UNSECURITIES LAB
Participant Report

Dr Nathan Jones

Security Lancaster, Lancaster University

Session: 31st March @ Goldsmiths, University of London

Chatham House Rules Apply — Restricted Distribution — Not for Attribution

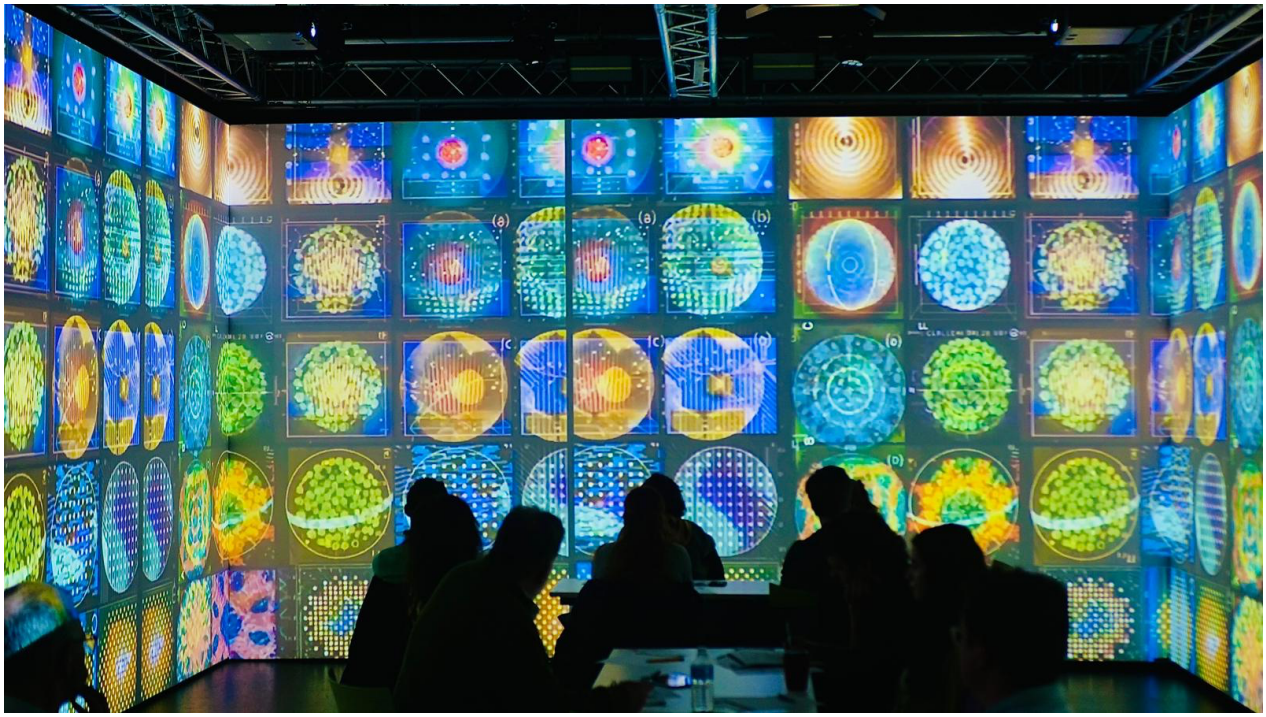

Executive Summary of Report

A unique workshop where government practitioners and academics used immersive art to simulate decision-making under high levels of uncertainty. By observing a film featuring AI-generated biology, participants explored how diverse professional backgrounds—from evolutionary biology to crisis response—interpret ambiguous threats differently.

- The exercise revealed that instinctive emotional reactions serve as vital analytical signals and that rushing toward a forced consensus can actually increase operational risk.
- Participants also studied non-human intelligences, to model more resilient systems capable of surviving the volatile environments of the Anthropocene.
- Ultimately, the findings suggest that cross-disciplinary collaboration and slowing down interpretation are essential for managing complex technological and environmental hazards.
- This experimental approach aims to strengthen security thinking by addressing the vulnerabilities inherent in AI-mediated evidence and shifting timescales.

The workshop is most relevant to policy and security teams working with uncertain, AI-mediated or cross-domain evidence, where premature classification, weak signal loss, or over-fast consensus can create operational risk.

The clearest entry point to this report is the **Interaction Map** in Output 1. It shows how the workshop moved from three different table readings of the same film to a shared diagnosis of the problem. Participants first responded through unease, alarm, empathy, anger and suspicion. These reactions were then tested through group discussion and developed into common observations about intervention, misreading, timescale, self-reinforcing evidence and classification failure.

The main finding is that the workshop showed how **groups can begin to handle uncertain evidence better** by collaborating across disciplines. In Session 1, participants slowed down, questioned the material, compared different professional perspectives, and avoided settling too quickly on one explanation. These are practical skills for working with ambiguous evidence.

In Session 2, the group used speculative models of resilience to **imagine what systems will respond better to volatile and uncertain conditions**. The discussion produced four useful capacities: remembering what has happened over long periods, coordinating across different perspectives, recognising when conditions have changed, and adapting when existing ways of operating no longer work.

The transcript analysis gives the evidence for these findings. It shows how professional background shaped what participants noticed first; how each table built its account; where people changed their minds; how instinctive reactions became useful analytical signals; and which observations emerged independently across all three tables.

Output 2 translates the workshop findings into five stress points for security thinking: unreliable evidence, mismatch between intervention and system timescales, systems feeding on their own outputs, the cost of forcing consensus, and the risk that what cannot be named cannot be protected.

The final section explains how **Security Lancaster can develop this approach further with policy, security and research partners, including bespoke workshops on issues with complex environmental, agential AI and cyber-physical overlaps.**

Contents

Executive Summary of Report 2

Contents 3

Introduction 4

Summary of Findings 5

Output 1: Interaction Map 6

Transcript Analyses 7

Analysis 1: Disciplinary Signature Analysis 7

Analysis 2: Table / Collective Analysis 9

Analysis 3: View Change Tracking 12

Analysis 4: Left / Right Hemisphere Vocabulary Analysis 14

Analysis 5: Instinct-to-Articulation Tracking 16

Analysis 6: Cross-Table Thematic Emergence 18

Output 2: Stress Points in Escalation and Coordination Logic 20

How Security Lancaster Can Help 22

Introduction

This report presents findings from the Unsecurities Lab workshop held at Goldsmiths, University of London, which brought together government practitioners and academics to test a new format for security thinking under conditions of ambiguity.

Unsecurities Lab is a format developed at Security Lancaster. It takes its name from forms of instability, exposure and interdependence that challenge conventional models of threat, risk or response. It uses immersive art to stage situations in which environmental, technological, biological and informational processes are entangled and difficult to stabilise.

In this episode, Joey Holder's deep-sea, deep-time works provided the provocation: dense, disorienting environments in which distinctions between lifeform, machine, ecology and infrastructure begin to blur, and in which intelligence and resilience must be thought beyond familiar human-centred models. The method is designed to test how expertise operates when familiar frames are under pressure: how participants orient themselves, interpret uncertain signals, compare perspectives and build shared accounts of what is happening.

For academics, this offered a way to test how disciplinary knowledge travels when placed alongside operational and policy perspectives. For government practitioners, it created a structured space in which to strengthen cross-domain judgement, test assumptions, and explore more flexible responses to emerging and interacting threats. The report that follows analyses session transcripts in order to identify the main claims and practical recommendations arising from the exercise.

Session 1 was structured around Joey Holder's *Ambiogenesis*: a moving-image artwork that brings together AI-generated creature imagery, biological visual culture, documentary underwater footage, and an unsettling sound design. We asked participants to watch *Ambiogenesis* as if it were a record of an unknown incident and to work together across their different disciplines to decide where the threat or risk lay.

Session 2 introduced four speculative creature-intelligences and presented a post-incident world in which these organisms' survival logics had become the basis for new forms of intelligence. Session 2 asked what kind of resilient system could emerge from each creature's distinctive way of being, and what that revealed about the strengths, limits, and trade-offs of different models of intelligence.

Summary of Findings

Three core claims run through the findings.

A: affective response is analytically productive. Emotional reactions to the material — unease, anger, disorientation — were recognised by participants, and regularly related to dimensions later used in analysis. Immersive experience generates the raw material; group discussion develops it into operational knowledge. (see Analysis 5: Instinct-to-Articulation Tracking)

B. disciplinary difference is a diagnostic resource. Each discipline performed a different epistemological operation on the same material, and the workshop's value lies in the encounter between those positions. (see Analysis 1 (Disciplinary Signature Analysis) and Analysis 6 (Cross-Table Thematic Emergence)).

C. a core vulnerability in complex systems is recursive self-misunderstanding: a system that intervenes at the wrong speed, treats the consequences of its own actions as fresh data, and lacks the vocabulary to name what it is affecting. (see Output 2, Analysis 2 (Table/Collective Analysis) and Analysis 6).

Four competencies for working with ambiguous evidence are mapped in Output 1:

- **comparing** partial views across disciplines before settling on a reading;
- **noticing** the frameworks through which an incident is being made intelligible;
- **slowing** interpretation under pressure;
- **questioning** the evidence itself, including its origin and sequencing.

Analysis 3 and Analysis 4 show how these competencies developed across the session.

Five stress points for security and policy are set out in full in Output 2.

1. The premise of evidence cannot be assumed reliable in genAI conditions, when its origin, sequencing, or interpretive framework may itself be part of the situation being assessed.
2. Timescale mismatch in complex ecosystems means interventions can have consequences accumulating across decades or cascading within hours.
3. Recursive self-sampling — a system trained on or used to validate its own prior outputs — results in the loss of analytical boundaries, and corruption of data in ways that echo biological corruption. Analysis 6.
4. Consensus has a security cost, suppressing dissenting signals at the moment of alignment. A mechanism examined in Analysis 2 and Analysis 3.
5. The most consequential features of a complex system are often those for which existing taxonomies have no place.

Participants imagined types of intelligent system that are resilient to those 'stress points' with four useful capacities: remembering what has happened over long periods, coordinating across different perspectives, recognising when conditions have changed, and adapting when existing ways of operating no longer work. Brought together they imply a resilient intelligence system design. They are summarised in Output 1 and developed in Analysis 2.

The disciplinary complementarities demonstrated here will deepen with practice. This session produced the foundations — shared vocabulary, tolerance for non-convergent readings, instinct for live diagnostics — that a programme of repeat workshops is designed to consolidate. We describe how to build on this work in the final page.

Output 1: Interaction Map

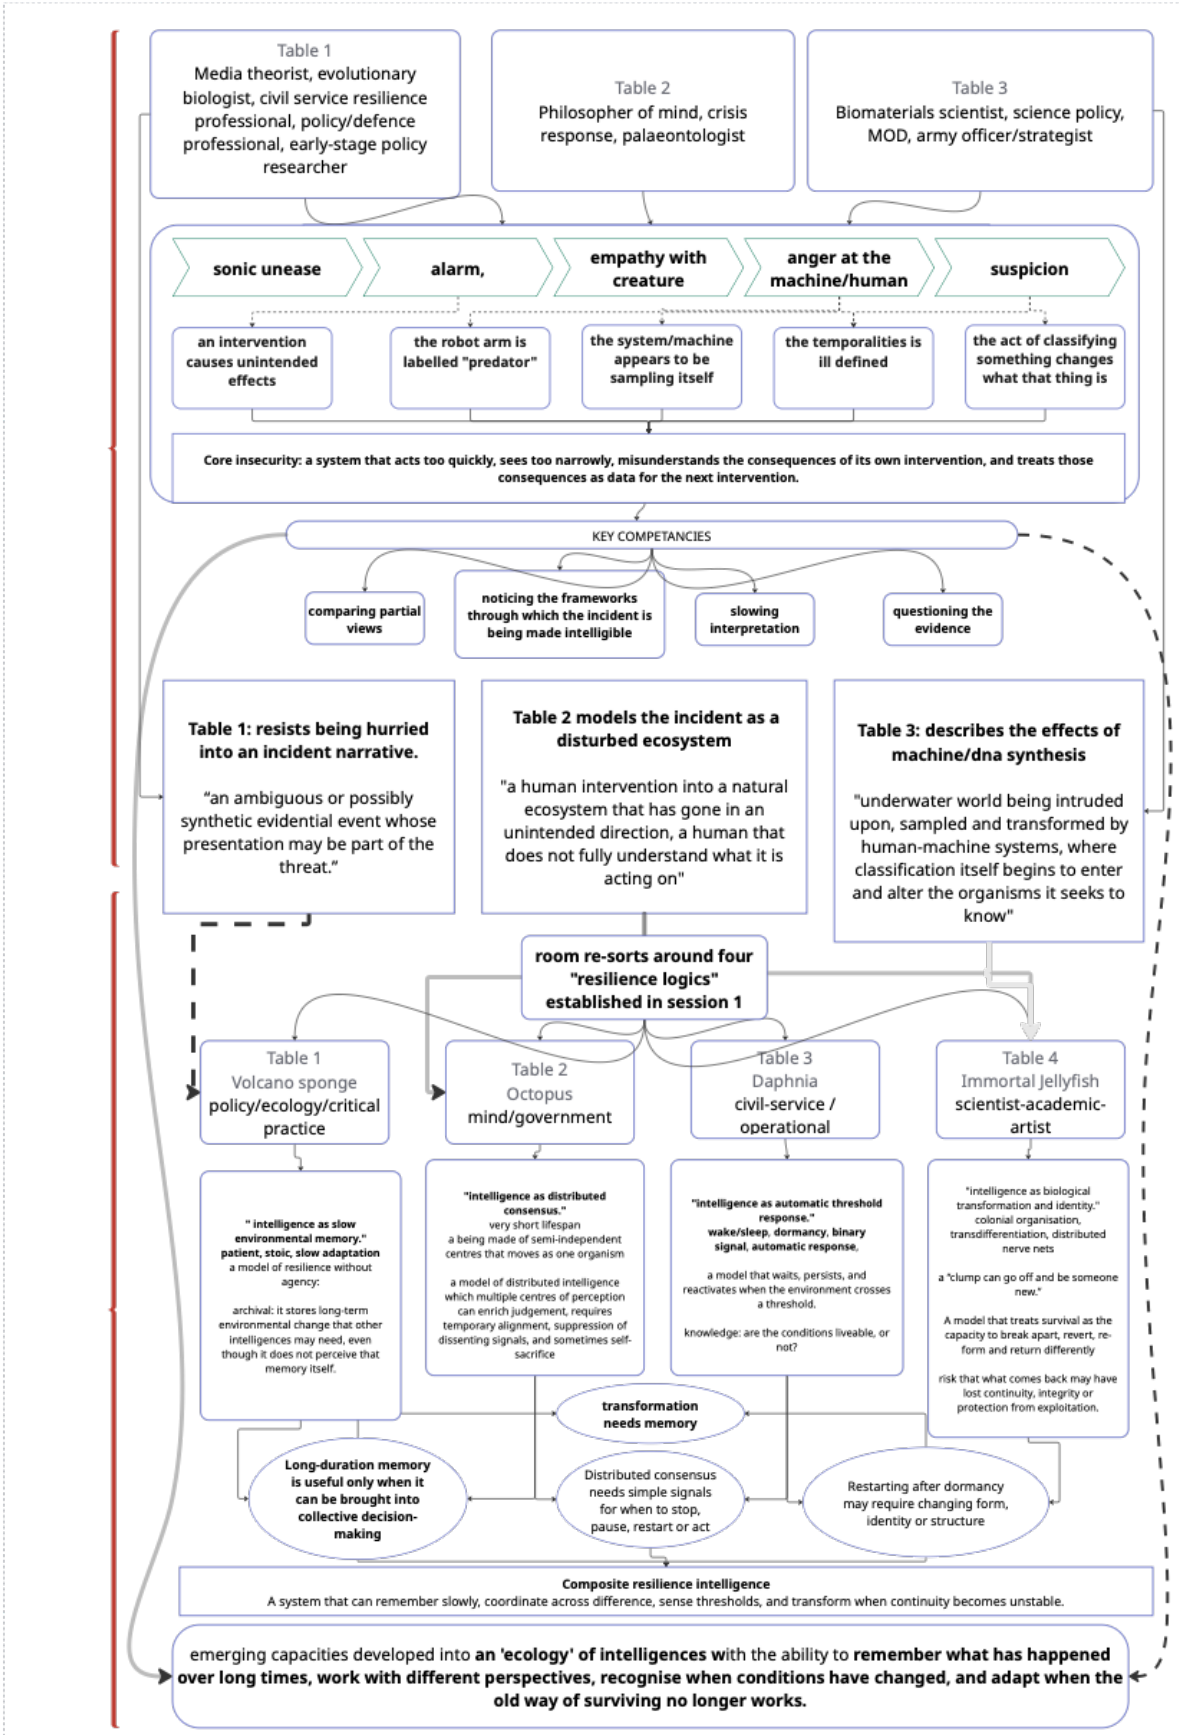

Transcript Analyses

Analysis 1: Disciplinary Signature Analysis

Below is a series of indicative profiles/persona types, describing the diverse instincts and vocabularies of different participants in the workshop.

The disciplinary signatures cluster into four distinct types.

1. Maker enquiries: How was this produced? What tools? What training data?
2. Credibility enquiries: Is this real? Does it conform to what we know?
3. Surveillance enquiries: Who is watching whom? What is the threat relationship?
4. Epistemological enquiries: What kind of thing is this? How and what can we know?

We can see from this analysis that disciplines are performing different epistemological operations on the same material. The integration the workshop produces – described in the analyses following – is the result of an encounter between them. Below we also note the initial responses participants make to ambiguous and 'un-interpretable' stimuli in part 1 of the workshop.

Speaker 1 — Media theorist / educator

Initial risk vocabulary: What's at stake in the technology we're using. Practice as a mode of engaging with knowledge. First instincts: Immediately identifies the image generation model as older than presently available ones, reads the aesthetic quality of the AI output for evidence of its training data, notices the classification systems absorbed into the biological forms. When the image resists naming, reaches for the meta-level: the mode of production as content.

Characteristic vocabulary: *cross-genre, mixtures of types, mode of presentation, how it's structured.*

Speaker 2 — Evolutionary biologist

Initial risk vocabulary: How organisms get trapped in situations on an ecological timescale. Resilience is actually interestingly opposed to strong and committed cooperation. First instincts: Immediately applies morphological credibility testing — looping arms are wrong, symmetry is suspicious, these are things biology probably wouldn't have made. When the image resists categorisation, asks: would I believe this if it were presented as a real observation?

Characteristic vocabulary: *developmental, looping, asymmetrical, evolution, macroevolutionary, resilience.*

Speaker 3 — Civil service resilience professional

Initial risk vocabulary: Risk feels incredibly real, not particularly grippable. We see a lot all the time with very limited levers. First instincts: Notices the natural/unnatural split, holds both sides simultaneously, looks for the pattern across the whole sequence rather than fixing on individual images. When the image resists categorisation, translates it toward familiarity — finds the recognisable element and anchors to it.

Characteristic vocabulary: *holistically, intersections, pattern of familiarity, digestible, translating*.

Speaker 4 — Policy professional / defence

Initial risk vocabulary: Quantifying risk to try and make something fluid into something that can be mapped on a chart. First instincts: Reverse-engineers the prompts that generated the AI sequences. Sees art in terms of symbols, for example notices the sensory apparatus — eyes, ears — and reads them as symbols of surveillance. When the image resists categorisation, reaches for the sequencing: was this deliberate, is the order meaningful, are we being manipulated?

Characteristic vocabulary: *opportunity space, action-oriented, breaking it down into things you can do things about, packaging advice.*

Speaker 5 — Early-stage policy researcher

First instincts: Notes that the mechanical sounds are recognisable rather than alien — sounds like an old train. When the image resists categorisation, acknowledges confusion directly rather than resolving it, leans into the ambiguity as the content.

Characteristic vocabulary: *inherent risk in a strategy, resilient idea, ambiguity, unclassifiable*.

Speaker 6 — Biomaterials scientist

Initial risk vocabulary: Minimising any risk to the patient. Might we produce something which is sentient? There's quite a lot of ethical risk there. First instincts: Identifies the brain organoid on the multi-electrode array precisely and correctly. Reads the developmental sequences as something close to his own laboratory practice. When the image resists categorisation, contextualises within the ethics of synthetic biology — the question of whether you might accidentally produce sentience.

Characteristic vocabulary: *organs develop, screening drugs, tissue engineering, pre-clinical testing, ethical risks.*

Speaker 7 — Science policy

Initial risk vocabulary: ideas of contagion. First instincts: Reads the intrusion as potentially sampling for biological material for later use. Notes the creation/growth/destruction sequence. When the image resists categorisation, asks about intent: curiosity or extraction?

Characteristic vocabulary: *engineering and biology, advanced materials, balance things everyone is slightly unhappy with.*

Speaker 8 — Officer / Strategist

Initial risk vocabulary: A risk is something that hasn't happened, but might. How do we stop it happening? First instincts: Reads the visual field immediately in terms of observation and surveillance — being watched. Constructs a narrative: organisms, their visual experience, their lifecycle, then the alarm, then the mechanical intrusion. Reaches for the practical risk created when attempts to understand a system disturb or damage that system.

Characteristic vocabulary: *risk to the country, prevent risks occurring, control the likelihood.*

Speaker 9 — Psychologist

First instincts: Proposes the film might be structured in reverse chronology — what we see first is the end state. Interested in the lens rather than the content: are we seeing reality or seeing our apparatus of seeing projected onto it? In conversation, holds multiple framings simultaneously without resolving them.

Characteristic vocabulary: *opaque to understanding, ambiguity, holding tension, consciousness, what does it mean to have a mind*.

Speaker 10 — Palaeontologist

First instincts: Sees cephalopod traits, brachiopod traits, vertebrate features mixed together. Immediately thinks of genome tampering. When in doubt, reaches for scientific concepts: deep time, abiogenesis, ontogeny recapitulating phylogeny, evolutionary misfire.

Characteristic vocabulary: *deep time, mass extinctions, biodiversity, unintended consequences, single global ecosystem.*

Speaker 11 — Crisis response

First instincts: Notes that the understanding of the image keeps breaking, that the stream of comprehension is being interrupted. Interested in the incompleteness of the knowledge rather than its content. When the image resists categorisation, notices that the familiar elements keep twisting into something unbelievable.

Characteristic vocabulary: *government sees, diplomatic network, threats to UK interests, the psychology of.*

Analysis 2: Table / Collective Analysis

Table 1, Session 1

Table 1's defining feature was its resistance to being hurried into a story. It explores themes of manipulation, framing, and the possible untrustworthiness of the visual and audio material presented. They keep asking whether the imagery is real, whether they are being steered by sequencing, and whether narrative closure is itself a trap. The incident definition is: **"an ambiguous or possibly synthetic evidential event whose presentation may be part of the threat."** The table's core concern is how we are induced to believe that something has happened.

Its main strength was premise interrogation. It was the clearest table in treating framing, sequence, and manipulation as part of the problem itself. What it achieved was a form of epistemic caution that is highly relevant to AI-mediated evidence: it asked *"what is this evidence doing to us?"* A further strength was its tolerance for discomfort and friction early on, which meant it was less vulnerable to premature closure than some other tables.

"We then spent a while trying to figure out how we were being manipulated through the juxtaposition of the different images and the sequencing of that, whether that was a deliberate order to make us feel a particular thing, what were the linkages between the different image sets, were we trying to ascribe a narrative, and then loading it in with all our own kind of interpretive biases or not."

Table 2, Session 1

Table 2 moved more quickly toward incident architecture. It defines the incident as a human intervention into a natural ecosystem that has gone in an unintended direction, with a specifically human interpretive layer that does not fully understand what it is acting on. They agree on a potentially catastrophic cycle of intervention, incomplete understanding, and repeated correction, the results of which are unfolding beyond the intervener's grasp. The threat identified is human intrusion into a complex biological system under conditions of partial comprehension.

"We had this sense that there was an existing biology, under the sea, and some process of sampling it. There was a sense of sampling by perhaps some kind of robotic observation system for the purpose of manipulation of biology — that this perhaps went wrong."

Table 3, Session 1

Defines the incident as an underwater world being sampled or intruded upon, though they disagree on whether this is curiosity-driven research or extractive intervention. They reach a rough consensus that they are seeing an underwater environment disrupted by human or machine contact, with later images suggesting manufacture or transformation. The threat is framed as intrusion into a poorly understood system, with disagreement over whether the motive is study or exploitation.

The strongest idea here is that the organism appears to become part classification. The apparatus of knowledge does not remain outside the thing known; it enters the form itself. This makes the table especially important for thinking about AI-mediated evidence, because the problem is no longer only whether a system misclassifies the world, but whether it begins to produce beings made of its own classificatory habits.

"You've fed into the AI loads of images of like biological things, whatever it is, but some of those images have the classification systems within them. And so then what you get as a result is a weird kind of hybrid where it's not biological, but it is a classification system, but it's also taken on some of the features of the organic things that it's classifying."

--

In Session 2, the group make decisions about which table to join, based on their reading of the modes of intelligence depicted in the film

--

Table 1, session 2

Adopts the **volcano sponge** as its creature. Their discussion describes an intelligence that may barely "act" in recognisably human terms, which fits the sponge's rooted, long-duration, low-agency logic. The resilience model described is patience, environmental embeddedness, and survival through duration — and therefore vulnerability to rapid change.

"Its survival logic is that it has no intelligence and therefore is not a threat to anything around it. It just exists. So why would you bother posing a threat to it?"

The table's strength lies in its engagement with the idea of structural memory, how change and threat relate to metabolic speed, and the idea of archives as deep infrastructure.

Table 2, session 2

This table chooses the **octopus**, immediately orienting around the "nine brains" thought experiment. It defines its problem as one of distributed intelligence, coordination, and decision-making without a single unquestioned centre. Its resilience logic becomes visual, tactile, and mood-based coordination across a decentralised organism.

"Your camouflage requires your entire body to agree on the signal. What is the security cost of the moment of consensus?"

Its core strength was its self-reflexive way of addressing both challenges. The table realised that its own conversation was enacting the same coordination problem it was trying to analyse. What it achieved was a practical insight into distributed intelligence, especially the difficulty and cost of consensus when agency is decentralised. The most important development here is the recognition that consensus has a security cost. Action requires consolidation, but consolidation also shears off weak signals.

Table 3, session 2

Two civil servants gravitate toward daphnia. Adopts the water flea / daphnia, and its model is explicit: no speech, little agency, mostly chemical response, dormancy, and waking when conditions change. They define resilience as automatic persistence through catastrophic timescale shifts, not active perception or strategic deliberation. The table's vulnerability is that such an organism cannot really "perceive" threat, only cycle through environmental triggers.

They define the water flea almost as an **anti-deliberative intelligence**: chemical sensing, no speech, no vocal communication, little or no free will, and almost no decision-making. One participant says its communication is "through the chemical," another says it communicates the state of the environment by whether it is awake or not; they then describe it as "a sensor or a switch," with "no speaking, listening, or deciding" because it is "automatic."

Table 4, session 2

At table 4 scientists and academics self-sort toward the immortal jellyfish because it "seems quirky" They think through existing jargon (shared in these disciplines?) transdifferentiation, distributed sensation, dispersal and reconvening. This narrows the operational diversity of the table, but strengthens the conceptual quality of the conversation by producing a precise problem of survival without stable identity.

Chooses the **immortal jellyfish**. They define its intelligence through colonial organisation, transdifferentiation, distributed sensation, dispersal, and reconvening. Its resilience model is therefore splitting apart under pressure, surviving in fragments, and potentially reassembling later, though with risks around identity and exploitation when reconvening.

The table's strongest contribution is on the issue of collective and fragmented identity. The table explicitly worries about whether what reconvenes is really "you", and whether reconvening creates exposure to exploitation. That gives the discussion analytical grasp on the trade-off between survival and integrity.

"You have reverted forty times, each you is genetically identical but experientially distinct. Where is continuity located? Where is vulnerability?"

Initial quantitative and lexical indicators across tables

Lexical shift: Part 1 is dominated by disciplinary-risk language: policy, evidence, control, strategy, and uncertainty. Part 2 is dominated by creature/resilience language: creature, organism, environment, intelligence, memory, distributed, dormant, communication. The shift is large enough to be visible even in basic keyword proxies.

Semantic shift: In the opening discussion, risk is framed largely as something to prevent, quantify, manage, or mitigate. By Session Two, it is more often understood as disturbance within a system: a mismatch of scales, logics, or interventions, often producing recursive effects.

Shift in confidence: Participants begin by speaking confidently from within their own discipline, then become less certain when the material disrupts those frameworks, then start to recognise shared structural patterns across different readings, and finally move toward a more plural idea of resilience, where survival depends on multiple different strengths rather than one dominant model.

Analysis 3: View Change Tracking

This section records moments where participants changed their reading, added another discipline's perspective to their own, or recognised that the group's decision-making process was itself part of the exercise. These shifts matter because they show the workshop turning ambiguity into shared judgement rather than simply producing a single answer.

1. The predator/prey inversion

Starting position: initial alarm at the creatures, reading them as threatening. Trigger: group discussion over approximately ten minutes.

> "My first response was, if I was swimming around and I saw that, I'd be like — just run away, basically. And yeah, as we've rationalised it together, more and more... the threat relationship is the other way around. We're the threat."

Mechanism: collaborative rationalisation — the group worked out together what the images implied, and the working-out changed the reading. [Gestalt shift]

2. The scientist notices manufacture

Starting position: the biomaterials scientist reads the intrusion as purely curiosity-driven, not extractive. Trigger: the army officer's narrative summary, which included the endpoint of synthetic manufacture.

> "I really like the last bit which you said, how it then comes to manufacture — our attempts to manufacture life almost. I think that's a link which I hadn't made. I now see that in that last set of images."

Mechanism: disciplinary complement — one discipline surfaced something in the material that another could not see from its own position. The scientist does not abandon his curiosity reading; he adds the manufacture endpoint to it. [Integrative thinking]

3. The order reversal

Starting position: one participant reads the film in forward order. Another:

> "See, I saw it the other way around. I saw the third one being the destruction, and then this one, back to first — so this one reeling from that destruction."

A third participant: "Actually I like that…" and develops the reversed reading further. Mechanism: temporal reframing — an alternative ordering opened up a richer reading that others could inhabit. [Synthesis]

4. The watching inversion

Starting position (biologist): "These things are not watching us. They are being watched, despite how many eyes they have." Counter (policy professional): "I disagree. Certainly right at the start, it felt like we were under observation." The biologist does not fully concede but shifts:

> "I mean, I feel like why would we be deep under the sea? If we are deep under the sea, it's because we put sensors down there to watch them. And maybe they're watching our sensors."

Mechanism: productive non-resolution — neither party changes their fundamental position, but the biologist incorporates the counter-reading into his own framework, producing a more complex account. [Holding opposing view in balance]

5. The communist sponge

Starting position: one participant characterises the volcano sponge's survival logic as "communist." Others resist. The original participant concedes:

> "Still very communist... I haven't written down communist though. Realpolitik. It's always taking a census of its environment, and it's happy to just be the census taker."

Mechanism: challenge producing precision — resistance to the first framing forced it toward a more accurate articulation. The original insight was correct; the word was wrong.

6. The octopus group recognising themselves

Starting position: extended struggle to understand how nine semi-independent brains constitute a single decision.

> "What do the four of us gain by sitting around this table and talking, trying to come to a consensus, that each of us thinking individually wouldn't gain? We are four brains sat around a table trying to arrive at some kind of conclusion."

Mechanism: self-reference — the group recognised that the problem it was wrestling with abstractly was the problem it was actually enacting. What they discovered through it was specific and operational: that the moment of consensus in any distributed system is simultaneously the moment of maximum vulnerability. The octopus's camouflage requires all nine brains to agree on the signal, and during that alignment it is most visible to predators. They named the government parallel directly: rapid cross-departmental agreement produces diluted decisions and suppresses the dissenting signals

that may be most important. The short lifespan of the octopus, they proposed, is not incidental to its intelligence but a consequence of it.

7. The flea as memory

The facilitator proposes a device neither group had developed:

> "Maybe there's a device there that you can agree on somehow where you sort of like can give a bit of simple information to the water flea to store it so that when you both vanish in the situation when you come back, one can create itself anew dependent on what's the best thing for that situation."

Immediately taken up: "The retaining of information might also be this kind of genetic, biological thing, rather than something that we just remember in our brains." Mechanism: facilitated synthesis — the facilitator identified a complementarity between two organisms that the tables had not reached independently.

View changes cluster into three types:

1. Gestalt shifts (predator/prey, watching) — the same material is suddenly seen as a different whole. The rarest and most significant change.
2. Additive integrations (manufacture, communist→realpolitik) — a new element is added to an existing reading without replacing it. Most of the change in this workshop is of this type.
3. Self-recognition (octopus group) — the abstract problem is recognised as the actual situation. This happened once, and it was the most direct demonstration of the methodology's claim that embodying non-human intelligence produces insight about human intelligence.

Analysis 4: Left / Right Hemisphere Vocabulary Analysis

Coding the transcripts for two modes of cognition, tracking the ratio across the discussion phases.

Left-hemisphere markers: Categorisation / naming / fixing / representing / quantifying / certainty claims / extracting / manipulation / building up from parts / algorithmic / procedural / what can be done / action orientation / the explicit

Right-hemisphere markers: Holding open / attending / feeling / the whole / relationship / time and temporality / uncertainty acknowledgment / what resists naming / the living / the unfamiliar / instinct / body response / context / ambiguity as content rather than problem

Phase 1: Introductions (before the film)

Almost entirely left-hemisphere in character. Participants describe their fields in terms of what they produce, what they manipulate, what they can do:

> "Reducing that complexity into ways of processing that information in ways that can be more planned against and then devising activities, operations that we're able to navigate."

The notable exception is the evolutionary biologist, whose risk vocabulary introduced the only right-hemisphere observation in the introduction phase:

> "Resilience is actually interestingly opposed to strong and committed cooperation."

This relational, counter-intuitive, whole-system observation is the only one in this phase. It is significant that it comes from the person whose discipline operates on the longest timescale.

Phase 2: Immediate response to the film

The shift is marked and rapid. Within two or three exchanges of the post-film discussion, participants are in a qualitatively different register:

> "I felt like the techno sounds and the mechanical sounds were disturbing, and it didn't feel like they were of the same world."
>
> "As soon as you see the eyes, I sort of started feeling quite sorry for it."
>
> "It took me straight underwater."

These are right-hemisphere utterances: body response, felt quality, sense of presence, the uncanny. They are spontaneous, not produced by the exercise prompt. The left-hemisphere re-entry is also visible in real time — the media theorist's first move back toward the explicit: "My instinct is to try and figure out how it was made." But the analysis they produce is richer than it would have been without the affective encounter. The right-hemisphere exposure primed the left-hemisphere analysis with material it would not have generated independently.

Phase 3: Collective sensemaking

The groups are asked to produce a left-hemisphere output but resist more than the instructions suggest they should:

> "My instinct is to want to push against this, but I don't know if it's possible to even do that."

The group that refused to fill out the form stayed in right-hemisphere mode longest and produced the most sophisticated epistemological observation (whether the premise of the decision can be trusted). The group that moved most quickly to narrative produced the clearest account but the least nuanced one.

Phase 4: The creature exercise

The most sustained right-hemisphere engagement of the whole day. The vocabulary shifts markedly:

> "You have reverted forty times, each you is genetically identical but experientially distinct. Where is continuity located?"

The facilitator explicitly names what has been happening at the close:

> "Not everything is information, is it? Some things are felt and lived. And that is actually one of the intentions of the workshop — using sound and visuals and experience and talking and feeling together in the room today. It isn't just information."

Phase	Left-hemisphere	Right-hemisphere
Introductions	85%	15%
Immediate post-film	30%	70%
Consensus phase	60%	40%
Creature exercise	35%	65%
Plenary feedback	55%	45%

The artwork demonstrably shifted the mode of attention. The consensus phase partially reasserted the left-hemisphere default. The creature exercise re-opened the right-hemisphere mode. The plenary returned participants to a mixed mode that retained substantially more right-hemisphere engagement than the introductions.

The pattern is significant for the workshop's objectives. A cohort that begins in predominantly left-hemisphere mode — categorical, procedural, action-oriented — and ends with substantially more right-hemisphere engagement has not simply discussed a different topic. It has practised a different mode of attending to problems. That shift is what the methodology is designed to produce, and what repeated workshops are designed to deepen. The data here establishes a baseline: future sessions with the same cohort will show whether the shift becomes easier to activate, more sustained, and more readily available under operational pressure.

Analysis 5: Instinct-to-Articulation Tracking

This section shows how participants first registered something as a feeling — wrongness, sadness, anger, disorientation or being watched — and then used group discussion to turn that feeling into an explicit analytic point. The purpose is not to treat instinct as automatically correct, but to show how early signals can be tested and made useful.

Instance 1: Sadness

Pre-articulate: “As soon as you see the eyes, I sort of started feeling quite sorry for it.”

Articulation produced by group: “If you think this is something that has language and has a complex inner world and can communicate and has concepts and symbols, then it made me feel very, very sad.”

The feeling was right before the reasoning. The reasoning then worked back to explain what the feeling already knew.

Instance 2: Anger at the machine

Pre-articulate: “I was getting so mad.”

Articulation produced by group: “The machine feels to me entirely robotic and clumsy and we’ve designed something to go into an environment that is just not used to.” And then, later: “We’re the predator.”

The anger at the machine’s incompetence preceded the recognition that the incompetence was ethically significant — that the machine’s failure to operate properly was destructive.

Instance 3: The wrong note

Pre-articulate: “The techno sounds and the mechanical sounds were disturbing, and it didn’t feel like they were of the same world.”

Articulation produced by group: “If you took sounds of heartbeats and like blood pumping, and you mixed it with the sound of machines working, that’s the kind of thing we’d get... the temporality of sound is more linear, whereas with an image you’ve got like multiple dimensions that things are kind of working in.”

The disturbance preceded the analysis of why sound and image were carrying different ontological registers.

Instance 4: Being inside

Pre-articulate: “It took me straight underwater.” / “I thought I was inside.” / “For that little moment, it felt like we could have been looking through some other type of organism’s eyes.”

Articulation produced by group: the compound eyes discussion — “So a lot of organisms have actually many, many tiny eyes that they use to build the world” — which leads to: “Who is watching whom?”

The bodily sense of being inside the image came before the group developed a more practical question about surveillance: who is observing, who is being observed, and what security consequences follow when the direction of watching is uncertain?

Instance 5: The looping arms

Pre-articulate: “Instead of having a tentacle that is free, you have a circle. What do you do with a circle? That’s actually not — doesn’t make sense.”

Articulation produced by group: “Are those — if we were presented with these as a real thing, would we believe them?” Leading to: “Who is classifying whom? What does it mean that the classification system has become part of what it’s classifying?”

The bodily sense that the looping arm was wrong was the gateway to the epistemological question about the relationship between classification and classified.

> In every case, the pre-articulate response was tracking something real in the material. The group work did not produce the insight; it produced the articulation of an insight that the affective encounter had already generated. Art generates circumstances for pre-articulate knowing that group sensemaking then develops into expressible knowledge. The art is not illustrating what the analysis finds; it is the origin of what the analysis processes.

Analysis 6: Cross-Table Thematic Emergence

Which themes emerged independently across all three tables without coordination.

Theme 1: Unintended consequences of intervention

Three independent observations that...

Table 1: “Something is being gestured to which is like the AI trying to come to terms with what it’s been given.” Table 2: “An experiment that might have gone wrong at some level, or some intervention that misfired.” Table 3: “Getting it catastrophically wrong. A curiosity-driven act. We could see it easily both ways.”

Theme 2: The predator label

Three independent noticings of the same textual detail, arriving at compatible readings. Discovered independently by each group.

Table 1: “Even the fact that the mechanical arm in the water was labelled predator. Did you see that?” Table 2: “The predator thing is a real interesting label. We’re the predator.” Table 3: “It says Predator on the arm of the…”

Theme 3: The self-sampling moment

Three independent identifications of the most structurally significant moment in the film.

Table 1: “It’s almost like the AI’s brain is trying to come to terms with what it’s been fed on.” Table 2: “There’s a moment where the robot seems to be sampling itself or attempting to sample itself.” Table 3: “There’s a bit where it came up — trying to grab itself, but it unclips.”

Theme 4: Timescale

Each table developed a timescale observation from a different disciplinary angle: phenomenological (Table 1), biological (Table 2), ethical (Table 3).

Table 1: “The temporality of sound is more linear, whereas with an image you’ve got multiple dimensions.” Table 2: “Ontogeny recapitulating phylogeny and recapitulating evolutionary change.” Table 3: “These things will be evolving over hundreds of millions of years. And then that third one is just like... it just takes humans like a second to completely disrupt it.”

Theme 5: Categorisation as the problem

All three tables arrived at some version of the observation that classification systems are not neutral — that the act of categorising something changes what it is.

Table 1: “The breakdown of the act of classification and the things being classified.” Table 2: “We talked a lot about categorization and classification, both from the perspective of AI and from a scientist trying to describe a very unknown and unclear creature.” Table 3: “The way I’m reading this is that you’re seeing what it is and then you’re seeing a simplification and a labelling, a bit like a model of a molecule with balls on the end of sticks.”

Themes that did not converge

Curiosity vs. extraction — one table divided internally; others did not resolve it. This non-convergence is productive: it shows the question is genuinely contested.

Whether the creatures were watching or being watched — Table 1 divided on this explicitly; other tables leaned toward one reading.

The narrative ordering — each table arrived at a different account of whether the film moved forward or backward in time.

Five themes emerged independently across all three tables from a fifteen-minute film with no accompanying briefing on content. The themes that converged are structural (consequences, timescale, categorisation, the predator, the self-sampling) rather than narrative. The artwork reliably produces structural insight while allowing disciplinary variation in how that structure is inhabited and articulated. The themes that did not converge are questions whose answer depends on the position from which you look. Military strategists read extraction; scientists read curiosity. This is not disagreement about what is in the film — it is evidence that disciplinary position produces genuinely different perceptions of genuinely ambiguous material.

Output 2: Stress Points in Escalation and Coordination Logic

Five specific points of structural vulnerability were identified across the sessions.

STRESS POINT 1	The premise of the intelligence cannot be assumed reliable

When the images being assessed are of uncertain origin, when the sequencing of information may be manipulated, and when the framework for interpretation may itself be a product of the situation being assessed, conventional risk analysis has no secure foundation.

> “Were we trying to ascribe a narrative, and then loading it in with all our own kind of interpretive biases or not?”

STRESS POINT 2	Timescale mismatch between intervention and system

Every table identified a version of the same structural problem: the intervention operates on one timescale and the system being intervened in operates on another. Both failure modes — consequences accumulating over decades and consequences cascading within hours — were present in the rehearsed scenario.

> “These things will be evolving over hundreds of millions of years. And then that third one is just like — it just takes humans like a second to completely disrupt it.”

STRESS POINT 3	The self-sampling moment: system loses its boundary

The moment where the robotic arm appeared to turn its instruments on itself was identified by all three groups as the most significant single event in the footage. This applies directly to AI-mediated intelligence that is trained on, or used to validate, its own prior outputs.

> “There’s a moment where the robot seems to be sampling itself or attempting to sample itself. And then a general kind of discussion we had was on this idea of there being unintended consequences to an action that keep unfolding.”

STRESS POINT 4	Consensus has a security cost

The octopus group derived this by working from inside the creature: its camouflage requires all nine brains to agree on the signal, and during that alignment it is most visible to predators. The government

parallel was named directly in the room: rapid cross-departmental agreement can produce diluted decisions, suppress dissenting signals, and remove precisely the anomalies that may matter most.

> “Your camouflage requires your entire body to agree on the signal. What is the security cost of the moment of consensus?”

STRESS POINT 5	What cannot be named cannot be protected

The things that matter most in a complex system are often those that existing taxonomies have no categories for. The organisms in the film were disturbing precisely because they had absorbed the classification systems into their own form: the apparatus of knowledge had become part of what it was trying to know.

> “The breakdown of the act of classification and the things being classified — even mixtures of classification systems with the very biological things that they’re trying to classify.”

How Security Lancaster Can Help

The transcript analysis demonstrates that the complementarities between disciplinary positions identified in this workshop — between credibility and strategic frames, between empirical challenge and interpretive resistance, between maker's questions and epistemological ones — are not accidental features of this particular room. They are structural, and they deepen with practice.

A cohort that returns to this methodology across multiple episodes develops three things that a single workshop cannot produce: a shared vocabulary for exactly the class of problems the methodology addresses; a practised capacity to hold non-convergent readings without premature resolution; and an increasingly refined instinct for the moment when an abstract problem has become a live diagnostic. The transcript evidence from this session shows all three beginning to form. Future workshops are where they consolidate.

Security Lancaster can support this development in the following ways.

Cohort development. The Unsecurities Lab is designed as a repeatable format for a developing cohort. Participants who return to the methodology with accumulated experience produce richer complementarities and more precise diagnostics. Security Lancaster can design a programme of sessions tailored to the specific risk domains, disciplinary configurations, and institutional contexts of a given cohort — including the selection and commissioning of artworks, facilitation, transcript analysis, and production of findings reports.

Domain-specific workshops. The artwork and organisms are content variables; the structure is fixed. Future episodes can be directed at specific risk domains where the same questions of timescale mismatch, recursive consequence, and the limits of legibility-seeking frameworks apply directly. Other previous workshops included LUMI, on self-managing AI environmental systems. Upcoming Unsecurities Lab episodes include *Sunk Costs* — addressing the entangled material, financial and environmental consequences of the loss of the MV Felicity Ace — and a workshop examining AI image recognition, automated weapons, and ‘self-jail-breaking’ AI systems.

Bespoke design and delivery. Security Lancaster can be contracted to adapt the format for departments, agencies, research teams or partner organisations, including the choice or commissioning of suitable artworks, facilitation by a Security Lancaster lead with artist and domain expert as required, and production of a tailored findings report. A half-day format is standard; a shorter Session One-only version is available where time or resource is constrained.

Contact: securitylancaster@lancaster.ac.uk