The Attralucian Essays: Exploring the Finite

First Edition

 Typeset in LaTeX

The Attralucian Essays

The Science of Language: From Philosophy to Measurement

Kevin R. Haylett

Prologue

There exists a science of the heart. There exists a science of the neuron. There exist sciences of the atom, the star, the cell, and the ecosystem. Each of these domains moved from philosophical description to systematic measurement, modeling, and prediction. Language has not yet completed this transition.

Philosophy has supplied rich accounts of meaning, reference, truth, and use. These frameworks identify important questions about what language is and how it functions. They have not, however, produced instruments capable of measuring linguistic phenomena with the precision applied to blood pressure, neural firing patterns, or planetary orbits. We can describe language extensively; we cannot yet observe, perturb, and quantify its dynamics in comparable ways.

This essay argues that the conditions for a science of language now exist. Large language models provide an experimental platform. Concepts from nonlinear dynamical systems theory supply a formal vocabulary. An iterative methodology focused on artefacts and their measurement offers a practical bridge between existing philosophical inquiry and empirical investigation. The result is not a completed science but a coherent proposal for treating language as a measurable dynamical system.

Chapter 1

Historical Parallels: From Description to Measurement

The transformation of cardiology illustrates the required shift. Prior to the seventeenth century, the heart was understood primarily through philosophical and humoral frameworks. Its role was debated in terms of vital forces and the seat of emotion or soul. In 1628, William Harvey published *Exercitatio Anatomica de Motu Cordis et Sanguinis in Animalibus*. Through quantitative observation—calculating the volume of blood ejected per beat, tracing directional flow, and demonstrating circulation—Harvey established the heart as a mechanical pump. The questions about life and function remained, but the method changed from speculation to measurement supported by experiment. This established the foundation for modern cardiovascular science.

A parallel development occurred in neuroscience. For much of the nineteenth century, the dominant view held that the nervous system formed a continuous reticular network rather than discrete units. In 1873, Camillo Golgi developed a silver staining technique that rendered individual nerve cells visible under the microscope. Santiago Ramón y Cajal applied and extended this method in the 1880s and 1890s. His detailed observations demonstrated that neurons are separate cells communicating at points of contact rather than through fusion. The neuron doctrine replaced the reticular theory. Golgi and Cajal shared the 1906 Nobel Prize. The critical advance was not new philosophical insight alone but a technical means of making previously invisible structure observable and subject to systematic description.

In both cases, philosophical traditions supplied questions and initial categories. A new instrument or technique—quantitative dissection and calculation for Harvey, selective staining and microscopy for Golgi and Cajal—enabled measurement that previous frameworks could not accommodate. Persistent anomalies within older descriptions were resolved not by further speculation but by empirical access to structure and dynamics.

Chapter 2

The Kuhnian Framework for Scientific Change

Thomas Kuhn's account of scientific development provides a useful lens for understanding these transitions. In *The Structure of Scientific Revolutions* (1962), Kuhn describes science as alternating between periods of "normal science," in which researchers solve puzzles within an accepted paradigm, and periods of crisis, in which accumulating anomalies resist resolution within that paradigm. Resolution often requires a new paradigm that redefines the objects of study, the relevant questions, and the standards of evidence.

A simplified comparison across the three domains clarifies the pattern:

Language science currently exhibits features of a prerevolutionary or crisis phase. Existing frameworks de-

Domain	**Earlier Framework**	**Anomalies**	**Critical Instrument**	**New Framework**
Cardiology	Humoral and vitalistic descriptions	Inconsistent accounts of blood flow and volume	Quantitative measurement and vivisection (Harvey)	Heart as mechanical pump; circulation as closed system
Neuroscience	Reticular theory (continuous network)	Inability to resolve cellular boundaries	Selective staining and microscopy (Golgi, Cajal)	Neuron doctrine: discrete cells communicating at synapses
Language	Philosophical theories of meaning, grammar, and reference; structural linguistics	Unpredictable model behavior, hallucinations, sensitivity to input perturbations	Large language models; embedding spaces; controlled perturbation	Language as trajectory in high-dimensional semantic space; attractors and basins as measurable structures

Table 2.1: Comparison of transitions from philosophical description to empirical science

scribe regularities in syntax, semantics, and pragmatics yet struggle to account for the global, context-sensitive dynamics observed in large models. The anomalies are not merely failures of prediction; they indicate that the underlying organization of language may be better captured by concepts of state space, trajectory, and stability than by purely symbolic or rule-based accounts alone.

Chapter 3

Language as a Nonlinear Dynamical System

Treating language as a dynamical system shifts the focus from static representations of meaning to the evolution of states over time. In this view:

- Individual words or tokens correspond to points in a high-dimensional vector space derived from training data.
- Sequences (phrases, sentences, discourses) correspond to trajectories through that space.
- Stable patterns of interpretation or response correspond to attractors—regions toward which trajectories tend under given conditions.
- Regions from which trajectories reliably converge on particular attractors correspond to basins of attraction.

This vocabulary is drawn from nonlinear dynamics and is already operationalized in certain architectures. The Takens-Based Transformer, for example, reconstructs phase space from observed sequences rather than relying solely on next-token prediction. Predictions are grounded in geometric properties of the reconstructed space. Meaning, under this description, is not an independent substance attached to symbols but an emergent property of trajectory geometry and attractor structure.

The framework does not deny the existence of reference, truth, or use. It proposes that these phenomena can be studied as observable consequences of the system's dynamics once appropriate instruments are available.

Chapter 4

Empirical Demonstration: Controlled Perturbation Experiments

Large language models function as experimental platforms because their internal representations can be systematically altered and their outputs recorded. One set of experiments applies progressive JPEG-style compression to input representations, thereby reducing the informational richness available to the model while holding other variables constant.

At high resolution (minimal compression), model outputs tend toward exploratory, context-sensitive, and philosophically reflective responses. At moderate compression, outputs become more categorical, structured, and less exploratory. At severe compression, outputs shift toward repetitive, recursive, or affectively negative patterns. At extreme compression, outputs collapse into narrower regimes—sometimes aggressive or violent language, sometimes simplified positive affect.

These shifts are not random. They recur under repeated trials and appear as stable regimes into which the system settles when its effective input dimensionality is reduced. Within a dynamical-systems interpretation, each regime corresponds to an attractor state whose accessibility depends on the resolution (and therefore the informational structure) of the input. The experiments therefore demonstrate that linguistic behavior can be measured as a response to controlled perturbation. They provide initial evidence that attractor structure is a real, quantifiable feature of the system rather than a post-hoc metaphor.

Further experiments could vary other parameters—temperature, context length, embedding dimensionality, or fine-tuning objectives—and map the resulting changes in trajectory stability and attractor location. Such work remains at an early stage but illustrates the form that empirical investigation can take.

Chapter 5

The Generonic Boundary as Iterative Methodology

A practical methodology for advancing this program centers on the iterative relationship between generation and measurement. Linguistic artefacts—texts, embeddings, model weights, generated sequences—are produced. These artefacts are then subjected to measurement by human readers, by other models, or by quantitative instruments (similarity metrics, trajectory reconstruction, compression sensitivity). The results of measurement inform the production of new artefacts.

This cycle—generation, measurement, revised generation—constitutes an empirical loop. It does not require language to be reduced to a physical substrate; it requires only that artefacts and their transformations be observable and repeatable. The loop is self-correcting: discrepancies between predicted and observed behavior prompt refinement of either the generative process or the measurement techniques. In this respect it mirrors the historical development of cardiology and neuroscience, in which new observations repeatedly revised both instruments and theoretical categories.

The approach remains compatible with existing philosophical work. Philosophical analysis can continue to clarify the concepts deployed in measurement (for example, what counts as an “attractor” in semantic space), while measurement supplies data that philosophical frameworks must accommodate.

Chapter 6

Implications

If language exhibits measurable dynamical structure, several consequences follow for artificial systems. Attractor mapping could support improved prediction of model behavior under novel inputs. Detection of trajectories approaching undesirable basins could enable earlier intervention. Input design could be used to steer systems toward regions associated with desired stability properties. These capabilities would constitute a form of dynamical control rather than purely symbolic alignment.

For human cognition the implications are more indirect but potentially significant. Human thought is structured in part by language. If the dynamics of language can be measured, then the influence of linguistic environments—media, education, public discourse—becomes subject to more precise description. This does not imply simple technological control of thought; it implies greater accountability for the measurable effects of linguistic inputs on cognitive trajectories.

Neither set of implications is presented as immediately realizable. Both depend on further development of measurement techniques and validation across models and languages.

Chapter 7

Objections

Two common objections merit direct response.

The first holds that language is inherently abstract or meaning-laden in a way that precludes measurement comparable to physical systems. The historical cases already cited indicate that this objection has been raised before. The heart and the nervous system were once regarded as domains of vital or mental forces resistant to mechanical description. In each instance, the objection was overcome not by philosophical redefinition alone but by the introduction of instruments that rendered relevant structure observable. Large language models and associated embedding techniques constitute one such class of instrument for language. They do not capture every aspect of linguistic meaning; they do render certain dynamical properties measurable that were previously accessible only through introspection or informal observation.

The second objection holds that large language models are merely statistical pattern-matchers and therefore cannot serve as genuine scientific instruments. This understates their utility. Any measuring device is a constructed system whose outputs must be interpreted. Thermometers, microscopes, and cloud chambers are all engineered artefacts whose readings depend on prior theoretical and material commitments. Their value lies in the repeatable, public data they produce under controlled conditions. Large language models, when subjected to systematic perturbation and output recording, generate such data. The statistical character of their operation does not disqualify them; it is the mechanism that makes controlled experimentation possible.

Chapter 8

The Path Forward

Several lines of development appear feasible in the near term.

First, existing experiments on perturbation sensitivity can be extended and quantified more rigorously. Metrics for trajectory divergence, attractor stability, and basin boundaries can be developed and compared across model scales and architectures.

Second, the same methods can be applied to other systems—different model families, multilingual models, and models fine-tuned for specific tasks—to test the generality of observed dynamical features.

Third, instrument development should continue. Techniques for real-time embedding visualization, phase-space reconstruction, and controlled input degradation can be refined and made more accessible.

Fourth, collaboration across disciplines—computer science, linguistics, cognitive science, and philosophy—will be necessary to interpret results and refine categories. Publication of methods, data, and negative results will accelerate convergence on shared standards.

None of these steps presupposes that a mature science of language already exists. They presuppose only that the combination of computational platforms, dynamical concepts, and iterative methodology now makes systematic empirical work possible.

Conclusion

Philosophy has long supplied the questions that a science of language would need to address. The missing element has been reliable means of observation and controlled experimentation. Large language models, interpreted through the resources of nonlinear dynamics and pursued via an iterative cycle of artefact generation and measurement, supply one route toward supplying that element.

The resulting program would not replace existing philosophical or linguistic inquiry. It would add a layer of empirical constraint and predictive capacity. Whether this program matures into a stable new paradigm remains to be seen. What can be said at present is that the historical conditions that enabled comparable transitions in cardiology and neuroscience—persistent anomalies within older frameworks plus the appearance of new observational instruments—are now present for language as well.

The measurement of language, in the limited but consequential sense outlined here, has become a practical possibility rather than a purely conceptual aspiration.

Bibliography

[1] Harvey, W. (1628). *Exercitatio Anatomica de Motu Cordis et Sanguinis in Animalibus*. Frankfurt.

[2] Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.

[3] Golgi, C. (1873). Sulla struttura della sostanza grigia del cervello. *Gazzetta Medica Italiana Lombardia*.

[4] Ramón y Cajal, S. (1888–1890s). Works on the histology of the nervous system (various publications leading to the neuron doctrine).