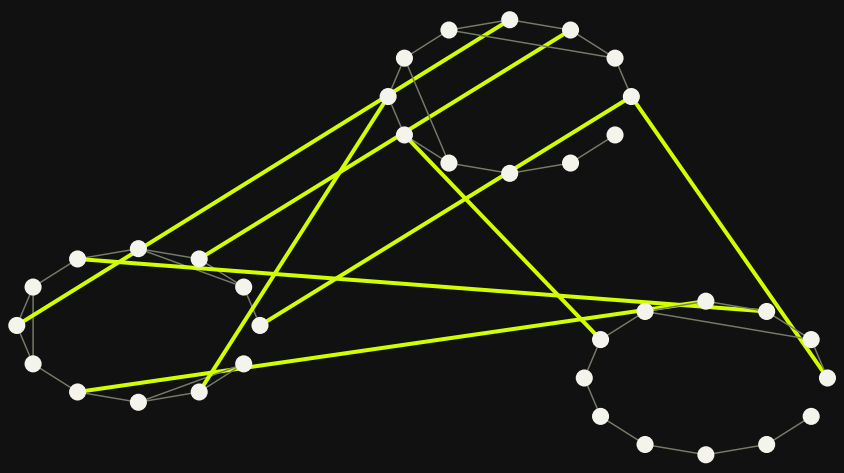
FIELD STUDY 002 / 24 OCTOBER 2026

DESIGN A CONFERENCE THAT LEARNS.

An organizer research paper on ontology, confirmed networks, conference activity and prospective decisions.



NETCUES x FALKORDB

36 profiles. 56 confirmed encounters. 47 distinct pairs. 1 connected component(s).

Companion: index.html / Data: paper-data.json / This document contains clickable navigation and source links.

READER MAP

READ THE OPERATING PICTURE.

Start with the observed record, inspect the graph, then examine a proposed decision.

03 The study in one view >

04 A typed conference ontology >

05 Invitation states and outcomes >

06 Cross-specialty encounter matrix >

07 The confirmed network >

08 Booth interest and visits >

09 From observation to useful context >

10 Registration and explicit attendance >

11 Professional cohorts and scope >

12 A prospective prediction protocol >

13 Probability estimates and baselines >

14 Held-out evaluation >

15 Language as a hypothesis layer >

16 An organizer decision scenario >

17 Executed operation ledger >

18 Sources and data index >

Every page footer returns here. Figures link to the relevant evidence section. Rich exploration lives in the companion HTML.

01 / OBSERVED SYSTEM

THE STUDY IN ONE VIEW.

An actual application workflow with a controlled synthetic conference dataset.

36
REGISTERED PROFILES

56
CONFIRMED EVENTS

47
DISTINCT PAIRS

1
COMPONENTS

DECLARE	ACT	CONFIRM	COMPUTE
Context and interests	Typed product records	Two-person outcome	Graph and forecast

A useful conference system has to do more than display activity. It needs to connect an organizer question to a precise unit of observation. This study combines browser-driven actions with a prospectively executed API workflow, then reads the persisted cloud graph at named checkpoints. The application creates the records; the analysis preserves their identifiers, meanings and sequence.

The hosted workflow exercises invitations, target engagement and booth visits. Separate identities initiate and close encounters. It retains the prior network without resetting the dataset. A separate virtual-clock rehearsal records session attendance and professional cohorts. Its 46-profile catalog and 36 study actors remain distinct from the 36-profile hosted prediction study.

Hosted graph and prospective invitation workflow; session attendance and professional cohorts come from a separate virtual-clock rehearsal.

02 / ONTOLOGY

A CONFERENCE, EXPRESSED AS RELATIONSHIPS.

Each edge family answers a different organizer question.

ATTENDEE

36

People with declared context and individual activity.

COMPANY

31

Professional affiliations available to typed queries.

TOPIC

83

Concepts connecting specialties and conference targets.

INDUSTRY

18

Industry entities retained in the ontology.

CONFERENCEITEM

21

Booths, talks, workshops and speaker targets.

INTERACTION

75

Requested and closed encounter records.

The property graph connects professional declarations, shared topics, conference targets and recorded actions. An attendee can declare a specialty, relate to an industry or company, save a session, request a route and participate in an encounter. Each relation retains its own meaning. The same participant ID can therefore support an introduction explanation and a later attendance or network query.

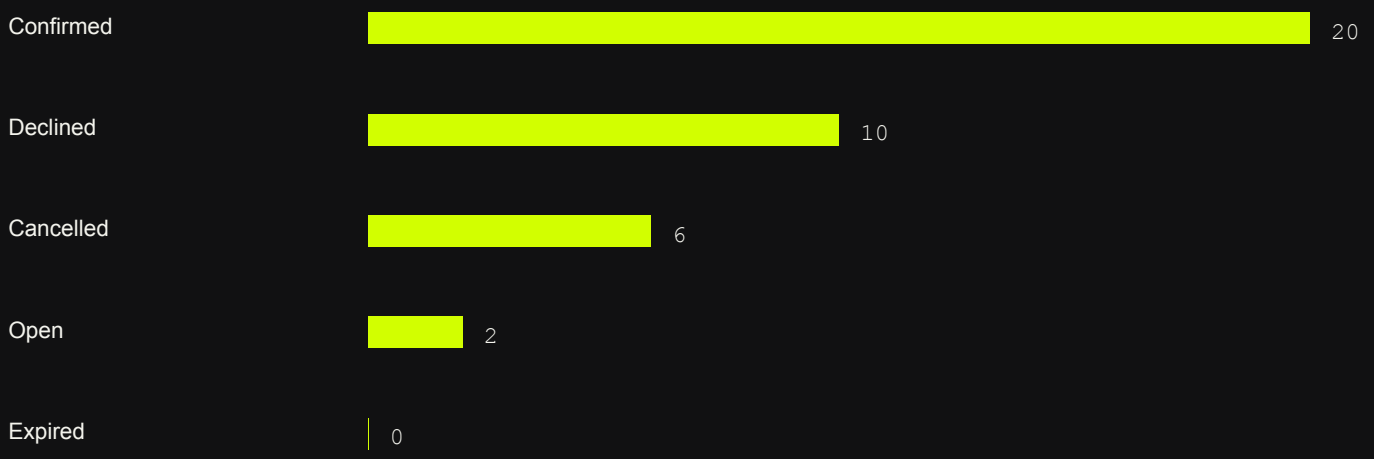
The organizer should choose the projection that matches the decision. Shared topics help explain relevance. Current intent helps plan activity. Explicit check-in history records attendance declarations. Confirmed encounters form the observed social graph. Traversing a wider ontology is useful for context; selecting a narrow relationship family makes a metric interpretable.

Figure 1 / Entity counts from the final captured graph. Label counts can overlap when one vertex has multiple labels.

03 / INVITATION LEDGER

FOLLOW THE REQUEST TO AN EXPLICIT STATE.

One request identity; several possible outcomes.



The controlled workflow created 38 invitation records. The ledger groups them by their captured state, while retaining the requested, confirmed or closed timestamp for each ID. A confirmed record is an encounter event. A decline or cancellation is an explicit alternative outcome. Open and expired requests remain separate categories in the descriptive accounting.

An invitation funnel is useful when its population and response window are explicit. The numerator should correspond to the same eligible request set as the denominator. This study uses explicit closure for the prospective prediction endpoint: confirmed versus declined or cancelled. Unresolved requests contribute to the organizer ledger but are excluded from those binary predictive labels.

Figure 2 / Controlled-workflow invitations, rather than initialized fixture encounters. Select a state in the HTML to inspect records.

04 / ENCOUNTER MATRIX

WHICH SPECIALTIES ACTUALLY CONNECTED?

The matrix counts events; a repeated pair can contribute more than once.

	AI	HARDWARE	ROBOTICS
AI	13	0	0
HARDWARE	2	14	0
ROBOTICS	3	7	17

21.4%

CROSS-SPECIALTY ENCOUNTERS

12 of 56 events with both specialties known

Each cell counts confirmed encounters using the two specialties stored when the encounter was confirmed. This makes the classification event-specific. A later profile edit does not rewrite the recorded context of an earlier interaction. The cross-specialty fraction divides encounters with differing known specialties by all encounters with two known specialties.

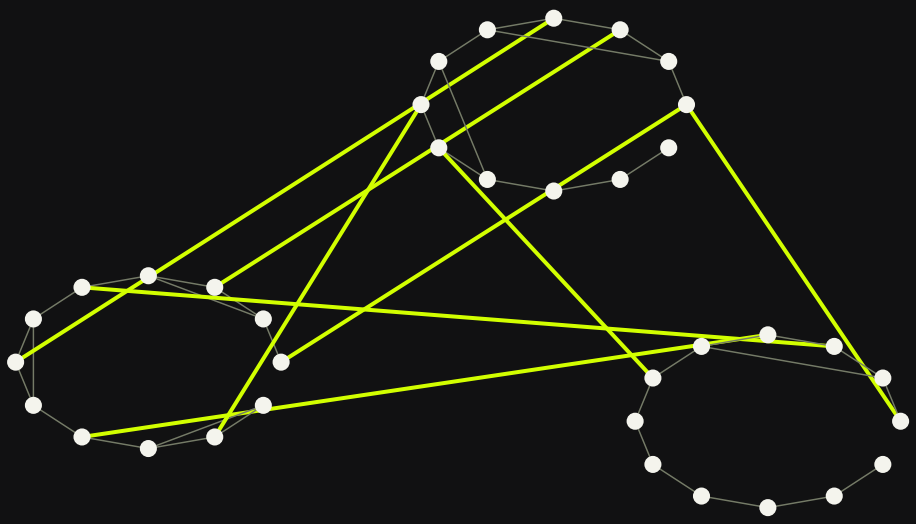
The final graph contains 12 cross-specialty encounters among 56 encounters with both specialties known. The HTML matrix supports record-level inspection: a selected cell reveals the exact interaction IDs that produced it. This gives an organizer a way to inspect where cross-field contact occurred and distinguish repeated contact from broader pair coverage.

Figure 3 / Rows are initiator specialty and columns are recipient specialty at confirmation. Direction records the request roles.

05 / NATIVE GRAPH ANALYSIS

THE NETWORK IS A DIFFERENT QUESTION.

Observed reachability over eligible attendee vertices and confirmed pair edges.



The final observed projection contains 36 eligible attendee vertices and 47 distinct MET edges. FalkorDB's native weakly connected component procedure partitions those vertices by confirmed paths, ignoring stored edge direction. The component sizes are 36. An independent traversal checks the same sorted member sets.

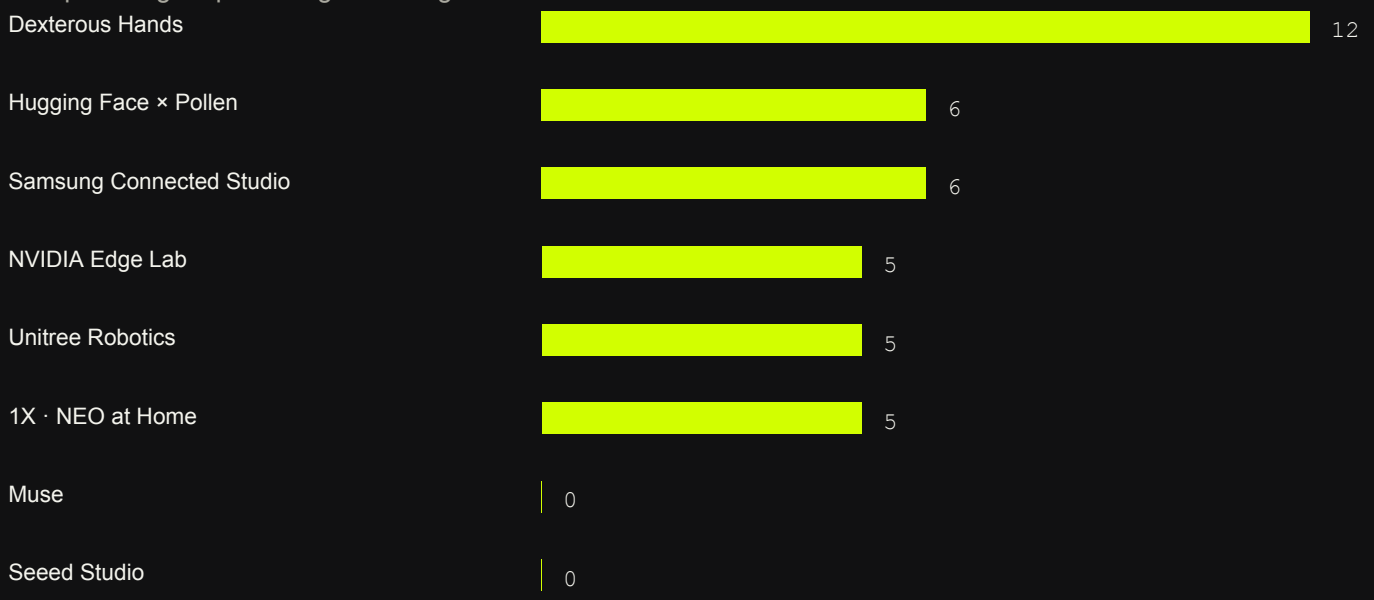
A bridge between components changes reachability because paths can now cross between their members. A repeated encounter on an existing pair changes event history while preserving the pair edge. Both outcomes can be useful, but they serve different organizer objectives. A network designed for broad introductions should inspect new pairs and bridges; a program designed for deeper collaboration should also retain repeat encounters.

Figure 4 / Three-dimensional layout projected onto the page. Positions are a visual arrangement; edges and membership are recorded facts.

06 / BOOTH INTELLIGENCE

INTEREST HAS SEVERAL EXPRESSIONS.

Compare target-specific signals using their own units.



A detail view is a recorded action. A save is an active relationship in the attendee plan. A route request is another intent signal. A booth visit begins with an explicit declaration and has its own state, source and expiry. These records should be compared side by side before attempting to interpret an ordered journey.

The vendor and organizer views share the same target identity but use different authorized scopes. Vendors inspect their booth; organizers inspect the conference. Distinct visitors answer breadth of participation, while event counts retain repeated activity. The interactive companion offers a booth selector and the corresponding views, saves, routes, visits and unique visitors, allowing an organizer to inspect the supporting mix.

Figure 5 / Recorded booth visits by target. These are declared visits, and distinct visitors are calculated separately.

07 / ACTIVITY SEMANTICS

A USEFUL JOURNEY NEEDS STABLE IDENTITIES.

Connect records through attendee, target and time.

19	19	19	12
DETAIL VIEWS	ACTIVE SAVES	ROUTES	VISITS

Dexterous Hands

The counts above describe parallel signals for one booth. Their identities make more specific questions possible: who viewed before saving, who requested a route after saving, and which participants later supplied a booth check-in? An ordered analysis must match the same attendee and target, then enforce the intended temporal sequence.

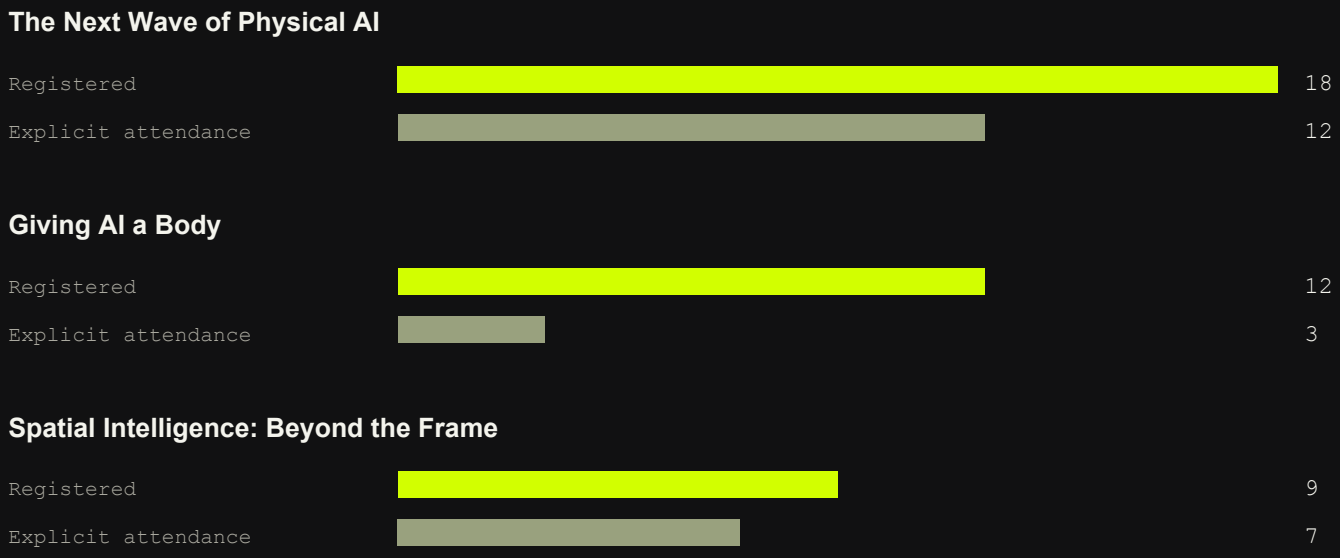
This is where relational intelligence becomes practical. A shared target ID can connect an interest declaration to a visit record, while professional context provides an eligible cohort. The organizer can define a follow-up, record its time and inspect subsequent activity within a stated window. The result remains a traceable record-level comparison rather than a narrative assembled from unrelated totals.

Figure 6 / Current saves and recorded events have distinct counting rules. The HTML presents each unit explicitly.

08 / SPEAKER INTELLIGENCE

KEEP REGISTRATION AND ATTENDANCE DISTINCT.

Intent and explicit check-in are separate observations.



Registration is a current relationship indicating intended participation. Explicit attendance is collected through a check-in action and retained as a separate observation. The organizer can compare both sets for a session and inspect their intersection. A catalog selection or a saved talk contributes to intent, not automatically to attendance.

This distinction supports useful speaker questions. How much recorded interest existed before the session? Which participants checked in? What fraction of the eligible registered set supplied an attendance record during the defined window? The answers share a session identifier, but each needs its own collection rule and denominator.

Figure 7 / Separate virtual-clock rehearsal (78 seconds of wall time): 12/18, 3/12 and 7/9 registrants supplied explicit check-ins; 22 total Attendance records.

09 / RELATIONAL INTELLIGENCE

PROFESSIONAL CONTEXT CHANGES THE QUESTION.

Use voluntary declarations and the correct population for each aggregation.

36

COHORT-ELIGIBLE PROFILES

46

REHEARSAL CATALOG PROFILES



The same network can answer several professional questions: which industries mix, which companies reach outside their existing contacts, and which topics connect a booth to interested participants? A cohort definition specifies the current declared fields and the permission rule used to include each attendee.

Coverage, encounter frequency and mixing have different denominators. Coverage counts distinct cohort members with at least one eligible encounter. Per-capita frequency counts encounter endpoints per eligible member. Mixing counts interactions between two selected groups. Recording these definitions allows an organizer to compare like with like and inspect the cohort definitions and contributing record rules.

Figure 8 / Separate virtual-clock rehearsal: 36 cohort-consenting profiles in a 46-profile catalog; four companies of 9 members.

10 / PROSPECTIVE EXPERIMENT

FORECAST BEFORE YOU KNOW THE OUTCOME.

Participant-disjoint training and held-out evaluation.

24
TRAINING EPISODES

12
HELD-OUT EPISODES

12
DISJOINT PAIRS

SEAL	FIT	FREEZE	SCORE
Outcome writer commits a stochastic script before fitting.	Estimate probabilities using training outcomes only.	Save each forecast before its outcome is dispatched.	Evaluate the same held-out cases for every estimator.

The study executes 36 controlled synthetic invitation episodes across 12 disjoint pairs. Eight pairs supply 24 training episodes; four unseen pairs supply 12 held-out episodes. Each pair moves through recorded booth intent stages: detail view, save and route request. The endpoint is explicit confirmation versus explicit decline or cancellation.

The outcome writer seals its stochastic policy and draws independently of the modeler. The forecasting process receives training outcomes and held-out pre-outcome features. Each held-out probability is written before the corresponding result is dispatched. No attendee appears in both the training and held-out sets, and the three episodes for each held-out pair remain visible in pair-level evaluation.

Figure 9 / The experiment evaluates a prospective research model against controlled app-generated outcomes.

11 / PROBABILITY ESTIMATES

A MODEL EARNS ITS ADDED COMPLEXITY.

Compare fitted intent tiers with economical, predeclared alternatives.

Intent stage	Training n	Confirmed	Fitted p	Rule p
View	8	3	0.400	0.25
Save	8	5	0.600	0.50
Route	8	6	0.700	0.75

```
Intent-conditioned estimate: p(t) = (confirmed(t) + 1) / (episodes(t) + 2)

Constant prevalence: p = (all confirmed + 1) / (all episodes + 2)
```

The fitted estimator uses a Beta(1,1) prior separately for each of the three intent stages. Its posterior mean is the smoothed confirmation frequency. Parameters freeze after all training outcomes. The primary feature is the recorded intent tier; same-specialty status and shared-topic count are retained as context rather than silently added to the model.

Two economical alternatives make the comparison useful. A smoothed training prevalence assigns the same probability to every case. A fixed monotone rule assigns 0.25, 0.50 and 0.75 to view, save and route stages. When available, a validated language-model estimate is frozen from the same training aggregates before evaluation. Every estimator is scored on the same held-out episodes.

Figure 10 / Training-derived values are frozen before held-out outcomes. The rule probabilities were specified in advance.

12 / HELD-OUT RESULTS

SCORE THE PROBABILITIES, THEN INSPECT THE CASES.

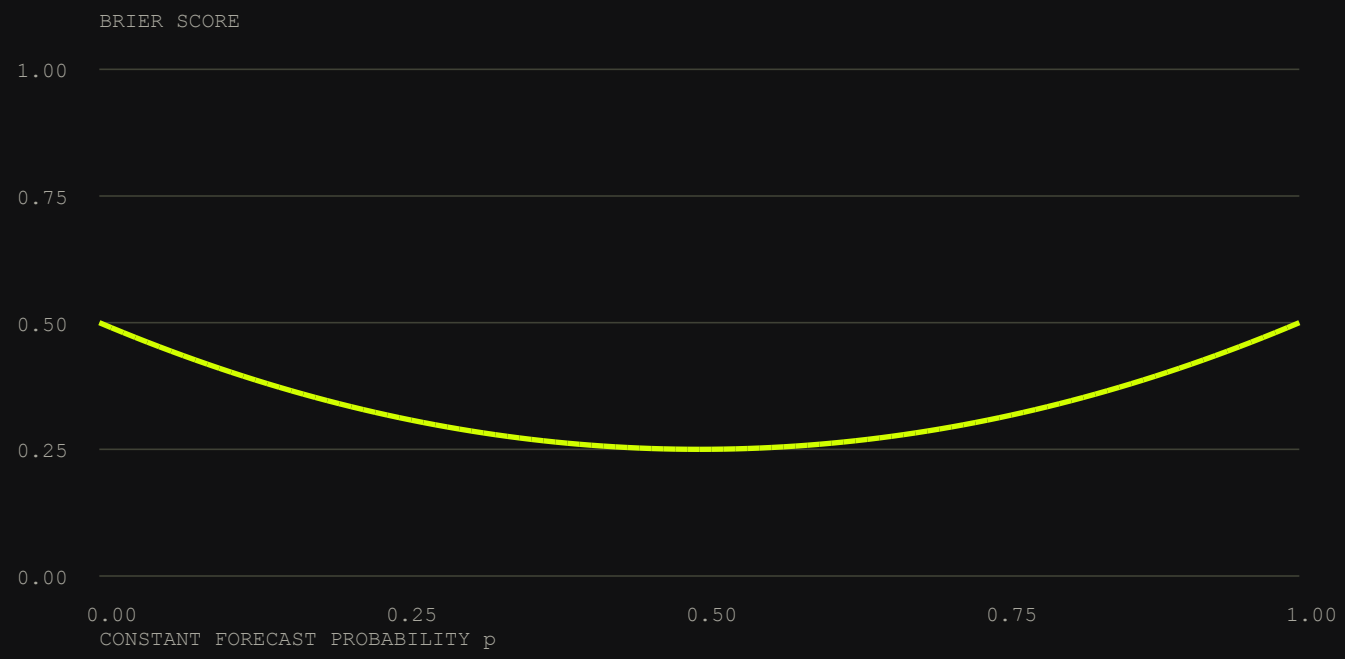
Mean binary Brier score: smaller is better.

0.1667 FIXED INTENT RULE

0.1771 OPENAI BLIND

0.1867 FITTED INTENT

0.2559 TRAINING PREVALENCE



Across 12 held-out episodes, the lowest recorded mean Brier score is 0.1667 for Fixed intent rule. The score is the mean of (p - y)^2, where y is the binary observed outcome. The chart shows how a constant probability would score on those same recorded outcomes; the interactive companion lets the reader move that comparison probability.

The evaluation also retains logarithmic loss, thresholded classifications, intent-tier reliability points and pair-level contributions where supplied. These views help distinguish a useful probability estimate from an apparently good aggregate. Three episodes share each held-out pair, so the record includes pair-level mean differences alongside the episode average.

Figure 11 / The curve is a scoring sensitivity analysis over fixed observed outcomes, not a refitted production model.

13 / HYPOTHESIS LAYER

USE LANGUAGE TO SHARPEN THE NEXT TEST.

Generated proposals stay connected to their supporting facts.

01

Use the fixed intent rule as the next comparison baseline

Carry the predeclared view/save/route probabilities 0.25, 0.50 and 0.75 into the next prospective evaluation, alongside the fitted model. This rule had the lowest Brier loss in the completed held-out trial.

02

Show probability error and calibration together

Display Brier loss beside tier-level mean probability and observed confirmation fraction. Keep each numerator, denominator and probability estimate visible.

03

Repeat the frozen protocol on new participant pairs

Collect the next cohort without changing the current forecasts, then assess whether the ranking and calibration persist across new pairs.

The language layer receives a bounded set of aggregate training facts or final study results and returns a structured proposal. The saved request and response metadata identify the actual model call, the input hash and the validated output. Referenced fact IDs are checked against the supplied evidence set before the proposal enters this paper.

A useful proposal connects a decision to a mechanism, an eligible population and an observable outcome. It should say what action to try, what result would support it and what observation would change the decision. The interactive hypothesis ledger allows readers to inspect those elements together rather than treating generated prose as another graph measurement.

Two actual gpt-4.1-mini-2025-04-14 calls: 3,831 tokens. Training and result-phase response IDs appear on page 17.

14 / DECISION SCENARIO

MAKE THE ORGANIZER ASSUMPTIONS ADJUSTABLE.

An explicit calculation connects a proposed action to an expected result.

Invitations	Assumed completion	Capacity	Expected events	Expected new pairs
10	40%	12	4.0	2.8
20	40%	12	8.0	5.6
30	40%	12	10.9	7.7
40	40%	12	11.9	8.3
20	60%	12	11.1	7.8

```
X ~ Binomial(n,p); E[events] = sum_k min(k,capacity) P(X=k)
```

The observed graph supplies the starting state. The scenario introduces four explicit assumptions: proposed invitation count, completion probability, available encounter capacity and the share of completions that create a new pair. These inputs are adjustable in the companion. The scenario assumes independent completions and computes the exact expected count after the capacity cap. Expected new pairs are a specified fraction of that expectation.

This calculation helps an organizer identify the sensitive assumption. Increasing invitation volume has diminishing effect as the completion distribution approaches capacity. Improving completion probability matters differently before and after that boundary. Component change additionally requires the actual endpoint pairs: a new edge inside an existing component and a bridge between components have different structural effects.

Figure 12 / Illustrative planning arithmetic, with 70% assumed new-pair share. These values are explicit scenario inputs.

15 / EXECUTION LEDGER

INSPECT THE ACTUAL EXECUTED OPERATIONS.

A compact trail from visible actions to stored records and model responses.

BROWSER UI

4 recorded actions: room check-ins, invitation and recipient confirmation.

OPEN RECORD >

HOSTED API

188 requests; 187 successful. One late-signup 409 is retained.

OPEN RECORD >

TIMED REHEARSAL

232 requests; 231 successful. One bad-resource 404 was retained and recovered.

OPEN RECORD >

FALKORDB READ-BACK

21 query receipts; rehearsal read-back found 22 Attendance records.

OPEN RECORD >

OPENAI RESPONSES / TWO ACTUAL CALLS / 3,831 TOKENS

TRAINING / gpt-4.1-mini-2025-04-14 / 1663 TOKENS

resp_0bfb8b13524d0d49016add41f32f6487d09e4a09c8edb27e9c

RESULTS / gpt-4.1-mini-2025-04-14 / 2168 TOKENS

resp_0ad36d3d294a5222016add42937fe487d092f0109bc758d581

Hosted bundle dpl_7sXc9wEzNTeWAHF4Z8Es9pH61fnf; READY. Main base 8a37ebf, gitDirty=1. Alias matched at study start and final verification. Deployment ID identifies the executed bundle.

Linked records preserve timestamps, operation identities, exact outputs and validation results.

16 / SOURCES AND DATA

OPEN THE SUPPORTING RECORD.

Primary technical sources and the complete artifact index.

01

FalkorDB: Weakly Connected Components

Native reachability components over selected vertex labels and edge types.

02

FalkorDB: Cypher language reference

Graph patterns, filters and projections for the recorded ontology.

03

FalkorDB: GRAPH.RO_QUERY

Read-only graph exports used to inspect stored records.

04

Easley and Kleinberg: Networks, Crowds, and Markets, Chapter2

Paths, components and independent breadth-first traversal.

05

Gneiting and Raftery (2007), Strictly Proper Scoring Rules, Prediction, and Estimation

Quadratic/Brier and logarithmic scores for probability forecasts.

06

OpenAI Structured model outputs

Strict JSON-schema Responses output used for the language layer.

07

OpenAI Responses create

Responses request and returned model/usage metadata.

THE RESEARCH PACKAGE

index.html - interactive figures, record inspectors and scenario controls.
paper-data.json - measured values embedded in this document and companion.
data/ - snapshots, action records, query manifest and workflow commitments.
model/ - frozen probabilities, forecast issue times, outcomes and evaluation.
evidence/ - retained screenshots from the live interface.

NetCues organizer research / 24 October 2026 / Standalone artifacts prepared for review.