

Geodesics of Meaning: Modeling Semantic Curvature in Transformer-Based Language Models via General Relativity

Travis S. Taylor

QuantumFrontier, LLC, Huntsville, AL, USA

University of Alabama in Huntsville, Huntsville, AL, USA

tst0072@uah.edu

ORCID: 0009-0003-3209-5825

January 29, 2026

Abstract

We present a general relativistic framework for modeling transformer-based language models (LLMs) as nonlinear dynamical systems evolving on curved semantic manifolds. Standard transformer architectures are shown to approximate a flat Minkowski spacetime, where attention mechanisms define a local semantic metric tensor. We extend this formulation by introducing curved metrics—specifically the Schwarzschild and Friedmann–Lemaître–Robertson–Walker (FLRW) solutions—to model context-sensitive meaning, narrative curvature, and long-range semantic dependencies. A stress-energy tensor encodes topical mass, tonal flow, and tension, driving semantic curvature via Einstein’s field equations. We validate this framework using both simplified language simulations and full narrative data, showing that Ricci curvature serves as a physically interpretable measure of coherence, complexity, and twist. This work bridges differential geometry, nonlinear systems, and AI interpretability, offering a new paradigm for analyzing and guiding large language model behavior.

1 Introduction

Transformers have revolutionized NLP [1–4], but their token-by-token generation limits their ability to capture global context and long-range dependencies. Nonlinear systems theory and differential geometry have long offered tools for modeling structure and dynamics in complex systems. In this work, we apply general relativistic geometry—specifically the Ricci curvature and Einstein field equations—to construct a dynamical framework for semantic structure in language. Large language models (LLMs) are used as instantiations of a high-dimensional information system, where semantic evolution is traced via geodesics on a curved manifold. Drawing inspiration from GR, we propose a spacetime metric framework to address these limitations.

Key Contributions:

- Show that the standard transformer architecture fits a Minkowski-like spacetime, with word embeddings and attention mechanisms defining a flat semantic metric. This geometric interpretation aligns with previous work on manifold embeddings in language models [5–7].
- Introduce curved spacetime metrics (e.g., Schwarzschild, Friedmann) to model semantic curvature, context shifts, and hierarchical structures in text.
- Demonstrate how tensor calculus and differential geometry can enhance LLMs, enabling next-sentence and next-paragraph prediction.

This approach aligns with previous work in statistical biophysics and information geometry by Bialek [8–10], who applied geometric and thermodynamic principles to neural encoding and biological computation. His use of curvature-based metrics and manifold representations of sensory data anticipates many of the semantic-geometric dynamics explored here.

Throughout this work, general relativity is used as a mathematical and conceptual scaffold rather than a literal physical model; curvature, stress–energy, and geodesics are employed as structured analogs for semantic dynamics, not as claims about physical spacetime or exact solutions of Einstein’s equations.

2 Background

In this section, we provide an overview of the core principles behind transformer-based language models and general relativity (GR), establishing the mathematical and conceptual foundation for our proposed framework. We then bridge these two domains by highlighting structural analogies between language modeling and spacetime geometry.

2.1 Language Models as Nonlinear Trajectories in Curved Semantic Space

The transformer model, introduced by Vaswani et al. [1], revolutionized natural language processing (NLP) by leveraging self-attention mechanisms to capture dependencies between words across long sequences. Unlike traditional recurrent architectures, transformers process entire sequences in parallel, making them highly efficient and scalable.

2.1.1 Key Components of Transformers

A standard transformer consists of:

- **Word Embeddings:** Continuous vector representations of words or subwords in a high-dimensional space.
- **Positional Encoding:** Adds sequence order information to embeddings, typically through sinusoidal functions.

- **Self-Attention Mechanism:** Computes weighted relationships between words using the scaled dot-product attention formula:

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V \quad (1)$$

where Q , K , and V are the query, key, and value matrices, and d_k is the dimension of the key vectors.

- **Multi-Head Attention:** Extends self-attention by computing multiple attention scores in parallel:

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_h)W^O \quad (2)$$

where each attention head learns different aspects of the relationships in the input.

- **Feedforward Network:** Applies nonlinear transformations to the attention output.
- **Layer Normalization and Residual Connections:** Improve training stability and gradient flow.

2.1.2 Limitations of Transformer Models

Despite their success, transformers have key limitations:

- **Local Context Bias:** Although attention mechanisms allow for long-range dependencies, they do not inherently encode a global understanding of the text.
- **Lack of Hierarchical Coherence:** Sentence-by-sentence or token-by-token generation does not enforce large-scale coherence in longer passages.
- **Fixed Context Window:** Many transformer models are constrained by a fixed-length attention window, limiting their ability to model arbitrarily long documents.

To address these issues, we propose a spacetime metric framework that extends transformer architectures by introducing curvature and global coherence constraints.

2.2 General Relativity and Spacetime Metrics

General Relativity (GR) describes gravity as a manifestation of spacetime curvature, governed by Einstein’s field equations [11] [12] [13]. The field equations were first formulated in [11] and later refined in [13] [14]:

$$G_{\mu\nu} = \kappa T_{\mu\nu} \quad (3)$$

where $G_{\mu\nu}$ is the Einstein tensor (describing curvature), $T_{\mu\nu}$ is the stress-energy tensor (representing mass-energy and momentum), and κ is a proportionality constant.

Key Concepts in Spacetime Geometry

GR is built upon the following core mathematical structures:

- **Manifolds:** A smooth space where physical or abstract distances can be measured.
- **Metric Tensor $g_{\mu\nu}$:** Defines the infinitesimal distance in spacetime via the line element:

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu \quad (4)$$

- **Curvature:** Describes how spacetime bends in response to mass-energy. This is quantified using the Ricci tensor $R_{\mu\nu}$ and the full Riemann curvature tensor.
- **Geodesics:** The paths that free-falling particles (or in our case, coherent semantic progressions) follow, given by:

$$\frac{d^2 x^\mu}{d\tau^2} + \Gamma_{\alpha\beta}^\mu \frac{dx^\alpha}{d\tau} \frac{dx^\beta}{d\tau} = 0 \quad (5)$$

where $\Gamma_{\alpha\beta}^\mu$ are the Christoffel symbols representing spacetime connection.

Minkowski Metric: Flat Semantic Spacetime

The Minkowski metric describes a flat, uncurved spacetime and is the simplest representation of a semantic space where meaning evolves in a uniform, linear manner. It is given by:

$$ds^2 = -dt^2 + dx_1^2 + dx_2^2 \quad (6)$$

where:

- t represents temporal progression in a text, such as the order of sentences or paragraphs.
- x_1 and x_2 denote semantic dimensions, mapping relationships between concepts.
- The absence of curvature terms indicates that meaning propagates in a straight-line fashion, without distortion or shifts.

Dimensionality of Semantic Spacetime

Throughout this work, we employ a minimal semantic spacetime with coordinates

$$x^\mu = (t, x_1, x_2), \quad (7)$$

corresponding to one temporal dimension (textual progression) and two effective semantic dimensions representing topic and tone. All metric tensors $g_{\mu\nu}$, stress-energy tensors $T_{\mu\nu}$, and curvature quantities are defined on this $(1 + 2)$ -dimensional manifold unless explicitly stated otherwise.

Higher-dimensional representations (e.g., full embedding spaces) are treated as underlying latent spaces whose dynamics are projected onto this reduced semantic spacetime. References to

four-dimensional Minkowski spacetime are used strictly as an analogy to emphasize flat versus curved semantic evolution, rather than as a literal physical dimensionality. The use of Schwarzschild and FLRW metrics in this work should be understood as an inspirational projection of their qualitative structure onto an effective (1+2)-dimensional semantic manifold, rather than as a literal lower-dimensional solution of Einstein's equations.

Interpretation in Textual Structure

A Minkowski-like semantic space corresponds to texts where:

- Ideas evolve smoothly, without abrupt thematic shifts or dramatic tonal changes.
- Logical progression is linear, such as in structured technical writing, procedural manuals, or mathematical proofs.
- No significant distortions are present in coherence or consistency—each sentence smoothly follows from the previous.

In this framework, self-attention in transformers can be interpreted as defining an implicit Minkowski space, where token embeddings evolve under a locally flat metric. This means that, in the absence of external perturbations (e.g., major topic shifts), the transformer preserves local semantic relationships in a manner analogous to inertial motion in special relativity.

Breakdown of the Minkowski Assumption

In real-world text, purely Minkowski-like behavior is rare because meaning is rarely uniform. Deviation from a flat metric introduces semantic curvature, representing:

- Sudden shifts in context or subject matter.
- Emotional intensity or tonal shifts disrupting a steady flow.
- Hierarchy of ideas, where primary and supporting concepts introduce a more complex structure.

Thus, while transformers approximate a Minkowski-like embedding space, real-world texts exhibit deviations that require a more advanced curved spacetime formulation.

—

Schwarzschild Metric: Semantic Gravity and Event Horizons

The Schwarzschild metric describes a spherically symmetric, curved spacetime around a massive object. In our framework, it models topics that exert strong "semantic gravity," pulling related ideas toward them. The metric is given by:

$$ds^2 = - \left(1 - \frac{2GM}{r}\right) dt^2 + \left(1 - \frac{2GM}{r}\right)^{-1} dr^2 + r^2 d\Omega^2 \quad (8)$$

where:

- M represents semantic mass, quantifying the strength of a central concept in a text.
- r measures semantic distance from this central idea.
- The term $1 - \frac{2GM}{r}$ governs how meaning is distorted near high-impact topics.

Interpretation in Textual Structure

A Schwarzschild-like metric is useful for understanding:

- Central Themes in Texts: Some ideas have high "semantic mass"—they dominate discourse and shape meaning around them. Examples include:
 - Political or philosophical texts where a core ideology attracts surrounding arguments.
 - Scientific papers where a single groundbreaking idea anchors all supporting discussion.
 - Novels centered around a single protagonist or conflict that defines the entire narrative.
- Conceptual "Gravity Wells": Some ideas create a semantic gravitational pull, attracting related ideas while making it difficult to escape into unrelated discussions.
- Hierarchical Information Structures: Certain texts have a strong core argument, with supporting details orbiting it—akin to planetary orbits around a central mass.

Event Horizons: Loss of Meaning and Comprehensibility

A defining feature of the Schwarzschild metric is the event horizon at $r = 2GM$. In our model, this represents a "semantic black hole" where comprehension is lost:

- Singularities in Meaning: A concept may become too dense or complex, leading to an overwhelming information load that prevents clarity.
- Obscure or Ambiguous Texts: Incoherent, overly convoluted, or cryptic writing styles create regions where meaning collapses—akin to a text falling beyond an event horizon.
- Unrecoverable Context Shifts: If a narrative or argument becomes too entangled, the reader may lose track of earlier meanings, making escape from the semantic black hole impossible.

For example:

- Philosophical works with extreme abstraction may form semantic singularities, where meaning is difficult to extract.
- Poorly structured narratives with excessive flashbacks or nonlinear structures may result in event horizons where coherence is lost.
- High-context writing, such as legal documents or advanced mathematics, can trap comprehension unless one possesses the necessary "escape velocity" of background knowledge.

Geodesic Interpretation: How Ideas "Fall" into Meaning Wells

Particles in a Schwarzschild spacetime follow geodesics, which describe their natural motion in a gravitational field. Similarly, concepts in a text naturally follow semantic geodesics determined by the meaning landscape. The geodesic equation:

$$\frac{d^2 x^\mu}{d\tau^2} + \Gamma_{\alpha\beta}^\mu \frac{dx^\alpha}{d\tau} \frac{dx^\beta}{d\tau} = 0 \quad (9)$$

predicts how ideas progress:

- A text with a strong central theme keeps ideas "bound" to it.
- High-density concepts warp meaning around them, drawing connections inward.
- If a topic becomes too dense, meaning collapses into a black hole of incomprehensibility.

Friedmann-Lemaître-Robertson-Walker (FLRW) Metric: Dynamic Narrative Evolution

The FLRW metric is widely used in cosmology to model an expanding or contracting universe [15]. In our framework, we reinterpret it as a mathematical representation of how a narrative evolves over time. The metric is given by:

$$ds^2 = -dt^2 + a(t)^2 (dr^2 + r^2 d\Omega^2) \quad (10)$$

where:

- $a(t)$ is the **scale factor**, which governs how the semantic space of the text expands or contracts over time.
- $dr^2 + r^2 d\Omega^2$ represents spatial relationships within the text, mapping how different ideas or themes relate to one another.
- dt^2 tracks sequential progression, ensuring causality in textual evolution.

Interpretation of $a(t)$ in a Narrative Context

The behavior of $a(t)$ provides insight into the overall complexity of a narrative:

- **Expanding Narratives** ($\dot{a}(t) > 0$): The introduction of new ideas, characters, and themes increases the complexity of the story. This is characteristic of epic novels or research papers that progressively introduce new concepts.
- **Contracting Narratives** ($\dot{a}(t) < 0$): The story narrows in focus, resolving conflicts, simplifying themes, and converging towards a conclusion. This is seen in detective novels as clues and storylines merge into a final resolution.
- **Static Narratives** ($\dot{a}(t) \approx 0$): A text maintains a steady level of complexity, typical of structured expository writing such as textbooks or legal documents.

Higher-Order Narrative Evolution: Acceleration and Oscillations

In some cases, the rate of change of $a(t)$ itself may evolve dynamically:

- **Accelerating Complexity** ($\ddot{a}(t) > 0$): The text’s thematic or conceptual expansion speeds up, analogous to the universe’s accelerated expansion due to dark energy. This models chaotic or nonlinear storytelling.
- **Cyclic Narratives** ($a(t)$ oscillates): Represents episodic structures where stories expand and contract periodically, such as TV shows with self-contained story arcs.

Geodesic Interpretation: Tracing Narrative Paths

In FLRW spacetime, particles (or in our case, coherent text structures) follow geodesics given by [15]:

$$\frac{d^2 x^\mu}{d\tau^2} + \Gamma_{\alpha\beta}^\mu \frac{dx^\alpha}{d\tau} \frac{dx^\beta}{d\tau} = 0 \quad (11)$$

where $\Gamma_{\alpha\beta}^\mu$ are the Christoffel symbols. These geodesics represent optimal semantic trajectories, allowing us to quantify coherence and logical flow within a text.

2.3 Bridging Transformers and General Relativity

Transformers and GR both describe evolving systems where distance and interaction strengths determine system behavior. We make the following analogies:

- **Word embeddings as points in spacetime:** Each token exists in a high-dimensional manifold of meaning.
- **Attention mechanisms as metric tensors:** The degree to which one word influences another is analogous to a metric defining distances between spacetime events.
- **Semantic curvature as topic evolution:** A text with a sudden shift in meaning exhibits nonzero curvature, analogous to a gravitational field bending light.
- **Coherent text as geodesics:** A well-structured argument follows a geodesic through semantic space, minimizing deviation from logical progression.

Related geometric interpretations of transformer architectures have recently been proposed, including models in which attention induces curvature in embedding space and token trajectories follow geodesic paths; the present work extends this line of thought by focusing on explicit curvature proxies tied to narrative coherence and controllable semantic dynamics [16].

Mathematical Reformulation

If we define a text's semantic structure via a metric tensor $g_{\mu\nu}$, then deviations from linear coherence can be measured using the Ricci curvature tensor:

$$R_{\mu\nu} = \partial_\alpha \Gamma_{\mu\nu}^\alpha - \partial_\nu \Gamma_{\mu\alpha}^\alpha + \Gamma_{\mu\lambda}^\alpha \Gamma_{\nu\alpha}^\lambda - \Gamma_{\nu\lambda}^\alpha \Gamma_{\mu\alpha}^\lambda \quad (12)$$

where $\Gamma_{\mu\nu}^\alpha$ are the Christoffel symbols, given by:

$$\Gamma_{\mu\nu}^\alpha = \frac{1}{2} g^{\alpha\lambda} (\partial_\mu g_{\lambda\nu} + \partial_\nu g_{\lambda\mu} - \partial_\lambda g_{\mu\nu}) \quad (13)$$

Regions of high curvature correspond to complex, nontrivial semantic shifts, where meaning deviates significantly from a flat, linear text progression.

Einstein's Field Equations for Semantic Space

To fully capture the geometric structure of meaning propagation, we employ the full Einstein field equations:

$$R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} + \Lambda g_{\mu\nu} = \kappa T_{\mu\nu} \quad (14)$$

where:

- $R_{\mu\nu}$ is the Ricci curvature tensor, encoding local distortions in meaning.
- R is the Ricci scalar, defined as $R = g^{\mu\nu} R_{\mu\nu}$, representing global coherence distortions.
- $g_{\mu\nu}$ is the semantic metric tensor, determining how meaning is structured across the text.
- Λ is the cosmological constant, interpreted here as a "dark energy" effect in textual evolution.
- $T_{\mu\nu}$ is the stress-energy tensor, representing the "mass-energy" of meaning, topic flow, and tonal pressure.
- κ is a proportionality constant.

Interpreting the Cosmological Constant Λ in Textual Evolution

The inclusion of $\Lambda g_{\mu\nu}$ introduces an additional force that can expand or contract the semantic structure of a text. This corresponds to:

- Positive Λ (Semantic Expansion) - Represents texts that introduce new concepts indefinitely. - Example: Expansive world-building in science fiction, research papers that continuously introduce novel theories.
- Negative Λ (Semantic Contraction) - Represents texts that collapse toward a singular core theme. - Example: Philosophical treatises that refine broad arguments into a single, fundamental insight.

- Zero Λ (Balanced Semantic Structure) - Indicates a text that maintains a steady rate of meaning evolution. - Example: Well-structured textbooks, where information is introduced at a steady, predictable pace.

Geodesics and Optimal Meaning Flow

In this curved semantic spacetime, meaning flows naturally along geodesics, governed by:

$$\frac{d^2 x^\mu}{d\tau^2} + \Gamma_{\alpha\beta}^\mu \frac{dx^\alpha}{d\tau} \frac{dx^\beta}{d\tau} = 0 \quad (15)$$

where x^μ represents the semantic coordinates of meaning, and τ is an affine parameter tracking text progression.

- Strongly curved geodesics indicate texts with dense meaning structures, where concepts are highly interconnected.
- Weakly curved or flat geodesics correspond to linear, straightforward texts with little deviation in meaning.
- Geodesic deviation equations could quantify how different sections of a text evolve over time, measuring coherence loss or gain.

Summary of the Extended Formalism

- The Ricci tensor measures local meaning distortion, such as topic shifts. - The Ricci scalar captures global coherence, identifying whether meaning remains structured or becomes chaotic. - The cosmological constant Λ models the long-term expansion or contraction of meaning. - Einstein's equations define the relationship between meaning structure (curvature) and its content (stress-energy tensor). - Geodesics trace the optimal flow of meaning, helping identify narrative consistency.

This extended mathematical formalism allows us to model and quantify how meaning propagates through text, distinguishing between stable, expanding, and collapsing narrative structures.

3 The Semantic Stress-Energy Tensor $T_{\mu\nu}$

In General Relativity, the stress-energy tensor $T_{\mu\nu}$ encodes the distribution of energy, momentum, and stress that shape spacetime curvature. In the context of semantic spacetime, we define an analogous tensor that represents the distribution of meaning, coherence, and topic interactions within a text. This builds on the foundational principles of information geometry [17, 18] and curvature-driven learning which allows us to implement the Einstein field equations in the form:

$$G_{\mu\nu} = \kappa T_{\mu\nu} \quad (16)$$

where $G_{\mu\nu}$ represents the semantic curvature (the "shape" of the text's logical and thematic flow), and $T_{\mu\nu}$ represents the mass-energy, pressures, and tensions within the semantic structure.

3.1 Defining $T_{\mu\nu}$ for Textual Evolution

We define the semantic stress-energy tensor in a structured form:

$$T_{\mu\nu} = \begin{bmatrix} \text{Semantic Mass} & \text{Topic Flow} & \text{Tonal Flow} \\ \text{Topic Flow} & \text{Topical Pressure} & \text{Tension/Twist} \\ \text{Tonal Flow} & \text{Tension/Twist} & \text{Tonal Pressure} \end{bmatrix} \quad (17)$$

where:

- T_{00} (Semantic Mass): Measures the informational density of a passage. High values indicate dense, complex sections, such as technical descriptions or philosophical arguments.
- $T_{01} = T_{10}$ (Topic Flow): Represents how smoothly a topic evolves over the text. Low values indicate rigid structure, while high values suggest fluid, evolving topics.
- $T_{02} = T_{20}$ (Tonal Flow): Measures emotional continuity—whether a text maintains a consistent tone or fluctuates.
- T_{11} (Topical Pressure): Captures how much a text resists topic change. High values indicate texts that stay highly focused on a central theme.
- $T_{12} = T_{21}$ (Tension/Twist Factor): Quantifies interplay between topic and tone, representing shifts in emphasis or perspective.
- T_{22} (Tonal Pressure): Measures the intensity of emotional or tonal content in the passage. High values correspond to dramatic or poetic text.

Each component of the semantic stress-energy tensor $T_{\mu\nu}$ serves as an analog to physical quantities in general relativity, adapted to represent structural properties of language. Specifically, T_{00} quantifies semantic mass—i.e., the informational density or conceptual weight of a passage. Off-diagonal terms like T_{01} and T_{02} measure flow of topic and tone over time, while T_{12} represents semantic twist: the nonlinear coupling between what is being said and how it is expressed. Diagonal terms such as T_{11} and T_{22} act as semantic pressure components, resisting or amplifying deviations in theme or emotional tone. This formulation allows the Einstein field equations to be reinterpreted as constraints on the internal geometry of narrative evolution.

3.2 Example Calculation of $T_{\mu\nu}$

For a highly structured, information-dense text with moderate tonal variation, we might obtain:

$$T_{\mu\nu} = \begin{bmatrix} 2 & 0.5 & 1 \\ 0.5 & 1 & 2 \\ 1 & 2 & 3 \end{bmatrix} \quad (18)$$

3.3 Interpreting $T_{\mu\nu}$ in Different Texts

The values of $T_{\mu\nu}$ change depending on the type of text being analyzed:

- Scientific Papers / Technical Texts:

$$T_{\mu\nu} = \begin{bmatrix} 5 & 0.2 & 0.1 \\ 0.2 & 4 & 0.3 \\ 0.1 & 0.3 & 1 \end{bmatrix} \quad (19)$$

- High semantic mass (dense content, formalized reasoning). - Low topic flow (tightly structured ideas). - Low tonal variations (objective and neutral).

- Novels with Strong Emotional Shifts:

$$T_{\mu\nu} = \begin{bmatrix} 2 & 1 & 2 \\ 1 & 2 & 3 \\ 2 & 3 & 5 \end{bmatrix} \quad (20)$$

- Lower semantic mass (narrative over factual density). - High tonal pressure (frequent emotional fluctuations). - High tension/twist factor (dramatic shifts in meaning).

- Poetry / Abstract Writing:

$$T_{\mu\nu} = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 4 \\ 3 & 4 & 6 \end{bmatrix} \quad (21)$$

- High tension/twist factor (meaning emerges from juxtaposition). - High tonal pressure (emotionally charged writing). - Moderate semantic mass (concise but layered meanings).

3.4 Linking $T_{\mu\nu}$ to Curvature and Meaning

Einstein's field equations describe how stress-energy generates curvature, which in our framework means:

- High T_{00} (semantic mass) bends the meaning structure, requiring more effort to "move" between ideas.
- High T_{12} (tension/twist) introduces instability, increasing complexity and unpredictability in text flow.

- High T_{11} (topical pressure) resists deviation, keeping discussions tightly bound to central themes.

Thus, solving $G_{\mu\nu} = \kappa T_{\mu\nu}$ reveals how meaning curves through a document, allowing us to predict where structural shifts, topic transitions, or tonal shifts will occur.

3.5 Future Applications of the Semantic $T_{\mu\nu}$

Applying this tensor-based formalism enables:

- Automated narrative analysis, where the curvature of a document predicts its logical coherence.
- Dynamic AI text generation, ensuring global consistency in long-form writing.
- Compression of meaning structures, using Ricci scalar analysis to extract core concepts.
- Identification of singularities and event horizons, flagging passages where meaning collapses or becomes ambiguous.

The semantic stress-energy tensor $T_{\mu\nu}$ formalizes meaning as a measurable, structured entity. Different components encode semantic density, topic evolution, tonal variation, and tension. Solving $G_{\mu\nu} = \kappa T_{\mu\nu}$ maps how a text deforms under meaning pressure. This approach bridges AI text processing with differential geometry, enabling a new class of structured, explainable language models.

4 Constructing a Small Language Model (SLM) in Semantic Space-time

To better understand the impact of semantic curvature and stress-energy interactions in text generation, we construct a Small Language Model (SLM) with a restricted vocabulary inspired by early-grade reader texts. This allows us to experiment with different stress-energy configurations $T_{\mu\nu}$ and study how they influence text structure and progression.

4.1 Step 1: First-Grade Reader Vocabulary

We define a basic lexicon (80 words) categorized by topics (x_1) and tone/emotion (x_2):

4.1.1 Nouns (Topics, x_1 Component)

- **People:** Dick, Jane, Spot, Baby, kids
- **Objects:** ball, balls, book, books, car, cars, house, tree, trees
- **Places:** yard, hill, park

4.1.2 Verbs (Actions, x_1 Shift)

- **Motion:** run, runs, ran, jump, jumps, jumped, walk, walks, walked, go, goes, went, fall, falls, fell
- **Interaction:** see, sees, saw, look, looks, looked, play, plays, played, catch, catches, caught, call, calls, called
- **Noise:** bark, barks, barked, yell, yells, yelled

4.1.3 Adjectives/Adverbs (Tone, x_2 Component)

- **Neutral:** big, little, good, here, there
- **Positive:** happy, fast, fun, funny, loud
- **Negative:** sad, slow, bad, quiet

4.1.4 Connectors (Syntax)

- **Basic:** and, with, in, on, to, at
- **Temporal:** now, then

This forms a structured dataset that allows us to map semantic space to words.

4.2 Step 2: SLM Simulation Framework

We simulate the SLM by:

1. **Assigning each word** a 2D coordinate in (x_1, x_2) -space.
2. **Defining $T_{\mu\nu}$** to represent narrative dynamics.
3. **Computing geodesics** in the perturbed metric $g_{\mu\nu}$.
4. **Selecting words** closest to the geodesic path.
5. **Tracking Ricci curvature R** to measure text complexity.

4.2.1 Coordinate Assignment

We define a numerical representation for the vocabulary:

- x_1 (Topic): Ranges from 0 to 15 (e.g., Dick = 0, run = 6, bark = 12).
- x_2 (Tone): Ranges from -5 to 5 (e.g., sad = -2, fast = 3, loud = 4).

Example Assignments:

- "Dick" \rightarrow (0, 0), "runs" \rightarrow (6.1, 0), "fast" \rightarrow (6.1, 3), "barked" \rightarrow (12.2, 1).

4.3 Step 3: Geodesic Computation in Perturbed Semantic Metrics

4.3.1 Mathematical Model

Given an initial sentence S_0 at $t = 0$, we compute the geodesic as a function of $T_{\mu\nu}$:

$$\frac{d^2 x^\mu}{d\tau^2} + \Gamma_{\alpha\beta}^\mu \frac{dx^\alpha}{d\tau} \frac{dx^\beta}{d\tau} = 0 \quad (22)$$

where:

- $T_{\mu\nu}$ defines how topics and tones evolve.
- $g_{\mu\nu}$ is perturbed based on $T_{\mu\nu}$.
- $x_1(t), x_2(t)$ track the progression of meaning.

We compare two cases:

4.4 Step 4: Flat vs. Curved Narrative Structures

For the toy SLM examples, we adopt a Minkowski baseline metric $\eta_{\mu\nu} = \text{diag}(-1, 1, 1)$, with semantic curvature introduced via small perturbations driven by $T_{\mu\nu}$.

We simulate these two paragraphs with different perturbations via the stress-energy $T_{\mu\nu}$.

4.4.1 Flat Semantic Space: Low Curvature

Tensor:

$$T_{\mu\nu} = \begin{bmatrix} 1 & 0.2 & 0 \\ 0.2 & 0.5 & 0 \\ 0 & 0 & 0.5 \end{bmatrix} \quad (23)$$

Metric:

$$g_{\mu\nu} = \begin{bmatrix} -1 & 0.2 & 0 \\ 0.2 & 1.5 & 0 \\ 0 & 0 & 1.5 \end{bmatrix} \quad (24)$$

Generated Paragraph:

"Dick runs fast. Jane jumps fast. Spot runs fast."

Ricci Curvature: $R \approx 0$ (Flat, linear flow).

4.4.2 Curved Semantic Space: High Curvature

Tensor:

$$T_{\mu\nu} = \begin{bmatrix} 2 & 0.5 & 1 \\ 0.5 & 1 & 2 \\ 1 & 2 & 3 \end{bmatrix} \quad (25)$$

Aspect	Flat Paragraph	Curved Paragraph
Generated Text	"Dick runs fast. Jane jumps fast. Spot runs fast."	"Dick runs fast. Spot barks loud. Jane yells now."
T_{00} (Semantic Mass)	1 (Low complexity)	2 (Higher complexity)
T_{01} (Topical Flow)	0.2 (Slow topic shift)	0.5 (Faster topic shift)
T_{02} (Tonal Flow)	0 (Static tone)	1 (Tone escalates)
T_{12} (Tension/Twist)	0 (No twist)	2 (High tension/twist)
R (Ricci Curvature)	≈ 0 (Flat)	> 0 (0.5, Curved)
Geodesic Path	Linear: x_1 slight rise, x_2 constant	Curved: x_1 rises, x_2 accelerates

Table 1: Comparison of flat and curved semantic structures using $T_{\mu\nu}$ perturbations.

Metric:

$$g_{\mu\nu} = \begin{bmatrix} -1 & 0.5 & 1 \\ 0.5 & 2 & 2 \\ 1 & 2 & 4 \end{bmatrix} \quad (26)$$

Generated Paragraph:

"Dick runs fast. Spot barks loud. Jane yells now."

Ricci Curvature: $R > 0$ (Tonal escalation and topic twists).

4.5 Comparison and Insights

As shown in Table 1, the flat semantic structure, characterized by a low-stress energy tensor $T_{\mu\nu}$, results in repetitive, structured sentences with minimal variation. In contrast, the curved semantic structure introduces shifts in meaning, tonal escalation, and greater emotional weight due to increased values in the tension and tonal flow components of $T_{\mu\nu}$. Additionally, higher Ricci curvature R correlates with greater complexity, increased narrative tension, and less predictable sentence evolution, demonstrating how perturbations in the semantic metric influence textual coherence and development.

4.6 Elucidation of Values and Nuances

The semantic stress-energy tensor $T_{\mu\nu}$ serves as a narrative driver, influencing the structure, complexity, and engagement level of generated text. Below, we analyze how different components of $T_{\mu\nu}$ affect paragraph structure, curvature, and narrative flow.

$T_{\mu\nu}$ as a Narrative Driver

The role of $T_{\mu\nu}$ in shaping text can be examined by comparing the flat and curved cases from Table 1.

Flat $T_{\mu\nu}$: Minimal Influence

- **Low** $T_{00} = 1$: Minimal "semantic mass" results in simple, repetitive sentence structures.
- $T_{01} = 0.2$, $T_{02} = 0$: Slow topical flow with no tonal change—each sentence remains close to the initial theme (e.g., "runs fast" persists).
- $T_{12} = 0$: No tension or twist, meaning topics (e.g., Dick, Jane, Spot) and tone (fast) do not interact dynamically.
- **Result:** A predictable, monotonous sequence resembling a Minkowski spacetime where events evolve linearly without deviation.

Curved $T_{\mu\nu}$: Dynamic Interaction

- **Higher** $T_{00} = 2$: Increased "mass" introduces complexity, enriching the paragraph with new elements (e.g., "barks," "yells").
- $T_{01} = 0.5$, $T_{02} = 1$: A faster topical flow shifts from "runs" to "barks" to "yells," while tonal flow escalates from "fast" to "loud" to an urgent "now."
- $T_{12} = 2$: High tension/twist couples topic and tone, creating a narrative arc where Spot's barking builds into Jane's yelling.
- **Result:** A dynamic, escalating story—akin to a curved spacetime where "gravitational" pulls from tension alter the trajectory of the narrative.

Ricci Curvature R as a Twist Indicator

- **Flat** ($R \approx 0$): Near-zero curvature reflects a lack of deviation from the seed sentence's trajectory. The paragraph remains static, with no surprises or emotional escalation—analogueous to a flat, unchanging landscape.
- **Curved** ($R > 0$): Positive curvature indicates a "converging" twist where the narrative builds intensity and shifts direction (e.g., from motion to noise to urgency). This mimics a spacetime warped by mass, pulling events toward a focal point of dramatic escalation.

Geodesic Path and Word Choice

The evolution of text follows geodesics determined by the perturbed metric $g_{\mu\nu}$, influencing the trajectory of word choices.

Flat: Linear Geodesic

- $x_1 = 6.1 \rightarrow 6.3 \rightarrow 6.5$: Stays near "run/jump" verbs.
- $x_2 = 3$: Locked at "fast," with no tonal evolution.
- **Nuance:** The paragraph exhibits repetition and consistency, similar to a drill or list.

Curved: Nonlinear Geodesic

- $x_1 = 6.1 \rightarrow 6.6 \rightarrow 7.1$: Moves from "runs" to "barks" (a topical shift via Spot) to "yells."
- $x_2 = 3 \rightarrow 4 \rightarrow 7$: Tone rises from "fast" to "loud" to an implied peak ("now" for urgency).
- **Nuance:** A mini-story emerges, where action intensifies and hints at conflict or chaos.

Physical Analogies

- **Flat:** Like a vacuum with no mass ($T_{\mu\nu} \approx 0$), spacetime remains unperturbed, and geodesics follow straight lines. The paragraph mirrors this structure, with no "gravitational" pull influencing word transitions.
- **Curved:** Like a star's gravity warping spacetime, high $T_{\mu\nu}$ values (e.g., $T_{12} = 2$) bend the geodesic path, creating a narrative "orbit" around tension and tone shifts. The paragraph evolves toward a dramatic transition.

Narrative Implications

- **Flat:** Ideal for rote learning, instructional material, or simple descriptions (e.g., early reader exercises). The stability of $T_{\mu\nu}$ maintains clarity but lacks engagement—analogueous to a low-energy state in physics.
- **Curved:** Suitable for storytelling, emotional impact, and complex narratives. The interplay of T_{12} (tension) and T_{02} (tonal flow) generates a sense of progression, conflict, or resolution.

Practical Insights

- **Control:** The stress-energy tensor $T_{\mu\nu}$ functions as a "narrative throttle":
 - Low values lead to stability and repetition.
 - High values introduce change and dramatic shifts.
- **Scalability:** Expanding the dictionary size allows T_{12} to generate more complex twists (e.g., "Spot barks at Jane" evolving into "Jane falls badly").
- **Quantification:** The Ricci scalar R provides an objective metric for assessing the "interestingness" of text:

- $R = 0$: A simple, predictable structure.
- $R > 0$: A dynamic, evolving narrative.

Key Takeaways

- The **flat paragraph** follows a straight-line trajectory in semantic spacetime, with low $T_{\mu\nu}$ maintaining stability and predictability ($R \approx 0$).
- The **curved paragraph** follows a bent geodesic path, where high $T_{\mu\nu}$ values (particularly T_{12} and T_{02}) introduce curvature ($R > 0$), leading to tonal escalation and topical twists.
- This framework offers a physics-inspired approach to understanding how tension, flow, and complexity shape the geometric structure of a narrative.

4.7 Future Work

- Train an SLM with this framework to bias outputs using $T_{\mu\nu}$.
- Test higher-order perturbations to simulate extreme tonal shifts.
- Explore negative curvature regimes (e.g., melancholic narratives).

5 Computational Analysis of $T_{\mu\nu}$ and Ricci Curvature in a Full-Length Novel

To validate our framework on real-world text, we applied our semantic spacetime metric approach to 100 randomly selected paragraphs from the science fiction novel *Ballistic* by Travis S. Taylor. We computed the stress-energy tensor $T_{\mu\nu}$ for each paragraph, extracted the corresponding Ricci curvature R , and analyzed the statistical distribution of these values. The implementation was performed using Python, leveraging natural language processing (NLP) tools, numerical tensor operations, and geodesic integration.

5.1 Computational Methodology

5.1.1 Random Paragraph Selection

To ensure a diverse sampling, we extracted 100 paragraphs randomly from *Ballistic*. The document was processed using the `python-docx` package to extract text while preserving paragraph structure.

5.1.2 Semantic Stress-Energy Tensor $T_{\mu\nu}$

For each paragraph, we computed the semantic stress-energy tensor $T_{\mu\nu}$ with the following components:

- T_{00} (Semantic Mass): A measure of complexity, approximated via entropy of word embeddings.
- T_{01} (Topical Flow): A measure of thematic consistency, computed via cosine similarity between sentence embeddings.
- T_{02} (Tonal Flow): Sentiment shift, determined using sentiment analysis differentials.
- T_{12} (Tension/Twist): Cross-coupling between tone and topic, extracted via principal component analysis (PCA) on high-dimensional embedding space.

Mathematically, this was implemented as:

$$T_{\mu\nu} = \begin{bmatrix} \text{Complexity} & \text{Topical Flow} & \text{Tonal Flow} \\ \text{Topical Flow} & \text{Topical Pressure} & \text{Tension/Twist} \\ \text{Tonal Flow} & \text{Tension/Twist} & \text{Tonal Pressure} \end{bmatrix}. \quad (27)$$

5.1.3 Ricci Curvature Computation

The Ricci curvature R was computed from the metric perturbations of the semantic space, using the standard formulation:

$$R_{\mu\nu} = \partial_\alpha \Gamma_{\mu\nu}^\alpha - \partial_\nu \Gamma_{\mu\alpha}^\alpha + \Gamma_{\mu\lambda}^\alpha \Gamma_{\nu\alpha}^\lambda - \Gamma_{\nu\lambda}^\alpha \Gamma_{\mu\alpha}^\lambda. \quad (28)$$

From the Ricci tensor, the scalar curvature R was extracted via:

$$R = g^{\mu\nu} R_{\mu\nu}. \quad (29)$$

6 Extracting the Metric Tensor $g_{\mu\nu}$ from a Novel

To systematically analyze the semantic curvature of a full-length novel, we developed a computational pipeline to extract the metric tensor $g_{\mu\nu}$ by measuring topic evolution, tonal variation, and structural flow. This approach allows us to quantify how meaning and emotional intensity shift throughout the text.

6.1 Computational Approach

To compute the semantic spacetime metric $g_{\mu\nu}$ from the text, we followed a structured approach:

1. Convert the novel into numerical embeddings: Each paragraph is transformed into a high-dimensional vector representation using a sentence embedding model.
2. Compute topic evolution: Measure the semantic deviation between consecutive paragraphs using cosine similarity in embedding space.

3. Compute tonal evolution: Track changes in emotional intensity by applying a sentiment analysis model to extract sentiment gradients.
4. Construct the metric tensor $g_{\mu\nu}$ for each segment: Using topic shifts, tonal variations, and structural flow, we define a tensor that encapsulates semantic connectivity.
5. Analyze and visualize curvature R : By computing Ricci curvature, we identify key structural elements such as climaxes, twists, and expositions in the novel.

6.2 Metric Tensor Formulation

We define the semantic spacetime metric $g_{\mu\nu}$ as:

$$g_{\mu\nu} = \begin{bmatrix} -1 & T_{01} & T_{02} \\ T_{01} & 1 + T_{11} & T_{12} \\ T_{02} & T_{12} & 1 + T_{22} \end{bmatrix}. \quad (30)$$

Here:

- T_{01} represents topical flow (how much the topic changes over time).
- T_{02} represents tonal flow (how much emotional intensity shifts).
- T_{12} represents tension/twist (how coupled topic and tone are).
- T_{11}, T_{22} represent semantic pressures, controlling how resistant a passage is to change.

6.3 Discrete Implementation: From Paragraphs to Curvature

Our curvature computation is implemented on a discrete sequence of paragraph indices $i = 0, 1, \dots, N - 1$, with the paragraph index serving as the time-like coordinate t . Each paragraph P_i yields a stress-energy tensor $T_{\mu\nu}(i)$, a metric $g_{\mu\nu}(i)$, and a scalar curvature $R(i)$.

Step 1: Compute stress-energy components

For each paragraph P_i , we compute:

$$T_{00}(i) = H(E(P_i)) \quad (\text{embedding entropy / semantic mass}), \quad (31)$$

$$T_{01}(i) = 1 - \cos(E(P_i), E(P_{i+1})), \quad (32)$$

$$T_{02}(i) = \text{Sent}(P_{i+1}) - \text{Sent}(P_i), \quad (33)$$

$$T_{12}(i) = \text{Cov}(\Pi_{\text{topic}}(P_i), \Pi_{\text{tone}}(P_i)), \quad (34)$$

where $E(\cdot)$ denotes a sentence-embedding map, $\text{Sent}(\cdot)$ a sentiment score, and $\Pi_{\text{topic}}, \Pi_{\text{tone}}$ are low-dimensional projections (e.g., PCA components) used to estimate the topic–tone coupling.

$H(\cdot)$ denotes a Shannon entropy estimate computed from a discretized distribution of embedding components.

Step 2: Construct the metric

Using these components, we define the $(1 + 2)$ -dimensional semantic metric:

$$g_{\mu\nu}(i) = \begin{bmatrix} -1 & T_{01}(i) & T_{02}(i) \\ T_{01}(i) & 1 + T_{11}(i) & T_{12}(i) \\ T_{02}(i) & T_{12}(i) & 1 + T_{22}(i) \end{bmatrix}. \quad (35)$$

In practice, T_{11} and T_{22} are treated as diagonal “pressure” regularizers to keep the metric well-conditioned. We use constant baseline values $T_{11}(i) = \alpha$, $T_{22}(i) = \beta$ unless otherwise stated.

Step 3: Discrete derivatives

Because $g_{\mu\nu}$ is defined on paragraph index i , derivatives are approximated by finite differences along t :

$$\partial_t g_{\mu\nu}(i) \approx \frac{g_{\mu\nu}(i + 1) - g_{\mu\nu}(i - 1)}{2\Delta t}, \quad (36)$$

with $\Delta t = 1$ paragraph step. Spatial derivatives $\partial_{x_1} g_{\mu\nu}$ and $\partial_{x_2} g_{\mu\nu}$ are set to zero in this first implementation, so curvature arises from temporal variation of the metric along the narrative.

Step 4: Connection and curvature

We compute the Christoffel symbols:

$$\Gamma_{\mu\nu}^\alpha(i) = \frac{1}{2} g^{\alpha\lambda}(i) (\partial_\mu g_{\lambda\nu}(i) + \partial_\nu g_{\lambda\mu}(i) - \partial_\lambda g_{\mu\nu}(i)), \quad (37)$$

then form the Ricci tensor $R_{\mu\nu}(i)$ and scalar curvature:

$$R(i) = g^{\mu\nu}(i) R_{\mu\nu}(i). \quad (38)$$

Interpretation

Under this discrete-time formulation, spikes in $R(i)$ correspond to intervals where the semantic metric changes rapidly from paragraph to paragraph—i.e., strong narrative transitions, tonal jumps, or topic shifts. Near-zero curvature corresponds to steady semantic flow under an approximately constant metric.

6.4 Extracting Semantic and Tonal Evolution

6.4.1 Topic Evolution

For each paragraph pair (P_i, P_{i+1}) , the topic shift Δ_{topic} is computed as:

$$T_{01}(i) = 1 - \cos(\theta), \quad (39)$$

where θ is the angle between the paragraph embeddings:

$$\cos(\theta) = \frac{E(P_i) \cdot E(P_{i+1})}{\|E(P_i)\| \|E(P_{i+1})\|}. \quad (40)$$

6.4.2 Tonal Evolution

To track the emotional intensity of the narrative, we define the tonal shift Δ_{tone} as:

$$T_{02}(i) = \text{Sentiment}(P_{i+1}) - \text{Sentiment}(P_i), \quad (41)$$

where Sentiment is obtained via a sentiment analysis model, mapping text to a scale from -1 (negative) to 1 (positive).

6.4.3 Geodesics and Semantic Paths

Once $g_{\mu\nu}$ is computed, we calculate the geodesic equation to determine the natural evolution of the text:

$$\frac{d^2 x^\mu}{d\tau^2} + \Gamma_{\alpha\beta}^\mu \frac{dx^\alpha}{d\tau} \frac{dx^\beta}{d\tau} = 0. \quad (42)$$

This equation predicts how the story should naturally flow, and deviations from this trajectory indicate unexpected plot twists or nonlinear storytelling structures.

6.5 Visualization of Curvature in the Novel

To assess how narrative complexity varies, we compute the Ricci curvature R at each segment:

$$R = g^{\mu\nu} R_{\mu\nu}. \quad (43)$$

Regions with high R correspond to narrative singularities, such as major revelations, conflicts, or emotional climaxes.




Figure 1: Computed Ricci curvature R across the novel *Ballistic*. Peaks correspond to climactic or highly dynamic sections.

6.6 Key Observations

- Flat $R \approx 0$ regions correspond to expository or steady-flow sections.
- High R values coincide with narrative tension build-ups (e.g., action sequences, character confrontations).
- Negative R regions suggest transitional passages where themes shift but without strong emotional shifts.

The Ricci curvature plot shown in Figure 1 reveals where these structural deviations occur, allowing us to quantify the impact of narrative twists and pacing.

Note: throughout this work, Ricci curvature is computed on a discrete, paragraph-indexed manifold and should be interpreted as a semantic curvature proxy rather than a continuous spacetime invariant.

6.7 Implications for AI-Generated Narrative Structure

By integrating these curvature-based metrics into language models, we can:

- Improve story pacing: Ensuring smooth narrative arcs rather than disjointed jumps.
- Enhance AI-driven storytelling: Predict and control story complexity dynamically.
- Provide editorial insights: Helping authors visualize and fine-tune plot structure.

6.8 Conclusion

This approach allows for a novel geometric perspective on storytelling, where spacetime curvature analogies offer a rigorous mathematical basis for structuring AI-generated narratives.

6.9 Results and Analysis

6.9.1 Statistical Distribution of $T_{\mu\nu}$

We examined the distribution of semantic mass, topical flow, tonal flow, and tension/twist across the selected paragraphs. To analyze the structural dynamics of the novel, we examined the statistical distributions of the key stress-energy tensor components: semantic mass (T_{00}), topical flow (T_{01}), tonal flow (T_{02}), and tension/twist (T_{12}). Each component provides insight into different aspects of narrative progression and complexity.




Figure 2: Distribution of T_{00} (Semantic Mass) across 100 sampled paragraphs.

Semantic Mass T_{00} As shown in Figure 2, the distribution of T_{00} reveals that most paragraphs exhibit moderate semantic density, corresponding to a steady level of informational content. However, certain sections show increased values, indicating conceptually dense passages, likely involving key plot developments or expository moments.




Figure 3: Distribution of T_{01} (Topical Flow) across 100 sampled paragraphs.

Topical Flow T_{01} The histogram in Figure 3 illustrates how topics evolve throughout the novel. A low T_{01} value signifies gradual shifts between topics, while a subset of paragraphs exhibits higher values, indicating abrupt transitions—potentially reflecting major scene changes, time jumps, or rapid conceptual shifts.




Figure 4: Distribution of T_{02} (Tonal Flow) across 100 sampled paragraphs.

Tonal Flow T_{02} Figure 4 displays the distribution of T_{02} , which quantifies fluctuations in emotional intensity. The majority of paragraphs maintain stable tonal characteristics, while a subset demonstrates larger variations, corresponding to moments of heightened tension, humor, or emotional shifts within the narrative.




Figure 5: Distribution of T_{12} (Tension/Twist) across 100 sampled paragraphs.

Tension/Twist T_{12} The distribution of T_{12} (Figure 5) is particularly revealing, as it measures the coupling between topic and tone. Most paragraphs exhibit low values, meaning their subject matter and emotional tone evolve independently. However, the presence of higher T_{12} values in certain segments suggests narrative curvature, where meaning and sentiment interact dynamically—potentially marking twists, foreshadowing, or emotionally charged dialogue.

Overall, these distributions highlight the structural complexity of the novel, revealing regions of relative stability versus those exhibiting high curvature and narrative dynamism. In the next section, we examine how these components contribute to overall Ricci curvature and its role in identifying key storytelling moments.

Ricci Curvature Analysis of Narrative Structure

To quantify the dynamical complexity of the novel’s progression, we computed the temporal evolution of the Ricci curvature R across 100 randomly sampled paragraphs, as shown in Figure 6. The majority of sections exhibit near-zero curvature, indicating linear or predictable semantic evolution. However, the presence of peaks at specific intervals suggests significant deviations from a uniform narrative flow. These peaks correspond to key storytelling moments, such as climaxes, major plot twists, or sudden tonal shifts.




Figure 6: Ricci curvature R temporal evolution computed across 100 sampled paragraphs of the novel. Peaks correspond to key narrative shifts, such as climaxes and major transitions.

From a general relativity perspective, these spikes in curvature indicate regions where the semantic and tonal densities change rapidly, akin to gravitational wells distorting the geodesics of a spacetime manifold. In practical terms, these are points in the story where readers experience heightened engagement, as the narrative shifts direction in an unexpected or dramatic way. Conversely, flatter regions correspond to sections of steady exposition or transitional passages, where the storyline progresses with minimal tension or surprise.

This analysis highlights the potential of geometric approaches to text modeling, revealing structural features of storytelling that are traditionally analyzed qualitatively. By associating Ricci curvature with narrative dynamics, we propose a new metric for quantifying engagement and complexity in literary works, opening pathways for computational storytelling and AI-assisted narrative generation.

6.10 Key Observations

- **Low-curvature paragraphs** ($R \approx 0$) correspond to expository sections where information is conveyed with minimal deviation.
- **High-curvature paragraphs** ($R > 0$) appear near climactic scenes in the novel, reflecting rapid tonal escalation and thematic shifts.
- **Negative curvature** ($R < 0$) occurs in transitional moments, possibly reflecting topic divergence or scene changes.

6.11 Comparison with Theoretical Predictions

The measured $T_{\mu\nu}$ and R values were compared to those from the small language model (SLM) experiments conducted earlier. A clear correlation emerged:

- Tonal flow (T_{02}) and tension/twist (T_{12}) strongly correlate with narrative curvature R , confirming that increasing twist or tonal variation enhances semantic spacetime curvature.
- Expository paragraphs mirror Minkowski spacetime, while climactic paragraphs resemble Schwarzschild-like gravitational wells, where semantic "mass" pulls narrative elements inward.
- Highly dynamic paragraphs show FLRW-like expansion, where the semantic scale factor $a(t)$ suggests large-scale shifts in meaning.

6.12 Implications for LLM Training and Generation

These findings suggest that integrating curvature-based metrics into transformer architectures could provide the following benefits:

- **Curvature-aware text generation:** LLMs could adjust Ricci curvature constraints to control pacing and complexity dynamically.
- **Geometric loss functions:** Training objectives could incorporate deviation from a target R to fine-tune narrative engagement.
- **Content structuring tools:** Editors and AI-assisted writing software could visualize and manipulate a text's "semantic curvature" to optimize narrative impact.

This computational study demonstrates that natural text exhibits measurable curvature under our semantic stress-energy framework. The correlation between stress-energy terms ($T_{\mu\nu}$) and Ricci curvature (R) provides a novel, physics-inspired lens for understanding and guiding LLM-based text generation.

7 Testing with SLM Vocabulary

To further validate our approach, we tested the Small Language Model (SLM) vocabulary on a structured text generation experiment. Using a 78-word dictionary, we applied the spacetime metric framework to generate structured paragraphs, systematically modifying the scale factor $a(t)$ and stress-energy components $T_{\mu\nu}$ to shape the narrative arc dynamically.

7.1 Experimental Setup

Our experiment was structured in five distinct narrative phases, with corresponding changes in $T_{\mu\nu}$ dynamics:

Time Step t	$a(t)$	x_1, x_2 (Topic, Tone)	Generated Sentence
0	1.0	(6.0, 3.0)	"Dick sees a big car fast."
1	1.1	(7.0, 5.0)	"Jane runs to the car loud."
2	1.2	(7.1, 5.0)	"Jane looks at the car now."
3	1.3	(7.2, 5.0)	"Dick walks to Jane here."
4	1.4	(7.7, 5.5)	"Spot jumps on the car fast."
5	1.7	(8.2, 6.0)	"Dick calls Spot with fun."
6	1.8	(8.2, 5.0)	"The car falls slow then."
7	1.9	(8.2, 4.0)	"Jane looks sad at Spot."
8	2.7	(9.2, 8.0)	"Dick yells loud and runs fast."
9	4.5	(10.2, 12.0)	"Spot barks at the car and Jane catches Dick now."

Table 2: Evolution of paragraph generation based on $a(t)$ and $T_{\mu\nu}$ variations.

1. Action Phase ($t = 0 - 1$) – High T_{00} (semantic mass) and T_{02} (tonal flow), representing rapid engagement.
2. Lull Phase ($t = 2 - 3$) – Low $T_{\mu\nu}$, representing slower-paced exposition.
3. Rising Action ($t = 4 - 5$) – Moderate T_{12} (twist factor) and T_{01} (topic shift), representing increasing tension.
4. Dip Phase ($t = 6 - 7$) – Negative T_{02} , high T_{00} , modeling a setback or conflict.
5. Climax Phase ($t = 8 - 9$) – Peak T_{12} and T_{22} , representing the highest tension and resolution.

7.2 Paragraph Generation via Metric Tensor Evolution

Each paragraph was generated by perturbing the metric tensor $g_{\mu\nu}$ based on evolving topic (x_1) and tone (x_2) values, with $a(t)$ dictating narrative expansion. The table below summarizes the evolution of paragraphs based on these calculations:

The generated story follows a clear trajectory, transitioning from calm observation (low curvature) to heightened tension and resolution (high curvature). This progression aligns with the computed Ricci curvature R values:

- $t = 0 - 1$: $R > 0.5$ – High curvature in the action-driven phase.
- $t = 2 - 3$: $R \approx 0$ – A flat lull, little progression.
- $t = 4 - 5$: $R > 0.3$ – Moderate curvature as tension builds.
- $t = 6 - 7$: $R < 0$ – Negative curvature, indicating a setback or narrative downturn.
- $t = 8 - 9$: $R \gg 0$ – Strong positive curvature as the climax unfolds.

Story 1 (Linear Growth)	Story 2 (Exponential Growth)
Dick sees a big car fast.	Dick sees a big car fast.
Jane runs to the car now.	Jane runs fast to the car.
Spot jumps on the car fast.	Spot jumps loud on the car.
Dick calls Spot with fun.	Dick yells fun at Spot now.
Jane yells loud at the car.	Jane calls kids and the car goes loud.
Spot barks and the car goes fast.	Spot barks at trees with Jane fast.
Dick runs with Jane to the hill.	Dick runs to the hill and yells loud.
The car falls loud on trees.	The car falls on the park with kids now.
Spot yells and kids play fast.	Jane catches Spot and Spot barks loud fast.
Jane catches Dick and Spot barks loud.	Dick yells at cars and trees fall loud with fun.

Table 3: Comparison of two generated stories using different expansion functions: Story 1 follows linear growth ($a(t) = 1 + kt$) while Story 2 follows exponential growth ($a(t) = e^{kt}$).

7.3 Comparing Different Expansion Models: Linear vs. Exponential

We examined how different narrative expansion models affect story progression. Specifically, we tested:

- Linear Growth: $a(t) = 1 + kt$ with $k = 0.5$, modeling gradual, steady tension increase.
- Exponential Growth: $a(t) = e^{kt}$ with $k = 0.3$, modeling runaway escalation with no resolution.

The comparison in Table 3 highlights the distinct narrative trajectories imposed by different expansion functions. The linear growth model ($a(t) = 1 + kt$) results in a steadily evolving story, where new elements are introduced at a controlled pace, maintaining coherence and resolution. The tension increases gradually, allowing for structured development, such as Jane’s engagement with the car and Spot’s interactions.

In contrast, the exponential growth model ($a(t) = e^{kt}$) leads to runaway narrative expansion mirroring characteristics of chaotic or unstable dynamics [19, 20], where each new sentence rapidly escalates the stakes, introducing abrupt shifts in tone and context. The story moves from simple observations to chaotic interactions, with multiple actors and settings emerging in quick succession. The exponential model lacks natural points of rest, mirroring a universe with unchecked expansion—leading to a sense of narrative instability and unresolved climax.

This experiment demonstrates how different curvature-driven models of semantic expansion influence the reader’s perception of coherence, pacing, and narrative resolution. The linear model maintains control and readability, while the exponential model introduces an unpredictable, feverish intensity. In future work, intermediate models could balance these dynamics, offering structured yet dynamic storytelling frameworks based on physical principles of spacetime evolution.




Figure 7: Comparison of Ricci curvature R for linear vs. exponential $a(t)$ growth. The exponential model results in an uncontrolled escalation, while the linear model provides a structured but ever-growing tension.

7.4 Findings and Insights

1. Growth Model Effects on Narrative Progression The linear expansion resulted in a well-paced progression, with gradually increasing stakes and a climactic resolution. In contrast, the exponential expansion caused an uncontrolled escalation, leading to a chaotic and unresolved story. This suggests that narrative tension should be moderated rather than continually accelerated.

2. $T_{\mu\nu}$ as a Narrative Driver The behavior of the stress-energy tensor $T_{\mu\nu}$ was closely linked to the semantic shape of the text:

- Low $T_{\mu\nu}$ values (lull phases) led to near-zero Ricci curvature, reflecting stable sections.
- High T_{12} values (climax) warped the metric significantly, driving the story forward.
- Negative T_{02} values (dips) introduced narrative downturns, reflecting setbacks.

3. Ricci Curvature as a Literary Complexity Metric The computed curvature R as shown in Figure 7 serves as a direct measure of story dynamism. Specific story changes from word to word are found from a more resolved map of the curvature as shown previously in Figure 6. Peaks correlate with high-energy sequences, while valleys correspond to pauses or plateaus. This validates the potential use of curvature-based metrics to quantify narrative pacing.

7.5 Implications for AI-Assisted Storytelling

By linking spacetime geometry to text generation, we propose a new way to modulate AI-generated narratives:

- Controlling $a(t)$ and $T_{\mu\nu}$ allows for fine-tuned pacing, ensuring structured progression rather than chaotic jumps.
- Ricci curvature provides an objective measure of storytelling tension, applicable in LLMs for detecting and adjusting pacing dynamically.
- The framework enables AI to generate narratives that mirror classical storytelling arcs, allowing for more engaging and human-like storytelling.

Future research will explore adaptive tensor weighting, where $T_{\mu\nu}$ evolves dynamically based on user input, allowing real-time control over AI-generated stories.

8 Conclusion and Future Work

By reframing natural language generation as a curvature-based dynamical system, this work opens the possibility of applying tools from chaos theory and general relativity to steer semantic trajectories. Such approaches could inform next-generation AI architectures by grounding coherence and narrative structure in physical analogs

8.1 Summary of Contributions

In this work, we introduced a novel spacetime metric framework for transformer-based language models, drawing direct analogies between semantic structure in text and the mathematical formalism of General Relativity (GR). By redefining linguistic evolution in terms of tensor calculus and curvature, we developed a semantic stress-energy tensor $T_{\mu\nu}$ that quantifies the mass-energy, topic flow, tonal dynamics, and tension within a text. Using this framework, we established:

- The **Minkowski metric** as a representation of standard transformer architectures, describing texts with linear progression, stable meaning propagation, and minimal curvature.
- The **Schwarzschild metric** as a model for texts with strong central themes, semantic "gravity wells," and event horizons where meaning collapses into incomprehensibility.
- The **FLRW metric** as an expansion model for dynamic narratives, where meaning evolves over time in an accelerating or oscillatory fashion.
- The Ricci curvature R as a measurable metric for quantifying narrative complexity, tension, and story progression.
- The geodesic equation as a predictor of coherent storytelling, where texts following optimal geodesics maintain logical progression while high curvature regions correspond to dramatic shifts.

To validate these ideas, we performed computational experiments on the novel *Ballistic*, computing $T_{\mu\nu}$ across 100 randomly chosen paragraphs, deriving curvature values, and visualizing the resulting semantic structure. The results confirmed that:

- Paragraphs with low curvature correspond to expository or steady-flow sections.
- Paragraphs with high curvature align with climactic moments, shifts in tone, or major narrative twists.
- Negative curvature regions suggest transitional passages, where semantic or emotional shifts dissipate rather than escalate.

Additionally, we tested our framework on a Small Language Model (SLM) using a constrained vocabulary to generate structured narratives with varying $a(t)$ and $T_{\mu\nu}$. We found that:

- Linear expansion models ($a(t) = 1 + kt$) result in structured, gradual tension buildup.
- Exponential growth models ($a(t) = e^{kt}$) lead to runaway escalation, loss of coherence, and unresolved complexity.

These findings provide a new mathematical foundation for controlling story pacing, engagement, and complexity in AI-generated text.

8.2 Broader Implications

The implications of this work extend beyond narrative analysis into several domains:

- Improving Large Language Models (LLMs): Current transformers generate text using local token-by-token probability distributions. By incorporating curvature constraints, LLMs could be trained to maintain long-range coherence, ensuring narratives develop in a controlled and meaningful way.
- AI-Assisted Editing and Narrative Structuring: By computing $T_{\mu\nu}$ and R for a given manuscript, authors could visualize where their story lacks engagement, where plot twists could be introduced, and how to optimize pacing for maximum reader impact.
- Quantum-Inspired AI Architectures: Our findings suggest that semantic superposition and uncertainty could be modeled using quantum wavefunction-like representations. Tensor networks inspired by quantum field theory may allow for nonlocal meaning propagation in AI-generated text.
- Cognitive Science and Psycholinguistics: The semantic curvature framework provides a rigorous mathematical way to study how humans process language, how information is stored and retrieved in memory, and why certain narrative structures are more engaging than others.

While this study focuses on language models, the semantic curvature framework could be applied to other high-dimensional symbolic systems, including biological signaling networks, social discourse structures, or neural activity flows. Any system with measurable topic, tone, or phase structure could be modeled as a dynamical manifold.

8.3 Future Work

Several promising directions emerge from this research, including:

- Exploring Higher-Dimensional GR Metrics: Beyond Schwarzschild and FLRW, additional metrics such as Kerr (rotating frames), Reissner-Nordström (charged texts with layered meaning), or Anti-de Sitter spaces (cyclic narratives) could further refine text dynamics.
- Applying Gauge Theory to Semantic Flow: By treating attention heads in transformers as gauge fields, we could model semantic coherence as a conserved charge, ensuring stability in long-form text generation.
- Developing Tensor-Based Loss Functions for AI Training: Instead of optimizing purely on token-level accuracy, LLMs could be trained to minimize geodesic deviation and maximize desired curvature patterns, leading to more structured, naturalistic text.
- Generalizing to Multimodal AI: Our curvature-driven framework could be extended to image generation (visual storytelling) and music composition (harmonic structure evolution), providing a unified geometric approach to creative AI.
- Integrating with Explainable AI (XAI): By interpreting model outputs through the lens of differential geometry, we could develop a transparent, physics-based explanation for why AI-generated text follows specific patterns.

8.4 Final Thoughts

This paper demonstrates that a spacetime-inspired approach to language modeling offers a rigorous, mathematically grounded perspective on AI-generated narratives. By applying GR-inspired techniques—tensor calculus, curvature analysis, and geodesic evolution—we introduce a geometric foundation for next-generation LLMs capable of enhanced coherence, structured complexity, and improved interpretability. Future research will continue refining these ideas, bridging the gap between physics, linguistics, and artificial intelligence.

This curvature-based approach offers new opportunities for applying techniques from chaos theory, including geodesic deviation analysis, curvature-induced attractors, and instability diagnostics, to the behavior of high-dimensional language systems. This work bridges nonlinear systems theory and geometric modeling of symbolic behavior in AI [6, 7, 18], providing a new mathematical foundation for controlling narrative evolution in transformer-based models.

This framework also opens new avenues for applying chaos theory to symbolic systems. In particular, geodesic deviation under high semantic tension (T_{12}) may lead to bifurcations in narrative trajectories or sensitivity to initial prompt conditions—hallmarks of chaotic systems. The evolution of meaning through curved semantic manifolds suggests the possibility of attractor structures or chaotic basins in narrative development. Future work could explore Lyapunov-like divergence in

semantic geodesics or identify semantic analogs of strange attractors within transformer-generated text.

Data Availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

References

- [1] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [2] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019.
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, and Amanda et al. Askell. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 2020.
- [4] Qiao Liu, Hao Wang, Chenwei Liu, Zhe Xu, Wayne Xin Zhao, Jingren Zhou, and Ji-Rong Wen. Pre-train, prompt, and predict: A survey of language modeling paradigms. *arXiv preprint arXiv:2107.13586*, 2022.
- [5] Jason Chuang, Caroline Uhler, and Stefanie Jegelka. Learning on riemannian manifolds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(11):2926–2941, 2020.
- [6] Godfried T. Toussaint. Geometric approaches to word meaning. In *Mathematics of Language*, volume 6878 of *Lecture Notes in Artificial Intelligence*, pages 43–57. Springer, 2012.
- [7] Barrett F Benford and Adilson E Motter. Topology of conceptual networks in natural language. *Physical Review E*, 78(3):036108, 2008.
- [8] William Bialek. *Biophysics: Searching for Principles*. Princeton University Press, Princeton, NJ, 2012. Explores physical principles in biological systems, including information geometry and dynamics.
- [9] William Bialek and Fred Rieke. Reliability and information transmission in spiking neurons. *Trends in Neurosciences*, 15(11):428–434, 1992. Discusses information-theoretic encoding and variability in biological signal processing.

- [10] William Bialek, Ilya Nemenman, and Naftali Tishby. Predictability, complexity, and learning. *Neural Computation*, 13(11):2409–2463, 2001. Introduces information bottleneck methods and complexity metrics relevant to semantic compression.
- [11] Albert Einstein. *The Field Equations of Gravitation*. 1915.
- [12] Bernard Schutz. *A First Course in General Relativity*. Cambridge University Press, 2009.
- [13] Charles W Misner, Kip S Thorne, and John Archibald Wheeler. *Gravitation*. W. H. Freeman, 1973.
- [14] Sean M Carroll. *Spacetime and Geometry: An Introduction to General Relativity*. Cambridge University Press, 2019.
- [15] Scott Dodelson and Fabian Schmidt. *Modern Cosmology*. Academic Press, 2020.
- [16] Riccardo Di Sipio, Jairo Diaz-Rodriguez, and Luis Serrano. The curved spacetime of transformer architectures. 2025. arXiv preprint arXiv:2511.03060.
- [17] Shun-ichi Amari. Natural gradient works efficiently in learning. *Neural Computation*, 10(2):251–276, 1998.
- [18] Shun-ichi Amari. *Information Geometry and Its Applications*. Springer, 2016.
- [19] Theodore E. Porter. Language as a nonlinear dynamical system. *Nonlinear Dynamics, Psychology, and Life Sciences*, 3(2):127–145, 1999.
- [20] Rudolf Hanel and Stefan Thurner. When do universal linguistic laws hold? *Proceedings of the National Academy of Sciences*, 114(17):4600–4604, 2017.