

Research systems, and the *evidence* they produced.

Two things live here. A research instrument built over the complete back catalogue of a documentary show — every claim extracted, classed as **fact**, **inference** or **speculation**, scored against an auditable rubric and checked against primary sources. And the products behind it: real interfaces, lifted from the repositories that shipped them, running on their own data with identities removed.

CANONICAL EPISODES	HOURS READ	CLAIMS CLASSED
136	362.1	801
RELATIONS	DOCUMENTS INDEXED	ACTUALLY RETRIEVED
624	122	83

THE RESEARCH INSTRUMENT

Built for a specific application. The corpus is the show's canonical episodes, and every citation deep-links to the exact second of source video. Nothing on this page is a figure someone typed — the numbers above are read live from the file the pipeline generates.

Findings

What the catalogue says, and what it never said out loud. The site-signature that splits into two disjoint families, the Cydonia coordinates in a 1984 tasking document, the evidence ladder, and the document index with its honest retrieval rate — including the documents that were named but could never be found.

9 tracks complete

generated from files on disk

no figure typed by hand

Evidence control room

The interactive graph. Every entity, claim and relation, filterable by class and confidence, with the claim-class breakdown drawn as a ring on each node.

live re-merge

timestamp deep-links

Anomaly atlas

Twenty-eight places on a globe, scored against one seventeen-point checklist instead of a vibe: the twelve the catalogue treats as anomalous, plus sixteen the show has never mentioned, scored the same way so the checklist can be seen failing as well as passing. Every point carries the episode and timestamp it came from, one site is a deliberate negative control, and five of the twelve are places the show itself never connected to the pattern.

87 cited observations

one checklist, all 28 sites

negative control included

Guest prep cards

Briefings for guests the show has not booked yet: every point where each person already appears in the catalogue, the questions only a full-catalogue reader could ask, and the claims a fact-checker should pre-check.

timestamp deep-links

live news verified

Mars — the adjudicated thread

Claims with their eight scoring components open to inspection, threads adjudicated against retrieved documents, and the distinction between a genuine session document and evidence about Mars.

documents retrieved

real archive identifiers

THE PRODUCTS

Each case embeds the interface that actually shipped. Where the real screen holds other people's data it is rebuilt with the same layout and mechanics on anonymised records, and the page says so. Nothing here is a mockup of a product that never existed.

Attendee research engine

People explorer, the hackathon event report, impact timelines and company ROI — all lifted from the shipped product, on anonymised records.

13,619 profiles scored

\$1,861.47 measured spend

Partner scoring

The approvals workspace, the real weights on the page, the automation workflow rendered from its own export, and one production prompt annotated section by section.

Reach 40 · Fit 35 · Engagement 15 · Identity 10

Investor research

The research dashboard, task creation and investor circles. Profile pages are labelled reconstructions on invented data rather than passed off as real.

995-row target set

3 delivered dossiers

Mars Arena — scored candidate atlas

The leaderboard and inspector, with the seed set and the full catalog counted separately rather than summed into one flattering number.

12 seed

1,062 catalog

Moon Arena — a board that measures its own ignorance

Eighteen lunar pits scored on four axes, with a second reading whose only job is to show how much of the ranking rests on measurements nobody has taken. The gap between the two readings gets its own column.

18 of 278 atlas pits

2 measured on all four axes

Aleksei Beliankin-Bauer · ab7.ai/american-alchemy · figures read from the generated findings file, not typed. Paper cannot be interrogated, so the instruments are live: all 850 entities at ab7.ai/graph/site, the twelve scored anomaly sites at ab7.ai/atlas.