



Tinder for Bio-Risks

Matchmaking Infrastructure for Surfacing Subtle Biohazard Threats

23.11.2025

Aayush Kucheria, Aditya Adiga, Kuil Schoneveld

[Groundless.ai](https://groundless.ai)

Colva Beach, Goa, India (currently, otherwise global)

Platform prototype hosted at: <https://live-infrastructure.vercel.app/>

Performed as part of: [Apart Research Defensive Acceleration Hackathon](#)

Abstract

This report describes our approach to building a [Live](#) Coordination system. This early warning system aims to identify key threat vectors from *uniquely subtle* conjunctions of medical abnormalities (with possible generalizations to other domains).

The central issue we aim to address is navigating the bureaucracy of sharing information that could lead to inferential cascade risks; where otherwise unimportant local occurrences become critical when combined, such as in the global ecosystem.

Abnormalities from various sources are given distributed attention to surface potentially relevant threat considerations, while automatically maintaining variable levels of privacy compliance. Our infrastructure acts as a just-in-time pairwise matchmaker between lone hospitals, medical labs, and health agencies without the need for a central authority, allowing them to find organisations they can collaborate with, enabling early detection and action towards the identified trends of abnormalities.

Central Constraints

The domains with which we are most concerned are shaped by the following constraints:

1. **Authority Vacuum** – There exists no central authority for dealing with distributed issues.
2. **Contextual Dependence** – Preserving rich contextual information is necessary and infeasible across boundaries. For example national borders, organizational norms, or varied compliance requirements.
3. **Informational Pragmatism** – Privacy preservation is necessary, but the extent to which a given body of information is useful is uncertain. This means human agency is crucial.
4. **(Twice) Bounded Temporality** – The speed of decision-making is critical. Naively fast reactions could cause panic, while slow reactions could miss the window to act.

Fundamentally, the chain of command is optimized for protecting individual privacy and institutional autonomy, not for enabling the cross-boundary information synthesis that pandemic response demands. A helpful step interpreting these constraints in terms of a [Live Theoretic](#) solution is articulated in [Appendix B](#) below.

Related Literature

Current Early Warning Systems and Their Limitations

Early warning of biological threats relies on diverse surveillance systems—syndromic surveillance tools like CDC's EARS monitor health indicators, while open-source platforms like ProMED-mail and BlueDot scan news and web data for outbreak signals. BlueDot famously flagged the Wuhan pneumonia cluster in December 2019 before official alerts, demonstrating the speed possible with AI-driven scanning.

However, **coordination across organizational boundaries remains the critical bottleneck**. Current technical solutions offer two inadequate extremes: centralized platforms like CDC's BioSense require extensive standardization and trust in a single authority, while federated approaches like Lyu et al.'s "federated epidemic surveillance" (2025) share only minimal statistics, preventing the rich context needed for early threat correlation.

The Coordination Problem

The core challenge is not technical but organizational. A US GAO review found that agencies hesitated to engage with NBIC's biosurveillance integration efforts because they "did not trust NBIC to use their information appropriately" and were unconvinced of the value they'd receive. During the 2003 SARS outbreak, China's initial withholding of information contributed to global economic losses of \$40-80 billion—countries face conflicting incentives where reporting triggers trade restrictions rather than assistance.

Research by Laxminarayan et al. (2014) demonstrated that **positive incentives work**: when an international mechanism provided free vaccines to countries reporting meningitis epidemics, previously silent countries began reporting more frequently. The lesson is clear: institutions share early warnings only when they trust how data will be used, see mutual benefit, and won't be punished for transparency.

What Makes Our Approach Novel

No widely-used system today offers **real-time, adaptive privacy controls** for outbreak information sharing. Current systems force binary choices—share vague summaries that lose critical context, or wait to share everything (often too late). Our Live Coordination Infrastructure enables graduated, context-preserving sharing where labs control what to reveal at each stage, with AI facilitating just-in-time structure generation from free-form inputs.

The Structure of Our Solution

Our Live Coordination Infrastructure addresses each central constraint through specific technical mechanisms. The system operates through three primary artefacts: Abnormality Bubbles, Distributed Matching, and Negotiated Communication Channels. Each component preserves the qualities we identified as essential while avoiding the failure modes of existing approaches. Look at [Appendix A](#) for workflow details.

Abnormality Bubbles: Postformal Artifacts for Distributed Meaning

Addressing: Authority Vacuum & Contextual Dependence

At the core of our infrastructure are *Abnormality Bubbles*—dynamic artifacts that represent local observations without requiring consensus on a shared representational schema. Each bubble originates from free-form input: a researcher's natural description of observations ("three patients with unusual respiratory patterns, all from the same flight..."), and our system extracts structured fields while preserving the original informal context.

This design addresses the **Authority Vacuum** by eliminating the need for standardized reporting formats. A hospital in India can describe symptoms using local diagnostic language, while a lab in Brazil uses genomic terminology—our infrastructure handles the translation just-in-time rather than forcing prior agreement.

The preservation of **Contextual Dependence** occurs through our privacy-adaptive architecture. When creating a bubble, the user selects one of three privacy levels that determines what information will be visible to others:

- **High Privacy**: Essential information only (symptoms, general location, timeline)
- **Medium Privacy**: Anonymized details with more specificity

- **Low Privacy:** Full contextual details including sample counts and genetic markers

Crucially, these levels are **calibrated choices** that the reporting party makes based on their assessment of the situation. The same underlying observation—captured in the original free-form input—can be shared at different levels of detail depending on the urgency, sensitivity, and potential audience. This allows the reporting organization to **preserve meaningful context while respecting institutional constraints**, rather than being forced into a binary "share everything" or "share nothing" decision.

AI-Powered Distributed Matching: Surfacing Inferential Cascades

Addressing: Authority Vacuum & Informational Pragmatism

The matching infrastructure solves a critical problem: how do you find relevant threats across organizational boundaries when you don't know what to look for, and there's no central clearinghouse that can see all data?

Our system uses AI to perform **semantic analysis** across multiple dimensions:

- Symptom similarity and temporal correlation
- Geographic proximity and travel patterns
- Detection method compatibility (clinical vs. genomic vs. environmental)
- Urgency alignment (both labs seeing escalation)

When a lab creates a bubble, the system suggests matches with **AI-generated explanations**: "Both locations show respiratory symptoms with rapid progression in adults 30-50; Singapore airport screening + Brazil clinical surveillance suggests possible travel vector."

This addresses **Informational Pragmatism** by making the relevance judgment computational and adaptive rather than requiring humans to manually triage thousands of reports. The AI doesn't decide what's *true*—it decides what's *potentially related*, preserving human agency for the critical verification step.

Importantly, the system also surfaces **counterfactual visibility**: labs can see not just recommended matches, but also browse the full "library" of all reported abnormalities; preventing false negatives from over-reliance on algorithmic matching while still providing the benefits of intelligent prioritization.

Communication Channels: Just-In-Time Negotiation Protocols

Addressing: Informational Pragmatism & Bounded Temporality

Once a lab identifies a potentially relevant match, they can open a communication channel. Rather than forcing users into either "share everything" or "share nothing," we provide three message types that enable **progressive unfolding** of information:

1. **Request Information:** "Could you share age demographics of affected patients?"
2. **Send Information:** Proactive sharing when urgency is clear
3. **Conditional Flow:** "I'll share our genetic sequencing data if you share travel history"

This structure addresses **Bounded Temporality** by enabling coordination at multiple speeds simultaneously. For rapidly unfolding threats, labs can share proactively. For uncertain signals, they can request clarification before committing to detailed disclosure. The conditional flows allow game-theoretic negotiation without lawyers—"I see value in coordination, but need reciprocity before revealing sensitive details."

The **Informational Pragmatism** constraint is met through this graduated disclosure. Privacy isn't binary; it's contextual and negotiated in real-time based on:

- What the other party has shared (reciprocity)
- How urgent the threat appears (risk-benefit calculation)
- How much trust has been established (history of interaction)
- What specific information would be mutually valuable (targeted exchange)

Critically, these negotiations happen in **minutes**, not weeks. The infrastructure does the work of generating sanitized versions, tracking reciprocity, and suggesting relevant information to share—freeing humans to focus on the semantic judgment of "does this connection matter?" rather than the bureaucratic work of "how do I legally share this?"

Conclusion

Our solution rejects three common failure modes in coordination infrastructure:

1. **Centralized Authority** (e.g., "report everything to the WHO and wait for their analysis")
 - Our system: Pairwise matching, distributed verification, no bottleneck
2. **Rigid Protocols** (e.g., "fill out this standardized form with these exact fields")
 - Our system: Natural language input, AI-extracted structure, postformal artifacts that adapt to context
3. **Premature Sanitization** (e.g., "strip all identifying details before sharing")
 - Our system: Context-preserving translation, privacy levels negotiated just-in-time, meaning preserved across trust boundaries

The key innovation is that **structure is generated when needed, not imposed in advance**. When a lab in India wants to coordinate with a lab in Brazil, the system creates a just-in-time "handshake"—a communication protocol tailored to their specific bubbles, trust level, and urgency. There's no universal schema they must both implement; the AI infrastructure does the translation work, allowing them to meet each other in their own language.

Appendices

Appendix A: System Workflow

This section outlines the step-by-step workflow for the "Live Coordination Infrastructure", illustrating how a local observation is transformed into a global signal.

1. Access and Dashboarding

The workflow begins with institutional authentication. Users log in to their specific agency node (e.g., National Institute of Virology) to view the internal dashboard.




Figure 1: Users select their agency node to ensure data sovereignty.

This dashboard tracks active "Abnormality Bubbles"—ongoing investigations shared previously by the agency.




Figure 2: The dashboard displays current alerts (e.g., "Potential rabies outbreak").

2. Reporting an Anomaly

When a researcher observes something unusual, they input the data using natural language. The system acts as a structuring tool, converting rough notes into a standardised format. This can accept structured, unstructured, or even voice notes, as well as other forms of information in the future.

- **Step:** The user types a description (e.g., "rapid progression," "negative tests").

- **Result:** The system extracts key fields (Location, Timeline, Findings) and lets the user set a privacy level (e.g., "Low (full details)") before saving.

← Back to Dashboard

Create Abnormality Bubble

Describe the abnormality in your own words. AI will help structure it for sharing.

Privacy Level: **Low (full details)**
Full details will be shared

Describe the abnormality or situation 🗺️ Roll Random Situation

We've got a situation at the main hospital. Three patients came in with what looks like atypical pneumonia but none of the usual tests are coming back positive. They're all from the same neighborhood. We're seeing rapid progression in two of them. The lab is running more tests but we're worried this might spread. We've isolated them but need to figure out what we're dealing with.

You might want to consider:

- Consider adding specific details about patient demographics and potential shared exposures in the neighborhood.
- Specify the exact medical tests that have been negative and what additional tests are currently being conducted.
- Include more precise information about the symptoms and clinical progression for each patient.
- Document any environmental or contextual factors that might be contributing to the unusual disease pattern.

Create Abnormality Bubble Cancel

Preview: Post-Formal Artifact HIGH

Click any section to edit. This is how recipients will see your abnormality bubble.

National Institute of Virology (NIV) Low Privacy

Pune, India
Clinical surveillance
+ Add clarification

SUMMARY
Atypical pneumonia with unusual clinical presentation
+ Add clarification

KEY INFORMATION

Location: Localized neighborhood
+ Add clarification

Timeline: Current/frequent
+ Add clarification

DETAILED FINDINGS
Rapid disease progression, negative standard tests
+ Add clarification

TECHNICAL DETAILS
Sample Count: 3
+ Add clarification

Privacy Level: Low Privacy
Full technical details are shared with recipients

Figure 3: The creation interface. The left panel captures raw text; the right panel shows the structured artefact ready for the network.

3. Contextual Analysis

Once the report is saved to the "Analysis Canvas," the system scans for patterns across other networked agencies.

- **Step:** The user reviews the canvas and selects the abnormalities that they feel are relevant to their case, both from the recommended AI suggestions and also from the other reported abnormalities from the library.
- **Result:** The system flags a "High" urgency from the National Centre for Infectious Diseases (Singapore), citing "Anomalous respiratory patterns" in airport screening data as relevant to this abnormality.




Figure 4: The Analysis Canvas highlights a suggested link between the local clinical case and international screening data.

4. Correlation and Mapping

Accepting a suggestion reveals the relationship between the data points.

- **Step:** The user accepts the match to visualise the connection and look at the reasoning for this connection.
- **Result:** A connection line links the NIV case to a UK Health Security Agency (UKHSA) report. A text block explains the reasoning for this connection: both locations are seeing respiratory issues with high urgency, despite different detection methods (clinical vs. genomic).




Figure 5: A node graph showing the confirmed link between the Indian and UK datasets, with a summary of the overlap.

5. Verification Channel

To confirm if the events are truly related without exposing sensitive patient data publicly, the user opens a direct line to the counterpart agency.

- **Step:** The user initiates a chat with UKHSA contextually tied to the bubble.
- **Result:** The user sends a specific request ("could you share information about the age group...") to verify the signal and start a communication channel between the groups that can be taken offline and continued.




Figure 6: The secure communication window used to validate the correlation.

Appendix B: Interpreting Constraints for a Live Infrastructure Approach

The design of our platform has been formed by direct reference to the central constraints we identified.

1. Authority Vacuum

There are multiple counterintuitive affordances we aimed to preserve in these spaces lacking authority for standardization and comparison. Under such regimes, we believe a pluralism of values is maintainable through the use of [postformal artefacts](#). This allows for new [logistics for distributing meaning](#) that does not stipulate a reduction of information for the sake of generalizing its application domain. Generalizing by inclusion (instead of exclusion) allows for collective sensitivity to localized problems.

To bypass this vacuum, we create exchangeable formats called Abnormality Bubbles, which intelligently combine to recommend potential threat vectors for human consideration. These represent the *appropriately structured* transformations of the informal observations of the

detection processes through the rigorous bureaucratic structures already mentioned.

2. Contextual Dependence

The rich context that is typically lost in crucial exchanges requiring privacy redactions is similarly a fault of generalizing by exclusion. This means that there is a particular usefulness to being able to conserve the informal aspects of information presented in optimally fluid user interactions. As such, the best solution would be one that can take in unstructured inputs and restructure outputs according to rigorous constraints dependent on the sender's relations with the receiver.

To address this, trustworthiness as a dynamic property can be reified with continual ratings between labs or agencies. Moreover, we can normalize these agencies' internal reporting with respect to the regional average baselines in order to more sensitively detect abnormalities.

3. Informational Pragmatism

Navigating the fine balance of privacy requirements and realizing information relevance requires careful progressive unfolding to combat the uncertainty of what information should be shared. This means that human centrality cannot be completely automated away, despite there being many incentives to computationally enhance this process.

Our infrastructure aims to be sensitive to the ability of the human doctors, diagnosticians or other personnel to realize what is relevantly abnormal about a situation. In so doing, it emphasizes their ability to make those observations which may save lives.

4. Bounded Temporality

Somewhat in contrast to the above, the incentives toward speeding up all bureaucratic processes in all domains not requiring human oversight will likely result in the (at least partial) automation of otherwise slow negotiation processes. This must occur in accordance with established organizational policies and incentive structures, but while respecting the nuanced observations unique to human agency.

For representing this in our prototype, we surface both presumed-to-be relevant Abnormality Bubbles and counterfactually display other Bubbles that were considered less relevant, in order to provide the users with a more transparent reasoning workflow.

Appendix C: Future Directions: Comparable Domains for Similar Solutions

Aside from healthcare, there are a few other notable situations that might benefit from similar infrastructural solutions. In particular, we claim that coordinating for problems such as financial crime, supply chain security, and AI safety coordination represent analogous domains across scales of varying complexity.

Otherwise, we would also intend to demonstrate the negotiation process in more detail aside from only representing this as a Communication Channel between affected labs. Moreover, correctly rewarding the act of contributing Bubbles for others to consider may be necessary in cases where the lack of heroes in prevention does not offer sufficient motivation.