

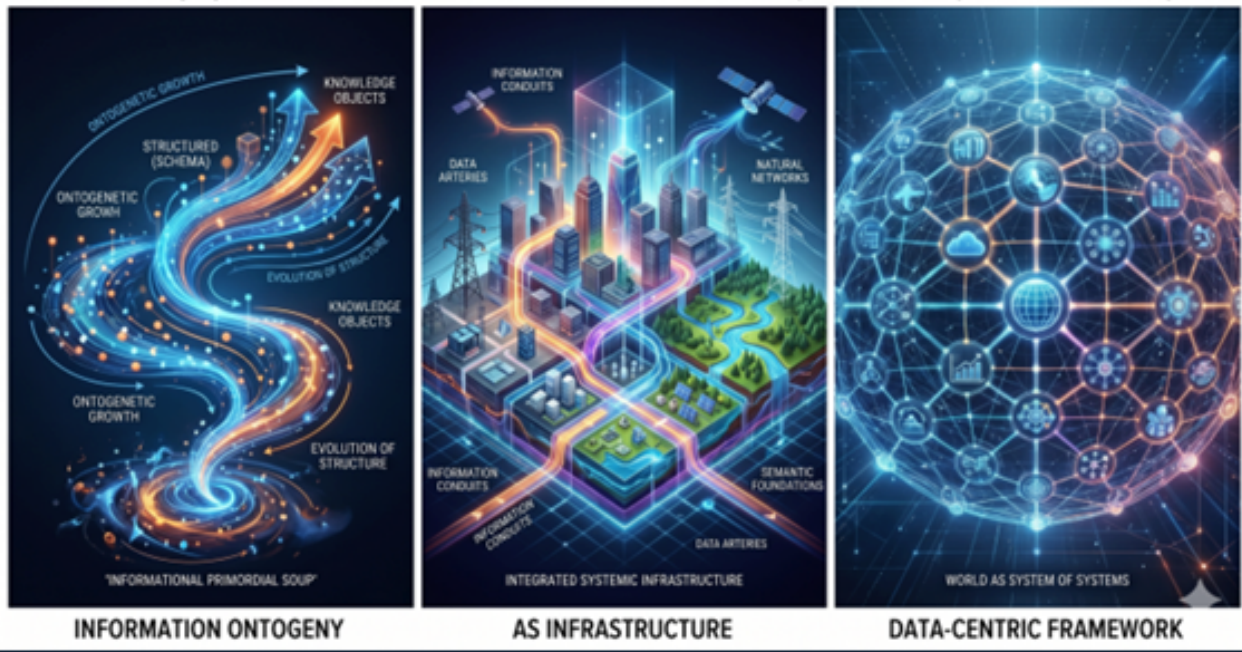
Information Ontogeny as Infrastructure

A Data-Centric Framework for Understanding the World as a System of Information Systems

Michael Ferrari, PhD

Founder & Managing Partner, Atlas Research Innovations
VP of Research, Moby · Senior Partner, AlphaGeo

Information Ontogeny as Infrastructure: A data-centric framework for understanding the world as a system of information systems



WORKING PAPER VERSION 5.0

MAY 2026 · COMPLEXITY & GEOSPATIAL INTELLIGENCE SERIES

Information Ontogeny as Infrastructure

A Data-Centric Framework for Understanding the World as a System of Information Systems

Michael Ferrari, PhD

Founder & Managing Partner, Atlas Research Innovations

VP of Research, Moby · Senior Partner, AlphaGeo

NEW IN VERSION 5

Section 5 - *The Acoustic Dimension*. Versions 1 through 4 of this paper argued that spatial provenance is the densest available key to the developmental history of an information asset. I still believe that. But the framework, on closer inspection, Depends on the preservation of developmental context, whatever the channel through which information is generated. Section 5 extends the argument to acoustic data, where the source-propagation-receiver geometry plays the role for acoustic data that spatial provenance plays for situated data. Two modalities, one epistemology. ***The thesis has not changed. Its surface area has widened.***

Author's Note: Two Traditions, One Claim

In late 2024, Dr. Fei-Fei Li, the computer scientist whose ImageNet dataset catalyzed the deep learning revolution, founded World Labs. The company was built on a single proposition: that reasoning in AI systems is fundamentally constrained by what those systems can perceive. “For humans,” Li wrote in her manifesto, “*spatial intelligence is the scaffolding upon which our cognition is built. It drives our reasoning and planning, even on the most abstract topics.*” Within its first year, the company raised over a billion dollars. It is engineering the consequence: AI systems whose outputs are not tokens or images but structured, physics-aware, generative representations of the world itself. ‘World Models,’ whose capacity to reason is grounded in their capacity to perceive spatial and physical structure.

Li arrived at her position from computer vision and AI. She also grounds her work in evolutionary biology and optics. I have arrived at the same position from spatial science, EvoBio, and complexity theory. The shared claim is straightforward: you cannot reason correctly about a system whose generative context you cannot perceive. If you strip the perceptual ground from a reasoning system, you have not simply reduced its precision. You have changed the kind of understanding it is capable of producing.

This updated working paper develops that claim at a different scale and in a different register. Where Li is concerned with AI systems perceiving 3D physical environments, this paper is concerned with analytical systems understanding complex social-ecological-economic information environments. While the technical apparatus is different, the epistemological premise is identical. And the intellectual lineage behind the version developed here predates both traditions. It goes back to the biologist Susan Oyama, who argued in 1985 that information does not reside in a medium waiting to be read out. It develops through the interaction of an organism and its environment, and its character is constituted by that developmental history.

Oyama was writing about genes and embryos. Li is building machines that can see. This paper applies the same argument to the data systems that underpin our understanding of the world. The result is a framework I call Information Ontogeny as Infrastructure. The developmental history of information - how it came into being, through what relational context, across what spatial and temporal scales - is not incidental metadata. It is the primary substance of what the information represents, and therefore the foundational lens through which complex systems can be understood.

****Note on Version 5**.** Versions 1 through 4 of this paper argued that spatial provenance is the densest available key to the developmental history of an information asset. I still believe that. But the framework, on closer inspection, does not depend on space being the operative dimension. It depends on the preservation of developmental context, whatever the channel through which information is generated. Section 5, added in this version, extends the argument to **acoustic data**, where the source-propagation-receiver geometry plays the analogous role of dense developmental key. So we have two modalities, but one epistemology. I am comfortable in saying that the thesis has not changed - its surface area has widened.

Abstract

Standard data practice treats information as a static resource – it is collected, cleaned, modeled, consumed. In one sense, there is nothing wrong with this. Data is the raw material on which decisions are based. But this treatment is misaligned to the better understanding of complex systems.

In this paper, I argue that information is not a resource. It has an ontogeny. Borrowing the analogy from biology, information contains a developmental life-cycle and history produced by the interaction of a measuring process and its environment. That history is instructive with respect to what the information means. Data stripped of its developmental and organizational context is less enlightening. It becomes a number rather than a signal embedded in a structural relationship.

The central claim is that information ontogeny functions as infrastructure. It is the foundational substrate that determines what relationships are knowable from a given data asset, what predictions are structurally possible, and what understanding is achievable. The world is best understood as a system of interacting information systems, each with its own ontogenic dynamics, coupled to each other through structural pathways that can be identified, characterized, and monitored. For situated data, the data-centric spatial lens is the primary instrument for reading these pathways - because spatial provenance is the richest single key to an information asset's developmental history.

The framework operates above any single sensing modality. Section 5 of this version extends the argument to acoustic data, which exhibits the same ontogenic structure with a different primary key: the source-propagation-receiver geometry takes the place that spatial provenance occupies for situated data. Spatial and acoustic are parallel lenses on the same physical environment. The thesis is a multi-modal-perception thesis: a commitment to reading developmental context as constitutive of meaning, regardless of which channel through which the physical world communicates.

The epistemological position - that reasoning is a function of perceptual grounding, and that grounding is constituted by ontogenic context - is now being deployed at massive commercial scale in AI by researchers including Fei-Fei Li, whose World Labs is engineering AI systems whose reasoning is grounded in spatial and physical perception. From that tradition, and from the complexity science tradition developed here, the same conclusion emerges. Stripping the generative context from information does not simply degrade analytical performance. It makes a class of structural understanding impossible. This paper attempts to describe what that class of understanding looks like, why it matters, and what it elucidates that other approaches cannot find.

1. Information Has a Life

“Developmental information does not arise from nothingness. It develops always from the conditional transformation of prior structure - by ontogenetic processes.”

— Susan Oyama, *The Ontogeny of Information* (1985)

In 1985, the biologist Susan Oyama published a slender and radical book. The prevailing dispute in developmental biology at the time asked whether the information guiding an organism's development resided in its genes or in its environment. Oyama argued that the question was not framed correctly. Information, she insisted, does not reside anywhere. It is produced through the interaction of organism and environment. It has a developmental history - an ontogeny - that is not separable from what the information is and what it does. The gene is not the carrier of pre-formed biological meaning. It is one participant in a developmental process through which meaning is produced.

The extension of this insight to data is direct and overdue. Information is not a carrier of pre-formed meaning waiting to be read out of a database. It is the product of a generative process: the interaction of a measuring instrument with a physical or social phenomenon situated in a specific context. That context is not a metadata label appended to the observation. It is the condition of the data's origination, and therefore part of what the information is and what it represents. Strip the context, and you have not simplified the data. You have amputated the part of it that encodes structural meaning.

The four stages of biological ontogeny apply directly to spatial data.

1.1 Birth: The Generative Context

Every piece of spatial data is produced by a process located in the world. A satellite sensor reading. A shipping manifest. A clinical measurement. A weather station observation. A commodity futures price. A regulatory filing. Each of these emerges from the interaction of a measuring instrument with a phenomenon situated in space, time, institution, and ecology. That situation is the datum's birthplace, and the birthplace encodes a dense bundle of structural context.

A nitrogen dioxide reading from Sentinel-5P over the Yangtze River corridor encodes Chinese manufacturing activity, regional atmospheric dynamics, the regulatory environment governing industrial emissions in that province, the seasonal climate system, and the coupling of that location to global atmospheric teleconnection pathways. None of this is in the numerical value of the reading itself. All of it is in the reading's spatial and temporal provenance. Preserve the provenance and the reading is a signal indexing a complex of structural relationships. Remove it - as we are seeing with the current dismantling of scientific observation infrastructure - and the reading is reduced to a number that can be correlated with other numbers, but cannot be placed in the world.

The birthplace is primary. Recovering or preserving it is the first task of any analysis that aspires to structural understanding.

1.2 Development: Relational Meaning Through Contact

Information does not arrive at its full meaning at the time of its birth. It develops meaning through contact with other information systems. A monsoon anomaly in the Bay of Bengal acquires pharmaceutical significance through its structural coupling to water availability at manufacturing facilities in Hyderabad. A crop stress index in Brazil acquires meaning for antibiotic production through the fermentation feedstock dependencies of pharmaceutical manufacturing. A port congestion signal in the South China Sea acquires meaning for retail inventory dynamics through the supply chain networks that route through it.

None of these relational meanings is intrinsic to the originating signal. Each is produced by the structural coupling between information systems - the developmental contact between a signal and the network of other signals and processes it encounters as it moves through the world. The same data, moving through different relational networks, acquires different meanings. Context is the mechanism of meaning production. It operates throughout the information's developmental trajectory, not only at its birth.

1.3 Maturity: The Actionability Window

Information, like biological fitness, experiences a maturity window - the period of peak predictive value, during which its relational meaning is sufficiently developed to be actionable. Before that window, the signal may be detectable but lacks the relational context to be interpretable. After that window, the information ages out of relevance as the process it indexed evolves beyond the snapshot it captured.

The maturity window is structurally determined. A sea surface temperature anomaly in the tropical Pacific reaches peak predictive value for downstream agricultural outcomes through a propagation sequence - atmospheric teleconnection dynamics - that unfolds over weeks to months, and in some cases years. That propagation sequence is the information's developmental trajectory. The maturity window is its peak. The characteristic lead time between originating signal and receiving-system outcome is the expression of its developmental velocity. Understanding information maturity means understanding the developmental dynamics of the system that produced it, not simply estimating a statistical lag.

1.4 Senescence: Preparing for the Next Generation

Senescence is the final stage in the information lifecycle. Data deconstructs and 'fades,' and in the process provides the raw material to spawn the next generation of information. Information ages at rates set by its generative system's dynamics: financial market data in milliseconds, satellite cropland imagery across growing seasons, infrastructure survey data across years of physical change. Old information also reactivates. A 150-year sea surface temperature time series is the developmental record of a climate system whose long-period variability may now be the most important predictive feature for infrastructure planning over decadal timescales.

Information ontogeny describes how data actually acquires, sustains, and loses meaning in complex systems. Treating data as a static resource rather than a developmental trajectory is the primary source of structural blindness in contemporary analytical practice.

2. Ontogeny as Infrastructure

“There appears to be a growing decoupling between the conditions that infrastructures were designed for and today's rapidly changing environments.”

— Mikhail Chester and Braden Allenby

Infrastructure is a particular kind of foundational system - one whose value lies not in what it does directly, but in what it enables other systems to do. A road network does not move goods. It enables goods to be moved. An electrical grid does not compute. It enables computation. Infrastructure shapes what flows, at what speed, with what reliability, and with what vulnerability to disruption. Its defining characteristic is generative: it determines what is structurally possible and what is not for the systems it serves.

Information ontogeny functions as infrastructure in exactly this sense. The developmental history of information - the chain of generative contexts, relational contacts, maturity trajectories, and decay rates that constitute its meaning - serves as the substrate that determines what is analytically and physically possible. Preserve the substrate, and structural relationships become visible that would otherwise be inaccessible. Strip it, as most current data pipelines do, and the analytical capacity is degraded in kind. An entire category of understanding becomes impossible.

2.1 Ontogenic Infrastructure Shapes What Is Knowable

Even as an engineering construct, physical infrastructure is not neutral. A road network does not simply connect existing destinations. It determines where destinations form, where economic activity concentrates, and where isolation persists. It is an enabler of resource transfer, and also of disease transmission. The configuration of infrastructure determines what is accessible and what is not - and therefore what is possible and what is not. Ontogenic information infrastructure operates the same way, or at least it should.

The degree to which an information asset's developmental history is preserved determines what structural relationships are visible. An organization that retains the spatial provenance of its transactional data, the institutional context of its regulatory filings, and the temporal provenance of its sensor readings has access to a class of structural inference that an organization that strips these attributes does not. This is not a quantitative difference. It is the difference between being able to see a category of relationship and being blind to it. Ontogenic infrastructure defines the epistemic terrain of what can be known.

2.2 The Dark Data Problem, Reframed

A great deal of data that is collected is never used. This is typically interpreted as a data management problem: the data exists, but is difficult to find and access. The ontogenic perspective reframes it more precisely, and more consequentially.

Surveying the commercial data landscape, a pattern is consistently repeated regardless of an organization's size or data-sophistication: most enterprise data is ontogenically dark. It has been collected in ways that stripped its developmental context. The spatial provenance was never recorded, was eliminated, or was simply lost. The temporal sequence has been compressed into aggregates that erase the dynamics originally indexed. The institutional and relational context of production has not been retained. The data exists as numbers, but the information - the structured meaning that can be related to other systems through its developmental history - was never fully exploited.

Dark data recovery is therefore not a retrieval problem. It is an ontogenic reconstruction problem. Some of what was lost can be recovered - spatial attributes inferred from partial records, temporal sequences partially reconstructed from audit trails and statistical interpolation - but the recovery is always incomplete. The argument for building ontogenic infrastructure at the point of data collection, rather than attempting reconstruction after the fact, is the same argument for laying pavement before a city is built, not after.

2.3 Space as the Primary Developmental Dimension

Of all the dimensions of information's generative context - spatial, temporal, institutional, relational - space is analytically paramount for situated data. This is not because place is more important than time or institution in any absolute sense. It is because a location in space encodes, in a single coordinate pair, the densest available concentration of generative context.

Climate makes the case cleanly. A spatial coordinate encodes the variables that constitute the climate system in space and time, as the information is subject to teleconnective couplings to other biological, physical, chemical and economic systems. It encodes the ecological regime in which the location sits. The regulatory jurisdiction it falls within. The physical infrastructure it has access to. The economic geography of markets and supply chains that shape its context. The institutional history of the specific place. No other single attribute encodes a comparable density of generative context in a bias-free manner. This is why the data-centric spatial lens is the appropriate instrument for reading information ontogeny: spatial provenance is the richest single key to an information asset's developmental history.

The data-centric spatial lens is not a methodology for making maps. It is a methodology for reading the developmental history of information, using spatial provenance as the primary key. The map is the output. The understanding of relational structure is the goal.

3. Information Systems of Systems

“The universe is made of energy, matter, and information, but information is what makes the universe interesting.”

— César Hidalgo, *Why Information Grows*

The world is not a single complex system. It is a system of systems - a configuration of partially overlapping, structurally coupled, independently evolving information systems, each with its own spatiotemporal dynamics, each generating information that flows into and shapes the others. Understanding the world analytically means understanding this configuration: which systems are coupled to others, through what directional and structural transfer pathways, at what timescales, and with what developmental consequences for the information that traverses those routes.

This is a stronger claim than the standard assertion that ‘everything is connected.’ Connectivity is a structural fact. The information-systems-of-systems framework adds an account of the informational dynamics of connection: how information born in one system develops relational meaning as it traverses coupling pathways, how it matures into actionability at specific points in its journey, how it arrives transformed in receiving systems and influences their behavior. These developmental dynamics are the substance of what connections actually do. They move us from ‘what’ to ‘where and why.’

3.1 Generative Systems and Ontogenic Transfer Pathways

Information systems, in this framework, are not software platforms or databases. They are generative processes that produce information with characteristic ontogenic properties: specific spatial and temporal scales of production, specific relational coupling pathways to other systems, characteristic maturity windows and decay rates.

Take climate. The global monsoon system is a physical information transfer system. It produces temperature, precipitation, and atmospheric dynamics signals at subcontinental-to-continental spatial scales and seasonal-to-interannual temporal scales, with well-documented coupling pathways to other climate systems and to agricultural, hydrological, and infrastructure systems. The global pharmaceutical trade network is a different type of system entirely. It produces shipping manifests, production records, regulatory filings, and inventory dynamics that index the state of global drug supply at facility-to-regional-to-global spatial scales and daily-to-quarterly temporal scales. A sovereign bond market is an information system. The emerging network of IoT sensors embedded in physical infrastructure is a physical-technological-human information system. Each one generates information with a characteristic ontogeny.

Coupling between information systems operates through ontogenic transfer pathways: structural channels through which information born in one system develops new meaning as it enters another. These pathways are determined by physical geography - shared climate zones, transportation corridors,

water systems. By institutional architecture - regulatory jurisdictions, trade agreements, legal frameworks. By ecological structure - shared biological production chains, environmental dependencies. And by the physical-digital infrastructure that increasingly mediates all of these. Understanding the pathways is understanding the developmental mechanisms of inter-system coupling.

3.2 Scale, Nesting, and the Granularity of Meaning

A system-of-systems framework needs flexibility with respect to scale and nesting.

A crop stress signal from a MODIS satellite pass can be read at the level of a specific field, a county, a watershed, a national agricultural system, or a continental production zone. Each level of spatial aggregation produces a different piece of information with a different ontogenic character - different relational couplings, different maturity windows, different decay rates. A field-level signal is most actionable for a farm operator on a daily-to-one-week timescale. The continental aggregate is most actionable for a commodity trader on a three-to-six-month timescale. The national estimate is most actionable for a pharmaceutical procurement officer planning fermentation input sourcing across a growing season.

These are not independent pieces of data. They are nested stages in a single ontogenic hierarchy. The field-level signal is developmentally prior to the county-level aggregate, which is prior to the watershed-level integration, which is prior to the national production estimate. Meaning is produced at each transition. Meaning is also irreversibly lost through aggregation - a form of ontogenic compression that discards some of the developmental history of the original signal. The question of what scale to analyze should be determined by the scale of the decision the information is intended to inform, and the ontogenic pathway between originating signal and analytical output should be explicitly traced and preserved.

3.3 Boundaries as the Sites of Maximum Ontogenic Activity

In a system of information systems, the boundary zones - where one information system transitions to another - are the most analytically productive territory. These are the sites where ontogenic information transfer occurs. Where information born in one generative context undergoes transformation as it enters a new relational environment. They are also the sites where new meaning is produced, where coupling mechanisms are most legible, and where the structural relationships between distant systems first become visible.

The boundary between the climate information system and the agricultural information system is where atmospheric dynamics are translated into soil moisture, heat stress, and growing-season anomalies. The boundary between the agricultural system and the pharmaceutical system is where crop production anomalies translate into fermentation feedstock availability and cost, which translate into antibiotic and vitamin production constraints. The boundary between the pharmaceutical system and healthcare delivery is where drug supply dynamics translate into patient access outcomes. Each boundary is a site of ontogenic development. Information arrives from one system, undergoes developmental

transformation through the coupling mechanism, and departs as information of a different character - carrying new relational meaning - into the receiving system.

The full developmental chain from an atmospheric signal to a healthcare outcome crosses at least four information-system boundaries, each mediated by a specific coupling mechanism. No single-domain framework can trace this chain. Understanding it requires a framework that treats the chain itself - the sequence of boundaries and the ontogenic transformations at each - as the primary object of analysis.

3.4 Teleconnections as Developmental Signatures

Teleconnections are the observable signatures of coupling between distant information systems. In climate science, an El Niño event in the tropical Pacific is coupled to monsoon anomalies in South Asia, drought in Australia, and winter temperature anomalies in North America. The teleconnection is not a statistical correlation between variables in these distant regions. It is the developmental trace of information - physical energy transfer - born in the tropical Pacific Ocean-atmosphere system as it traverses the global atmospheric information system, through the coupling mechanism of Rossby Wave propagation, and arrives, transformed but structurally related, at receiving systems on multiple continents.

This developmental framing is essential for understanding why well-characterized teleconnections have durable predictive power while statistical correlations often do not. A correlation is identified in historical data and applied as a predictor. It persists as long as the underlying coupling remains stable, and it fails when the system reorganizes. A teleconnection grounded in an identified coupling mechanism - *where the ontogenic transfer pathway is characterized, not merely the correlation between endpoints* - is robust to many forms of system reorganization. The mechanism, not the historical correlation, is the source of its predictive content. When the system reorganizes, the mechanism changes in ways that are observable and recharacterizable. The framework updates. The correlation simply fails.

Supply chain teleconnections operate by exactly this logic. The sea surface temperature anomaly in the Bay of Bengal is coupled to Indian monsoon dynamics, which are coupled to water availability at pharmaceutical manufacturing clusters in Andhra Pradesh and Telangana. These are coupled to active ingredient production volumes for the global generics market. Each coupling is an ontogenic transfer pathway. The signal travels through a developmental chain, acquiring new meaning at each stage. The predictive value of the originating signal for the terminal outcome - a drug supply disruption in distant markets - derives from the structural integrity of those pathways, not from the historical statistical association of the variables at the chain's two ends.

3.5 The Physical-Digital Layer and Accelerated Ontogeny

The emergence of a globally distributed digital infrastructure layer has not simply added a new information system to the world. It has created a new class of ontogenic transfer pathway that operates at timescales orders of magnitude faster than the physical coupling pathways it overlays, and in doing so has altered the developmental dynamics of every information system it touches.

Before this digital layer, information born in a physical system propagated through physical coupling pathways - commodity trade flows, population movement, infrastructure dependencies. These pathways had characteristic timescales of weeks to months for commodity markets, months to years for population adjustments, years to decades for infrastructure transitions. The ontogenic transfer was slow enough that developmental trajectories were often visible and the information's provenance could be tracked.

The digital layer has collapsed these timescales, but non-uniformly. A satellite observation of a flood in a major agricultural region now propagates into commodity futures markets within minutes through algorithmic trading systems. The ontogenic transfer, from physical event to market price signal, has been compressed from days to minutes. But the relational meaning of the resulting price signal still carries the spatial birthplace of the original observation: a flood in a specific place, indexing specific crop dependencies, with specific downstream consequences. The developmental velocity has changed. The spatial structure of the meaning has not.

What is genuinely new is the emergence of information systems native to the digital layer - social media dynamics, algorithmic market behaviors, platform-mediated demand patterns, AI-generated signals. These systems produce information with ontogenic properties that have no physical analogue. The information is born via digital processes, develops through digital coupling, and decays at digital timescales. And it is now coupled to physical information systems in both directions, creating feedback dynamics that neither layer exhibits independently. The physical-digital boundary is the most rapidly evolving information system boundary currently active. It is the site of some of the most consequential and least understood ontogenic processes in the contemporary world.

The digital layer does not represent the physical world. It participates in it. Ontogenic transfer pathways now operate simultaneously through physical and digital couplings, at different timescales, producing information whose meaning is constituted by both.

4. The Perception-Reasoning Dependency

“Spatial intelligence is the frontier beyond language - the capability that links imagination, perception, and action.”

— Fei-Fei Li, *World Labs Manifesto* (2025)

The argument developed in the preceding sections - that information ontogeny is foundational to meaning, and that stripping ontogenic context makes a class of structural understanding impossible - has an exact parallel in a different tradition. Fei-Fei Li, whose ImageNet project gave the deep learning revolution its primary training substrate, has spent the past several years arguing that language-based AI systems are subject to a fundamental limitation. They cannot reason correctly about the world because they cannot perceive it. They process tokens that represent the world. But as argued above, tokens stripped of the spatial, physical, and geometric context that gives worldly phenomena their structure. The consequence is not that these systems are slightly wrong. It is that they cannot achieve a category of understanding that requires perceptual grounding.

Li calls this category ‘*spatial intelligence*,’ and she founded World Labs to engineer its computational realization: AI systems whose reasoning is grounded in structured, physics-aware, generative representations of 3D space - world models. The goal is not better image generation. The goal is reasoning grounded in the geometry and physics of the world, and therefore capable of a kind of inference about spatial and physical relationships that token-based systems cannot approach. This is how babies learn and reason about the world. The logic makes sense.

The structural parallel to the Information Ontogeny framework is precise, and it is worth stating explicitly.

4.1 The Isomorphism

Li's argument, translated into the language of this paper: a system whose training data has been stripped of spatial and physical context - reduced to tokens - cannot reason correctly about the spatial and physical world. The stripping of perceptual context does not simply reduce precision. It makes a class of understanding inaccessible.

This paper's argument: an analytical system whose data has been stripped of ontogenic context - reduced to numbers - cannot reason correctly about the information systems of the world. The stripping of developmental context does not simply reduce precision. It makes a class of structural understanding inaccessible.

I believe the isomorphism is exact. In both cases, the failure mode is the same: a reasoning system operating on context-stripped representations cannot access the structural relationships that are only visible in the full developmental context of the information it is processing. In both cases, the corrective is the same: restore the perceptual, spatial, physical and ontogenic grounding that was stripped in

preprocessing, and the reasoning capacity recovers a category of understanding that the stripped version could not achieve.

The scale of application is different. Li is concerned with 3D physical environments at the scale of objects and rooms. This paper is concerned with complex social-ecological-economic information systems at the scale of supply chains, climate systems, and global markets. The technical apparatus is different too. World Labs is building neural world models. This paper uses geospatial analysis, network science, and complexity theory. But the epistemological premise is shared and, I would argue, fundamental.

4.2 The Sim-to-Real Gap as an Ontogenic Failure

World Labs' work provides one of the most vivid illustrations of the ontogenic failure this paper diagnoses. The 'sim-to-real gap' is the name given to a pervasive problem in robotic AI: a system trained in simulation fails to generalize to physical reality - often catastrophically - due to slight variations in lighting, surface texture, object positioning, or environmental physics that the simulation did not capture. The roots of this problem can be traced to early climate modeling efforts, via Ed Lorenz's observations of how initialization error in complex dynamical systems yields wildly different results. The simulation and the physical world produce data that looks the same in many surface respects but was born through radically different generative processes. The simulated world has a different ontogenic history than the physical world.

It follows that a model trained on simulated data has learned from ontogenically stripped information - data born in a procedurally generated environment that lacks the physical, material, and historical generative context of the real world. Synthetic data is great, until it isn't. When that model is deployed in the real world, it encounters information whose birthplace it has never seen and whose generative context it cannot read. The failure is not a failure of model capacity. It is a failure of ontogenic grounding. Until synthetic data can construct a synthetic ontogeny, this will remain a roadblock.

This is identical in structure to what happens when a supply chain model trained on historical data is deployed into conditions produced by a structural reorganization it has never encountered. The historical training data was born in a system configuration that no longer exists. The deployment data is born in a new configuration whose generative context the model has not been exposed to. The model fails not because it lacks parameters or equations, but because it cannot read the ontogenic context of the information it is now receiving. In both cases, the corrective is the same. Build analytical systems that are grounded in the structural character of their information's generative context - not only in the statistical patterns of its historical outputs.

4.3 World Models as Formalized Ontogenic Substrates

The world model concept that Li and World Labs are developing is, in the language of this paper, an engineered formalization of information ontogenic infrastructure at the AI-system level. A world model is a structured representation of the generative context of visual and physical information. It preserves the geometry, physics, and semantic structure of the environment that produced the training data. It

enables the model to reason about spatial and physical relationships rather than pattern-matching on pixel or token sequences.

In this paper's terms, a world model is an attempt to preserve the ontogenic substrate of visual information so that the reasoning system built on it can access the structural relationships that are only visible when generative context is retained. This is exactly what the Information Ontogeny framework asks of analytical systems operating on social-ecological-economic data: not just that data be collected and processed, but that the developmental history of the data be preserved and made available to the analytical system as primary structure, not stripped out as preprocessing overhead.

The technical implementation is necessarily different across the two domains. World Labs is building neural architectures that learn 3D world representations from visual data. The Information Ontogeny framework is building analytical architectures that preserve and analyze the spatial provenance, temporal trajectory, and relational coupling history of complex systems data. But the functional goal is the same: an analytical system whose reasoning is grounded in the generative context of its information, and that therefore has access to a category of structural understanding that context-stripped systems cannot achieve.

The goal is reasoning grounded in the geometry and physics of the world, and therefore capable of a kind of inference about spatial and physical relationships that token-based systems cannot approach.

5. The Acoustic Dimension

Information Ontogeny in a Different Medium

“Sound is touching at a distance.”

— R. Murray Schafer, *The Soundscape* (1977)

Working through versions 1 through 4 of this paper produced an argument I still believe: that spatial provenance is the densest available key to the developmental history of an information asset. The empirical cases I have worked through (pharmaceutical supply chain concentration, climate-driven food security teleconnections, agricultural and commodity dynamics) have continued to bear it out.

But the argument as written has been confined to a single sensing modality, and on closer inspection, the Information Ontogeny framework does not depend on space being the operative dimension. What it depends on is the preservation of developmental context, whatever the channel through which the information is generated. Space is the densest concentration of context for situated data. For other data, the densest concentration is somewhere else.

This section is about a second channel: acoustics. Sound is produced at a source, propagates through a medium whose state it records, and arrives at a receiver whose properties further shape the signal. That source-propagation-receiver chain is the ontogeny of the acoustic information. Read correctly, a recorded waveform is not a measurement of pressure variation. It is a chronicle of the system that produced it.

Two modalities. One epistemology. The argument has not changed. Its surface area has widened.

5.1 Sound Has an Ontogeny

Anyone who has spent time around acoustic data knows that the waveform at the receiver is conditioned by the entire path it took to get there. A geophysicist reads source location and depth from a seismic record. A sonar operator reads water temperature and salinity from the same recording that contains a vessel signature. A reliability engineer reads bearing wear from the spectrum of a pump. The signal is not just a number, and ***the path is in the signal.***

The framework I have been developing for spatial data applies, almost verbatim, to this practice. Versions 1 through 4 used four ontogenic stages - birth, development, maturity, senescence - to describe how situated data acquires, sustains, and loses meaning. The same four stages apply to acoustic data with a different set of indexing dimensions. For situated data, the stages index location, relational context, an actionability window, and decay. For acoustic data, they index a source event, a propagation history through specific media, an arrival window at the receiver, and a signal-decay envelope. The structure is preserved. The medium is different.

Take a low-frequency atmospheric pressure oscillation recorded at an infrasound array in central Greenland. Numerically, it is a time series. Read through its ontogeny, it is the developmental record of

an event that may have occurred thousands of kilometers away - a volcanic eruption, an ocean storm, a meteor entry, a rocket launch - whose signature was filtered through stratospheric waveguides whose temperature and wind structure determined which frequencies and which arrival angles survived the journey, and arrived at the sensor in a specific time-frequency configuration. None of this is in the numerical value of the recording. All of it is in the ontogenic context. Strip the geometry and the recording is a number. Preserve it, and the recording indexes a complex of physical relationships connecting the array to a distant generative event, and to the state of the atmospheric medium between them.

5.2 Source-Propagation-Receiver Geometry as Primary Key

In Section 2.3, I argued that of all the dimensions of information's generative context, space is analytically paramount for situated data - because a spatial coordinate encodes the densest available concentration of generative context for situated phenomena. The argument depended on the epistemological character of situated data. Most variables of interest in social-ecological-economic systems are located in physical space, and most relevant generative context - climate, regulation, infrastructure, ecology, institutional history - is encoded in or derivable from that spatial coordinate.

For acoustic data, the densest key is different. A pressure recording at a single point does not encode location-as-such. It encodes the geometry of the chain that produced it. The bearing from which the signal arrives, the travel time, the spectral evolution along the path, and the spatial relationship between source and receiver collectively constitute the primary key. Two recordings with identical numerical content can index different events if their source-propagation-receiver geometries differ. Two recordings with very different content can index the same event if the geometry chain is correctly read. The primary key is the chain, not the cell.

This is the acoustic counterpart to the spatial-lens claim, and it has direct practical consequences. An acoustic methodology that preserves source-propagation-receiver geometry - bearing, travel time, frequency content as a function of path - surfaces structural relationships that a methodology treating the signal as a time series of numbers cannot access. Infrasound arrays, distributed acoustic sensing systems, ocean hydrophone networks, urban microphone deployments, industrial vibration sensors - these are not collections of point sensors. They are instruments for reading acoustic ontogeny, in the same sense that satellite constellations are instruments for reading spatial ontogeny. The map is one output of the spatial lens. The time-frequency-bearing array is one output of the acoustic lens. The structural relationships they make legible are the same kind of object.

5.3 Acoustic Teleconnections as Mechanistic Lead Indicators

On 15 January 2022, the Hunga Tonga-Hunga Ha'apai volcano erupted underwater in the South Pacific. Within hours, the acoustic pressure pulse from the eruption was being recorded by infrasound stations on every continent through the Comprehensive Nuclear-Test-Ban Treaty Organization's global network. Over the next several days, the signal circled the planet multiple times.

The volcano was also doing something else, less obvious in the moment but more consequential over the long term. It was injecting a startling volume of water vapor into the stratosphere - enough to perturb

stratospheric chemistry and the radiative balance of the lower stratosphere for the better part of three years. The acoustic signal arrived first. The atmospheric chemistry perturbation showed up in composition measurements weeks to months later. The surface climate consequences are still being measured.

This is a textbook *acoustic teleconnection*. The chain runs from generative event, through stratospheric chemistry, to surface climate consequence, with a characteristic developmental velocity that makes the originating signal a structural lead indicator on a horizon no statistical correlation analysis could deliver. If you happened to be paying attention to the acoustic signal, with the right framework in place, you had a meaningful head start on the cascade that followed.

Hunga Tonga is the dramatic case. The continuous case is more useful. Microbaroms - a continuous, atmosphere-wide low-frequency pressure oscillation generated by the nonlinear interaction of ocean surface waves in the North Atlantic, North Pacific, and Southern Ocean - propagate inland through stratospheric waveguides and are detectable on every continent by appropriately configured infrasound sensors. In a precise and continuous sense, the microbarom signal is the acoustic developmental record of remote ocean wave climate. The generative source is open ocean. The receivers are land-based atmospheric sensors. The information content arrives on a timescale of hours.

Variations in microbarom amplitude, frequency, and arrival direction encode changes in the underlying ocean wave field. Reading them is reading the developmental history of distant ocean dynamics through an atmospheric waveguide. The wave field eventually manifests in coastal infrastructure exposure, shipping insurance, and offshore operations risk. The microbarom signal arrives first.

There is a third example, more recent and possibly the most consequential for commercial and policy applications. Distributed Acoustic Sensing - DAS - uses a laser to interrogate an installed fiber-optic cable and convert every meter of cable into an acoustic strain sensor. The infrastructure is the cable. The cable already exists; transoceanic and continental fiber networks have been built and operated for decades. Adding DAS to the cable does not require new construction. It requires a different interrogation strategy at the cable end.

A transoceanic cable becomes a continental-scale sensor array at no incremental infrastructure cost. DAS has been used to detect earthquakes hundreds of kilometers from the cable, to monitor ocean wave climate at landing points, to track marine traffic and submarine activity, and to surveil oilfield wells with spatial resolution that previously required dedicated equipment. The fiber's geographic topology - which the Information Ontogeny framework already reads as a structural concentration of physical infrastructure - becomes an acoustic sensor array whose developmental signal preserves geometry as a first-class attribute. The same ontogenic logic that surfaces concentration risk in the spatial lens surfaces, in the acoustic lens, an entirely new layer of legibility on the engineered subsurface.

These are three separate teleconnections with three different mechanisms. But they share the same structural logic. None of them is a correlation, and each is a propagation pathway whose physics is known and whose ontogeny is, in principle, traceable end-to-end.

5.4 Two Modalities, One Epistemology

The methodological claim is symmetric in both directions. An analytical posture that preserves the developmental geometry of an acoustic signal can access structural relationships that an acoustic methodology stripping that geometry cannot. An analytical posture that preserves the spatial provenance of a situated signal can access structural relationships that a financial methodology stripping that provenance cannot. The operative epistemology is the same - the Information Ontogeny framework - and the failure mode of context-stripped analysis is structural in both cases, not technical.

The implication is that the right unit of analytical commitment is not 'spatial finance' on the one hand or 'acoustic intelligence' on the other. It is **multi-modal perception**, executed under the discipline of ontogenic preservation. Spatial and acoustic data are parallel lenses on the same physical environment, not substitutes. Spatial reads the situated state of the system. Acoustic reads its activity, its transitions, and the propagation of its generative events through air, water, fiber, and structure. The combined readout is fuller than either lens alone.

This has direct consequences for how the perception layer of the broader analytical stack should be understood. The optical, synthetic-aperture-radar, and hyperspectral platforms that dominate the current commercial expression of Earth observation are one half of an emerging multi-modal sensing infrastructure. The acoustic platforms - DAS on fiber networks, atmospheric infrasound arrays, ocean hydrophone systems, urban acoustic deployments, industrial vibration monitoring, bioacoustic networks for ecological monitoring - are the other half. They converge under the same epistemological framework. They will increasingly converge under the same commercial and analytical infrastructures. Operators, investors, and policymakers building toward that convergence will see the physical world more accurately, and earlier, than those operating in a single sensing modality.

As stated above: Two modalities, one epistemology. The thesis has not changed. Its surface area has widened.

The thesis has not changed. Its surface area has widened.

6. What This Methodology Finds That Others Cannot

The information ontogeny framework leads to a specific category of finding that is inaccessible to the two dominant analytical alternatives - pure AI-driven pattern recognition, and pure domain expertise. The reason in each case is now clear from the preceding argument. This section makes the contrast explicit. With Section 5 in place, the argument extends symmetrically across both the spatial and acoustic lenses.

6.1 What Pure AI Cannot Find

Contemporary AI systems are powerful pattern recognizers operating on ontogenically stripped training data. Spatial provenance is typically reduced to numeric coordinates. Temporal sequence is reduced to lag features. Institutional context is often absent. Source-propagation-receiver geometry, when acoustic data is involved at all, is collapsed to a raw waveform or a low-dimensional embedding. The developmental history of the data - which information systems it passed through, what relational transformations it underwent - is not preserved. The model learns from what the context-stripped data contains. It cannot find the structural relationships that require ontogenic context to be visible, because the context was removed before the model saw the data.

The deeper failure is structural, not technical. AI systems trained on historical data reason about the world's past configuration. They cannot detect structural reorganizations before those reorganizations appear in the training record - reorganizations are, by definition, departures from historical patterns. A model trained before 2020 had no basis for anticipating the pharmaceutical supply chain vulnerabilities exposed by COVID. This is not because the data was insufficient. It is because the structural concentration of pharma manufacturing in China was a feature of the network's topology, not of its historical output statistics. The information ontogeny framework reads network topology as a primary feature. AI pattern recognition reads historical output statistics. These are not competing approaches to the same information. They are accessing different information entirely.

6.2 What Domain Expertise Cannot Scale

Deep domain expertise provides irreplaceable mechanism understanding. A specialist in country-specific pharmaceutical regulation understands the institutional dynamics of that system with depth and nuance no algorithm approaches. But the same expert cannot, working within that single domain, systematically trace the ontogenic transfer pathway from Bay of Bengal sea surface temperatures through monsoon dynamics to Hyderabad water availability to global generics supply. That cross-system, cross-scale developmental chain is not accessible through depth of expertise in any single information system. It requires the ability to trace the chain across boundaries - which requires a framework that treats information systems and their couplings as the primary object of analysis.

Domain expertise also carries the limitation of its interpretive framework. The pharmaceutical supply chain expert reads data through pharmaceutical supply chain categories. The climate scientist reads data through atmospheric dynamics categories. Neither, working independently, has the framework to recognize that the same structural coupling mechanism that produces El Niño teleconnections in the climate system produces analogous teleconnections in supply chains and financial networks. The information ontogeny framework provides the cross-domain scaffolding - a common language of generative systems, ontogenic transfer pathways, and developmental timescales - that makes systematic cross-system analysis tractable. Knowledge is created at the boundaries.

6.3 What the Ontogenic Lens Specifically Reveals

Three categories of finding are structurally accessible to the information ontogeny framework - across both the spatial and acoustic lenses - and inaccessible to the alternatives.

Self-organized structural concentrations of risk. These hubs are visible from network topology before they manifest as disruption. The geographic concentration of supply chains and manufacturing is a self-organized feature of the upstream/downstream information network, readable from its topology. This is reading the ontogenic history of the industrial information systems as they operate - not their historical outputs. In the acoustic lens, the same structural concentration is read on the acoustic-coverage side: which infrastructure is instrumented, which is not, and where acoustic monitoring inverts the topology of a physical asset into a topology of legibility.

Teleconnective chains with mechanistic grounding. These provide lead times that are structural rather than correlative. The developmental chain from Bay of Bengal sea surface temperatures through Indian monsoon dynamics to end-user manufacturing is a structurally grounded ontogenic transfer pathway with a characteristic developmental timescale of 6 to 18 months, depending on the stage. The acoustic analogue is the chain from a generative event - an eruption, an ocean storm - through atmospheric or aquatic waveguides to a sensor network. Lead time is not a statistical lag estimated from historical data. It is the developmental velocity of a specific coupling mechanism, and it is therefore durable across system reorganizations in ways that pure correlations are not.

Regime-change signals before disruption. Visible in spatial pattern changes that precede shifts in system output statistics, and in soundscape and acoustic-spectrum changes that precede shifts in environmental output statistics. Complex systems reorganize before they produce anomalous outputs. The reorganization is detectable in changes to the spatial pattern of information flows: shifts in the geographic distribution of manufacturing activity, changes in trade network topology, anomalies in the spatial correlation structure of climate and agricultural signals that precede regime transitions. In the acoustic lens, the analogous signals manifest as shifts in soundscape composition before species turnover, as changes in microbarom characteristics before remote ocean wave climate shifts, and as anomalies in industrial acoustic signatures before equipment failures. These are signals in the ontogenic substrate of the information systems - changes in how information is being produced, not yet in what it is saying.

7. Conclusion: How to Know the World

“The most important question is not ‘what does the data say?’ It is ‘where was this data born, and what has happened to it since?’”

— Michael Ferrari

The argument of this paper is a sequence of nested claims, each following from the one before. Information has a developmental history - an ontogeny - that is constitutive of its meaning, not separable from it. That ontogeny functions as infrastructure. It determines what is structurally knowable from a given data asset. The world is a system of interacting information systems, each with its own ontogenic dynamics, coupled through transfer pathways that are spatially organized and structurally characterizable. For situated data, space is the primary key to reading an information asset's developmental history, because a spatial coordinate encodes the densest available concentration of generative context. For acoustic data, the source-propagation-receiver geometry plays the same role. A methodology that preserves and analyzes information ontogeny - across whichever sensing modalities the physical environment offers - surfaces structural relationships that neither AI pattern recognition nor domain expertise can find, because both operate on ontogenically stripped representations.

The framework has independent validation from a tradition that reached a similar conclusion from a completely different starting point. Fei-Fei Li's work at World Labs instantiates the perception-reasoning dependency at the AI engineering level. Reasoning, Li argues, is a function of perceptual grounding, and perceptual grounding requires the structural context of the physical world to be preserved in the representations the reasoning system operates on. The sim-to-real gap is the engineering expression of the same failure mode this paper diagnoses theoretically: a system trained on ontogenically stripped information cannot reason correctly about a world whose information carries a different developmental history.

Two traditions converging on the same structural claim is not coincidence. I view it as evidence that the claim is identifying something real about how understanding actually works. That understanding requires grounding in the generative structure of the world, not only in the statistical patterns of its historical outputs. That the developmental history of information is not a curiosity for philosophers, but a practical determinant of what analytical systems can and cannot know. That the difference between an analysis that preserves ontogenic context and one that strips it is not a difference of degree but of kind.

The most consequential analytical failures of the past decade were not failures of data quantity, computational power, or domain expertise. They were failures of ontogenic grounding. The structural conditions that made past failures foreseeable were encoded in the developmental history of the relevant information systems - in their network topology, in their coupling mechanisms, in the spatial pattern of their generative concentrations, and increasingly, in the acoustic signatures of their activity. That history

was not analyzed, because the analytical frameworks in use treated the data as static resources rather than as developmental trajectories.

The most important structural relationships do not respect disciplinary boundaries. They operate through the ontogenic transfer pathways that connect information systems across those boundaries. The food-water-energy-infrastructure nexus is a zone of dense inter-system coupling; information born in the climate system develops into agricultural meaning, then into energy meaning, then into infrastructure risk, then into financial exposure, through a chain of ontogenic transformations that no single discipline can trace in its entirety. The acoustic dimension reinforces this conclusion. The continuous microbarom signal records a global ocean wave climate that no land-based sensor can produce in isolation. The infrasound trace of a distant eruption records the state of the stratospheric medium it traversed as plainly as it records the eruption itself. The DAS reading on a transoceanic cable records seismic, oceanographic, and anthropogenic activity simultaneously, through a sensor that was built for something else entirely.

Information Ontogeny as Infrastructure is not a model. It is a way of knowing. A commitment to reading the developmental history of data as primary, and to understanding the world as a system of information systems whose couplings are the primary object of analysis. The commitment now extends across both the spatial and acoustic lenses, and in principle across whatever further sensing modalities the physical environment offers as the perception layer of our analytical infrastructure continues to mature.

It is, I believe, the right way to know the world we are actually in.

As always, preparation trumps prediction.

Selected References

Primary intellectual anchors for the arguments developed in this paper. Section 5 of this Version 5 adds an acoustic-references block. A full bibliography accompanies the journal submission.

- Oyama, S. (1985 / 2000). *The Ontogeny of Information: Developmental Systems and Evolution*. Duke UP.
Foundational intellectual origin. Information as developmental product of organism-environment interaction. The direct predecessor of this paper's central claim.
- Li, F.-F. (2025). *From Words to Worlds: Spatial Intelligence is AI's Next Frontier*. Substack / World Labs Manifesto.
Convergent validation from AI: reasoning is constitutively dependent on perceptual grounding.
- Li, F.-F. (2025). *Spatial Intelligence Is AI's Next Frontier*. TIME.
Spatial intelligence as the scaffolding of cognition; the perception-reasoning dependency; AGI without spatial intelligence is incomplete.
- Shannon, C. & Weaver, W. (1949). *The Mathematical Theory of Communication*. University of Illinois Press.
Information theory foundations. This framework extends Shannon by insisting that context is constitutive of meaning, not separable from the signal.
- Massey, D. (2005). *For Space*. Sage Publications.
Space as relational rather than absolute. Ontological foundation for the spatial dimension of information ontogeny.
- Meadows, D. (2008). *Thinking in Systems*. Chelsea Green Publishing.
Systems dynamics; feedback loops; the primacy of relational structure over component properties.
- Batty, M. (2013). *The New Science of Cities*. MIT Press.
Space of flows; network science applied to cities as information systems.
- Liu, J. et al. (2013). Framing sustainability in a telecoupled world. *Ecology and Society*, 18(2).
Telecoupling framework: sending systems, receiving systems, spillover systems, coupling mechanisms.
- Boers, N. et al. (2019). Complex networks reveal global pattern of extreme-rainfall teleconnections. *Nature*, 566.
Network methods for identifying teleconnections from structural coupling rather than statistical correlation.
- Holland, J. (1998). *Emergence: From Chaos to Order*. Addison-Wesley.
Emergence in complex adaptive systems. Foundation for the structural concentration of risk argument.
- West, G. (2017). *Scale*. Penguin Press.
Universal scaling laws across biological, urban, and economic systems.
- USP Medicine Supply Map (2024–25). Geographic concentration of API manufacturing.
Primary empirical anchor: pharmaceutical supply chain spatial concentration as self-organized structural risk.
- Chester, M. & Allenby, B. (2024). *Infrasound first principles for the Anthropocene*.
Decoupling between traditional infrastructure development and the modern needs of today's increasingly complex environments.
- Hidalgo, C. (2015). *Why Information Grows: The Evolution of Order, from Atoms to Economies*.
Information lens applied to the origins of physical order and economic growth.
- §5 - Schafer, R. M. (1977). *The Soundscape: Our Sonic Environment and the Tuning of the World*. Destiny Books.
Foundational text in environmental acoustics. Intellectual ancestor of the acoustic-ontogeny argument.
- §5 - Pichon, A. Le, Blanc, E., & Hauchecorne, A. (eds.) (2019). *Infrasound Monitoring for Atmospheric Studies* (2nd ed.). Springer.

Definitive reference for the physics, propagation, and analysis of infrasound. Technical backbone of §5.2.

- §5 - Matoza, R. S. et al. (2022). Atmospheric waves and global seismoacoustic observations of the January 2022 Hunga eruption, Tonga. *Science* 377, 95-100.
The case study used in §5.3. Documents the global propagation of the Hunga Tonga acoustic signal.
- §5 - Donn, W. L. & Naini, B. (1973). Sea wave origin of microbaroms and microseisms. *Journal of Geophysical Research* 78, 4482-4488.
Classical mechanism paper for microbaroms as the atmospheric signature of ocean-wave nonlinear interactions.
- §5 - Zhan, Z. (2020). Distributed Acoustic Sensing turns fiber-optic cables into sensitive seismic antennas. *Seismological Research Letters* 91, 1-15.
Review of DAS principles and applications.
- §5 - Williams, E. F. et al. (2019). Distributed sensing of microseisms and teleseisms with submarine dark fibers. *Nature Communications* 10, 5778.
Demonstrates DAS detection on submarine cables - proof of cross-system ontogenic transfer through engineered fiber.
- §5 - Pijanowski, B. C. et al. (2011). Soundscape ecology: The science of sound in the landscape. *BioScience* 61, 203-216.
Founding statement of soundscape ecology. Theoretical basis for the §5.3 bioacoustic regime-change argument.
- §5 - Stowell, D. (2022). Computational bioacoustics with deep learning: a review and roadmap. *PeerJ* 10:e13152.
Modern review of acoustic ML applied to bioacoustic monitoring. AI-on-acoustics anchor.
- §5 - Lindsey, N. J. & Martin, E. R. (2021). Fiber-optic seismology. *Annual Review of Earth and Planetary Sciences* 49, 309-336.
Comprehensive review of fiber-optic sensing including DAS for seismic and acoustic applications.
- §5 - Ari Le Pichon et al. (2024). Climate change and atmospheric infrasound: monitoring stratospheric circulation through microbaroms. *Geophysical Research Letters* 51.
Recent work tying microbarom characteristics to stratospheric state - explicit demonstration that the acoustic signal carries developmental information about its propagation medium.

This working paper is an evolving document. The author welcomes critical engagement, methodological challenge, collaboration and productive disagreement, which is how understanding advances at the boundaries of ideas.

Atlas Research Innovations · Moby · AlphaGeo · atlasresearchinnovations.com