL LAB FUNDING

Capital is backing model companies

Model companies have raised $100M to over $1B at several stages.

Model company	Date	New capital and round	Disclosed company value
General Intuition [27,44,45,46]	Oct 2025 Jun 2026 Sep 2026	$133.7M seed $320M Series A $220M further funding	Not disclosed $2.3B* $6.2B*
Luma [56,57]	Nov 2025	$900M Series C	Not disclosed
Physical Intelligence [59,60,61]	Nov 2025	$600M Series B	$5.6B post-money, reported
Skild AI [58]	Jan 2026	$1.4B Series C	More than $14B*
Runway [54]	Feb 2026	$315M Series E	Not disclosed
World Labs [50]	Feb 2026	$1B new funding	Not disclosed
AMI Labs [28]	Mar 2026	$1.03B initial financing	$3.5B pre-money
Ineffable Intelligence [29]	Apr 2026	$1.1B seed	$5.1B post-money
Odyssey [55]	Jun 2026	$310M Series B	$1.45B*
TwelveLabs [47]	Jul 2026	$100M Series B	Not disclosed

*The cited announcement does not specify pre- or post-money. Physical Intelligence's amount and stage are investor-confirmed; its valuation is reported by Bloomberg. These are dated records, not claims that no later financing occurred.

What the stages tell us

General Intuition is the closest structural example: parent data, specialist researchers and a separate model business. Its game inputs provide action labels. Drako's footage and event records must earn their own technical case.

AMI and Ineffable show that exceptional research teams can raise large initial rounds before mature products. The other listed rounds came after research or product progress. There is no single financing sequence for every lab.

Selected 2025–26 announcements, not a market total. New capital, cumulative funding and company values are different measures. The financings do not establish a price per petabyte or an inception value for Drako.

11 CAPITAL AND STRATEGY

What the funding movement means

A strong research team, useful models and a route to customers can support a large raise.

OCT 2025		JUN 2026		SEP 2026
$133.7M	→	**$320M**	→	**$220M**
Seed		Series A		Further funding
Value not disclosed		$2.3B company value*		$6.2B company value*

General Intuition financing announcements. *Pre- or post-money basis is not stated.

Parent data can help launch a model lab

Niantic Spatial launched in **May 2025 with $250M**: $200M from Niantic and $50M from Scopely. This is parent and transaction-linked capital, not an independently priced $250M seed. [30,48]

Strategic buyers value more than datasets

AMD announced an approximately **$8.2B all-stock acquisition of World Labs** and closed it on **13 November 2026**. The deal includes the team, models and products. It is separate from World Labs' February $1B financing. [50,51,52]

The first channel helps the business case

Brivo and Cobalt can provide deployment settings, reviewed outcomes and potential first model customers. Other security platforms are data owners, potential buyers and bundled alternatives. They belong in the commercial context around the model business.

What Drako should take to investors

Name the research leaders. Show a working training-data system and the first signal that paired data improves models. Set the rights and economics for external model sales. Cost both research programs and delivery. A rare team can earn funding before a benchmark; without that team, early results matter more.

General Intuition values are quoted company valuations with pre/post basis unspecified. The World Labs deal is an acquisition, not funding. Brivo and Cobalt are proposed channels and buyers; no paid Drako model contract is assumed.

Sources [27] General Intuition seed [45] General Intuition A [46] General Intuition follow-on [30] Niantic launch [52] AMD closing

12 THE ENTRY STRATEGY

Where Drako could win

Understand events across cameras and records as they happen.

The first model product

Sell a video model that returns events, evidence and uncertainty. Pair it with a model that updates what is happening at a site and predicts useful next events. Start with ambiguous access episodes, changing evidence and unresolved conditions.

Software developers are the buyers. Brivo and Cobalt are proposed first license customers and test environments. Their workflows can supply reviewed examples without Drako rebuilding the customer products.

Why the data could matter

Access records anchor video in time and place. Later outcomes can turn uncertain episodes into checked training examples. The advantage must show up in model quality and cost. TwelveLabs already offers time-aligned metadata fusion. [20]

Test	Evidence to earn the next step	What would weaken the thesis
Learn more	Beat a strong video model using the same records and retrieval	A general model and database perform as well
Work elsewhere	Hold out entire customers and later time periods	Gains require hand setup at every site
Run affordably	Measure cost per monitored hour and resolved episode	Cost to run models consumes the value created
Build a model business	Paid licenses and durable rights to sell models	Value stays entirely with one distributor

The wider opportunity

Expand from embedded models to developer APIs and private deployments. Passive video can teach useful dynamics. [89] Prediction and simulation are research options; control claims need suitable action data and tests.

Proposed wedge and tests. Compare each model and the pair with a matched fine-tuned baseline. Equal results on Drako data test the need for a new architecture, not the value of its data. Also test available models and commercially available data. Remove pairings to test their contribution. Prevent future evidence leaking into earlier decisions. No paid model contracts or cross-industry demand are assumed.

Sources [20] Marengo [89] Praxis-1