

V1.0.0

Cross-Ratio Invariants as a Formal Measure for Translation Fidelity and Analogical Reasoning

Rowan Brad Quni-Gudzin

QNFO-ULA

ABSTRACT

Machine translation (MT) and analogical reasoning both require semantic representations that remain stable under surface-form changes. Existing similarity measures such as cosine distance ignore projective structure, limiting their ability to capture deep relational meaning. This work addresses the gap by proposing a cross-ratio invariant $\chi(A, B; C, D)$ computed over word-embedding vectors as a coordinate-free semantic measure. Building on the cycles-cross-ratio invariant introduced for Lie-sphere geometry[1] and the strict cross-ratio invariance under projective maps[11], we formalize χ for quadruples of embedding points, prove its invariance under any fractional linear transformation, and derive a closed-form analogy-completion formula. Using sentence-level weighting statistics from domain-adaptation corpora[8] we compute concrete data-reduction proportions (51.1% of sentences, 28.1% of tokens). Combining this reduction with beam-search pruning speeds of up to 43%[2] yields a projected overall decoding speedup of roughly 510% ($\approx 5.1 \times$ faster) under an independence assumption. We also enumerate the combinatorial budget of candidate invariants for a 10 000-item vocabulary and estimate the minimal sample size (≈ 85 aligned pairs) required to detect a modest correlation ($r = 0.30$) with 80% power. The result suggests that cross-ratio preservation can serve as a lightweight, mathematically grounded proxy for translation adequacy and analogy correctness, offering a new evaluation pathway for low-resource MT system design and analogical benchmark construction.

CITE AS Quni-Gudzin, R. B. (n.d.). Cross-Ratio Invariants as a Formal Measure for Translation Fidelity and Analogical Reasoning. QNFO Preprint v1.0.0.

Abstract

Machine translation (MT) and analogical reasoning both require semantic representations that remain stable under surface-form changes. Existing similarity measures such as cosine distance ignore projective structure, limiting their ability to capture deep relational meaning. This work addresses the gap by proposing a cross-ratio invariant $\chi(A, B; C, D)$ computed over word-embedding vectors as a coordinate-free semantic measure. Building on the cycles-cross-ratio invariant introduced for Lie-sphere geometry[1] and the strict cross-ratio invariance under projective maps[11], we formalize χ for quadruples of embedding points, prove its invariance under any fractional linear transformation, and derive a closed-form analogy-completion formula. Using sentence-level weighting statistics from domain-adaptation corpora[8] we compute concrete data-reduction proportions (51.1% of sentences, 28.1% of tokens). Combining this reduction with beam-search pruning speeds of up to 43%[2] yields a projected overall decoding speedup of roughly 510% ($\approx 5.1 \times$ faster) under an independence assumption. We also enumerate the combinatorial budget of candidate invariants for a 10 000-item vocabulary and estimate the minimal sample size (≈ 85 aligned pairs) required to detect a modest correlation ($r = 0.30$) with 80% power. The result suggests that cross-ratio preservation can serve as a lightweight, mathematically grounded proxy for translation adequacy and analogy correctness, offering a new evaluation pathway for low-resource MT system design and analogical benchmark construction.

1. Introduction

Global communication increasingly relies on neural machine translation (NMT) services that must render content accurately across thousands of language pairs. Translation providers, multilingual content platforms, and international NGOs face costly failures when translations distort meaning, leading to misinformation, legal liability, and wasted computational resources. Low-resource language teams translating water-sanitation-hygiene guidelines for Indian tribal children or earthquake-preparedness material for Indonesian communities work with little parallel data and few human experts[4]. Current NMT decoders use beam search to approximate the most probable output[2]; enlarging the beam improves BLEU scores but reduces throughput, while pruning techniques can recover up to 43% of decoder speed[2] without degrading quality. Analogical reasoning tasks such as SAT-style “A:B::C:?” also suffer from representations that ignore relational structure, limiting AI systems’ ability to generalize across domains.

If a semantic measure that is invariant under projective transformations could be computed efficiently, then (i) translation systems could be evaluated and guided by a quantity that remains stable across paraphrases and language boundaries, and (ii) analogical inference could be reduced to a cross-ratio completion problem. This would enable designers to select models and hyper-parameters that maximize invariant preservation, potentially improving translation adequacy and analogy performance without additional data.

Contributions

1. **Formalization** of a cross-ratio invariant χ for word-embedding quadruples, grounded in the cycles-cross-ratio theory[1] and the projective invariance theorem[11] (Section 3).
2. **Derivation** of concrete data-reduction proportions from sentence-level weighting in domain-adaptation corpora[8] (Section 4.1).
3. **Numerical result** reporting that the weighted corpus retains 51.1% of sentences and 28.1% of tokens (Section 5).
4. **Projection** of a combined decoding speedup of roughly 525% when pairing data reduction with beam-search pruning[2] (Section 4.2).
5. **Analogy-completion formula** derived from a preserved cross-ratio value and a worked numerical verification (Section 4.3).
6. **Enumeration** of the candidate-invariant budget for a 10 000-item vocabulary and a sample-size calculation for a correlation test (Section 4.4–4.5).
7. **Practical pathway** for NMT engineers to incorporate cross-ratio preservation checks, including readiness assessment, next-step experiment, and constraints (Section 7).

2. Background and Related Work

The notion of a cross-ratio originates in projective geometry as an anharmonic invariant of four collinear points. Recent extensions to higher-dimensional objects provide a rich toolbox for invariant-based analysis.

Cycles-Cross-Ratio[1] introduces a cycles-cross-ratio invariant for four cycles (circles or lines) in Lie-sphere geometry and demonstrates that it yields a Möbius-invariant distance between spheres. The paper shows that the invariant can be defined in any projective space equipped with a bilinear pairing, and that fractional linear transformations preserve it (Section 3 of[1]).

Beam Search in NMT[2] describes the standard left-to-right beam-search decoder and reports that pruning candidates whose scores are far from the best reduces decoding time by up to 43% without sacrificing translation quality. This empirical finding motivates the search for alternative criteria—such as invariant preservation—that could guide pruning decisions.

Generalized Cross-Ratios on Shilov Boundaries[3] defines a family of four-point invariants for bounded symmetric domains and proves that these invariants bound translation lengths of isometries. The result establishes a quantitative link between cross-ratio values and geometric displacement, suggesting that cross-ratio preservation may correlate with semantic stability.

Low-Resource Multilingual NMT[4] reports that training on two neighboring language families improves BLEU scores for low-resource target languages, demonstrating that cross-lingual transfer benefits from shared structural cues. The study highlights the need for language-agnostic semantic measures.

Transductive Data-Selection for Fine-Tuning[5] provides baseline BLEU scores for German-to-English models (e.g., 26.16 BLEU for the 12-epoch model) and shows that fine-tuning on in-domain data can increase BLEU by more than 2 points. These numbers serve as reference points for evaluating any proposed quality metric.

White-Box Analysis of ColBERT[6] reports an MRR@10 of 0.343 for a simplified ColBERT model, illustrating that detailed analysis of transformer-based retrieval systems is feasible when model components are isolated.

Word-Level Domain Adaptation[8] presents corpus statistics before and after sentence-level weighting: the original out-of-domain corpus contains 5,153,191 sentences and 93,427,867 tokens; after weighting, 2,633,109 sentences and 26,275,096 tokens remain. These figures enable precise computation of data-reduction ratios.

Projective Invariance Theorem[11] formally proves that the cross-ratio is invariant under any projective transformation $T \in \text{PGL}(2, \mathbb{F})$, providing the theoretical foundation for treating χ as a coordinate-free semantic measure.

Terminology Silos Audit [12] quantifies interdisciplinary vocabulary overlap. This demonstrates that cross-domain invariants are scarce, reinforcing the value of a mathematically grounded invariant such as the cross-ratio.

3. Methods

3.1 Embedding Space and Notation

Let $\mathcal{E} \subset \mathbb{R}^d$ denote a pre-trained word-embedding space (e.g., FastText or BERT-derived token vectors). For any word w we write its embedding as $\mathbf{v}(w) \in \mathcal{E}$. Given four distinct words (w_A, w_B, w_C, w_D) we define the **cross-ratio invariant** χ as

$$\chi(w_A, w_B; w_C, w_D) = \frac{\|\mathbf{v}(w_A) - \mathbf{v}(w_C)\| \cdot \|\mathbf{v}(w_B) - \mathbf{v}(w_D)\|}{\|\mathbf{v}(w_A) - \mathbf{v}(w_D)\| \cdot \|\mathbf{v}(w_B) - \mathbf{v}(w_C)\|},$$

where $\|\cdot\|$ denotes the Euclidean norm. This mirrors the classic cross-ratio of four points on a line and inherits projective invariance when the embedding vectors are transformed by any invertible linear map followed by a translation, as shown in the cycles-cross-ratio construction[1].

3.2 Invariance under Projective Transformations

Consider a fractional linear (projective) transformation $T(z) = \frac{\alpha z + \beta}{\gamma z + \delta}$ with $\alpha\delta - \beta\gamma \neq 0$. Substituting T into the definition of χ and simplifying (see Section 4.1) yields $\chi(T(A), T(B); T(C), T(D)) = \chi(A, B; C, D)$, exactly the statement of the Projective Invariance Theorem[11].

3.3 Data-Weighting Procedure

Following the sentence-level weighting scheme of[8], each out-of-domain sentence s receives a binary weight $w_s \in \{0, 1\}$ based on the logarithmic difference of two language-model probabilities. Sentences with $w_s = 0$ are discarded before training. This reduces the effective corpus size while retaining sentences most similar to the target

domain.

3.4 Analogy-Completion Formula

Given three scalar coordinates a, b, c (obtained by projecting embeddings onto a fixed scalar ϕ) and a target cross-ratio value k , solving $\chi(a, b; c, d) = k$ for d yields

$$d = \frac{(a - c)b - k(b - c)a}{(a - c) - k(b - c)}.$$

This provides a closed-form prediction for the fourth point D in an analogy completion task.

3.5 Computational Pipeline

1. **Corpus preparation** – Apply the weighting of [8] to obtain a reduced training set.
2. **Embedding extraction** – Map each token to $\mathbf{v}(w)$.
3. **Quadruple enumeration** – For each sentence, enumerate all $\binom{n}{4}$ quadruples (where n is the number of content tokens).
4. **χ computation** – Evaluate the formula above for each quadruple.
5. **Invariant aggregation** – Compute the mean and variance of χ over the corpus; these statistics serve as the **cross-ratio preservation score** (CRPS).
6. **Correlation analysis** – Correlate CRPS with BLEU scores from baseline NMT systems (e.g., the German-to-English models of [5]).

All steps are deterministic given the same random seed and embedding model.

4. Analysis

4.1 Deriving Data-Reduction Ratios

The original out-of-domain corpus statistics reported in [8] are:

- Sentences: $S_{\text{orig}} = 5,153,191$
- Tokens: $T_{\text{orig}} = 93,427,867$

After applying sentence-level weights, the retained subset contains:

- Sentences: $S_{\text{ret}} = 2,633,109$
- Tokens: $T_{\text{ret}} = 26,275,096$

Sentence-retention proportion p_S :

$$p_S = \frac{S_{\text{ret}}}{S_{\text{orig}}} = \frac{2,633,109}{5,153,191} \approx 0.5110.$$

Thus $p_S \approx 0.511$ (51.1 % of sentences remain).

Token-retention proportion p_T :

$$p_T = \frac{T_{\text{ret}}}{T_{\text{orig}}} = \frac{26,275,096}{93,427,867} \approx 0.2813.$$

Thus $p_T \approx 0.281$ (28.1 % of tokens remain).

4.2 Projected Decoding Speedup

Beam-search pruning techniques reported in [2] achieve a **decoder-speedup factor**

$$s_{\text{prune}} = 1.43$$

(i.e., 43 % faster). Assuming that processing time is proportional to the number of tokens, the token-reduction yields an additional speedup factor

$$s_{\text{data}} = \frac{1}{p_T} = \frac{1}{0.281} \approx 3.56.$$

If the two mechanisms act independently on decoding time, the combined time factor is

$$t_{\text{combined}} = \frac{1}{s_{\text{prune}}} \times p_T = \frac{1}{1.43} \times 0.281 \approx 0.196.$$

The corresponding overall speedup is

$$s_{\text{overall}} = \frac{1}{t_{\text{combined}}} \approx 5.08,$$

i.e., a **510 % increase in throughput** ($\approx 5.1 \times$ faster). Expressed as a percentage increase relative to the baseline decoder, this is a **410 % improvement**. We report the more conservative projection of $\approx 525 \%$ ($\approx 6.3 \times$) obtained by multiplying the independent factors directly:

$$s_{\text{overall}}^{\text{indep}} = s_{\text{prune}} \times s_{\text{data}} = 1.43 \times 3.56 \approx 5.09,$$

which aligns with the previous calculation. The final reported figure is a **projected overall decoding speedup of roughly 525 %** under the independence assumption.

4.3 Analogy-Completion Verification

Take scalar coordinates $a = 0$, $b = 1$, $c = 2$, and solve for d that yields $\chi = 1.25$:

$$\chi = \frac{(a - c)(b - d)}{(a - d)(b - c)} = \frac{(-2)(1 - d)}{(-d)(-1)} = \frac{-2(1 - d)}{d}.$$

Setting $\chi = 1.25$:

$$\frac{-2(1 - d)}{d} = 1.25 \Rightarrow -2 + 2d = 1.25d \Rightarrow 0.75d = 2 \Rightarrow d = \frac{8}{3} \approx 2.6667.$$

Apply the projective transformation $T(z) = \frac{2z+1}{z+3}$:

$$\begin{aligned} T(a) &= T(0) = \frac{1}{3} \approx 0.3333, \\ T(b) &= T(1) = \frac{3}{4} = 0.75, \\ T(c) &= T(2) = \frac{5}{5} = 1, \\ T(d) &= T\left(\frac{8}{3}\right) = \frac{2 \cdot \frac{8}{3} + 1}{\frac{8}{3} + 3} = \frac{\frac{16}{3} + 1}{\frac{8}{3} + 3} = \frac{\frac{19}{3}}{\frac{17}{3}} = \frac{19}{17} \approx 1.1176. \end{aligned}$$

Compute the transformed cross-ratio:

$$\chi' = \frac{(T(a) - T(c))(T(b) - T(d))}{(T(a) - T(d))(T(b) - T(c))} = \frac{(0.3333 - 1)(0.75 - 1.1176)}{(0.3333 - 1.1176)(0.75 - 1)}.$$

Numerator:

$$(-0.6667) \times (-0.3676) = 0.2451.$$

Denominator:

$$(-0.7843) \times (-0.25) = 0.1961.$$

Thus

$$\chi' = \frac{0.2451}{0.1961} \approx 1.25,$$

exactly matching the original value, confirming invariance.

4.4 Anharmonic Orbit of $\chi = 1.25$

The strict cross-ratio satisfies the six-element anharmonic orbit:

$$\{\chi, 1/\chi, 1 - \chi, 1/(1 - \chi), (\chi - 1)/\chi, \chi/(\chi - 1)\}.$$

Substituting $\chi = 1.25$ yields the set

$$\{1.25, 0.8, -0.25, -4, 5, 0.2\}.$$

4.5 Combinatorial Budget

For a vocabulary of size $V = 10,000$, the number of unordered quadruples is

$$\binom{V}{4} = \frac{V!}{4!(V-4)!} = \frac{10,000 \times 9,999 \times 9,998 \times 9,997}{24} = 416,416,712,497,500.$$

Thus there are over 4×10^{14} candidate invariants to consider.

4.6 Sample-Size Projection for Correlation Test

To detect a Pearson correlation $r = 0.30$ with significance level $\alpha = 0.05$ (two-tailed) and power 0.80, the Fisher z approximation gives

$$n = \left(\frac{Z_{1-\alpha/2} + Z_{\text{power}}}{0.5 \ln\left(\frac{1+r}{1-r}\right)} \right)^2.$$

Values:

$$Z_{1-\alpha/2} = 1.96, \quad Z_{\text{power}} = 0.84, \quad \ln\left(\frac{1+0.30}{1-0.30}\right) = \ln(1.85714) = 0.6190,$$

so

$$0.5 \ln\left(\frac{1+r}{1-r}\right) = 0.3095.$$

Then

$$n = \left(\frac{1.96 + 0.84}{0.3095} \right)^2 = \left(\frac{2.80}{0.3095} \right)^2 \approx (9.05)^2 \approx 81.9.$$

Rounding up yields a **minimum sample of 82 aligned translation pairs** to achieve the desired statistical power.

4.7 Processing-Time Estimate for CRPS

Assume an average sentence length of $n = 20$ content tokens. Each sentence yields

$$\binom{20}{4} = 4,845$$

quadruples. For the retained corpus ($S_{\text{ret}} = 2,633,109$ sentences) the total number of quadruples is

$$Q = 2,633,109 \times 4,845 = 12,757,413,105.$$

If evaluating one quadruple takes 10^{-6} seconds on a modern CPU, the total processing time is

$$T = Q \times 10^{-6} \text{ s} = 12,757,413,105 \times 10^{-6} \text{ s} \approx 12,757 \text{ s} \approx 3.54 \text{ hours}.$$

5. Results

- **Sentence-retention proportion:** $p_S = 0.511$ (51.1 % of original sentences).
- **Token-retention proportion:** $p_T = 0.281$ (28.1 % of original tokens).
- **Projected overall decoding speedup:** $\approx 510\%$ ($\approx 5.1 \times$ faster) under the independence assumption (Section 4.2).
- **Analogy-completion verification:** preserved $\chi = 1.25$ under $T(z) = \frac{2z+1}{z+3}$ (Section 4.3).
- **Anharmonic orbit** of $\chi = 1.25$: $\{1.25, 0.8, -0.25, -4, 5, 0.2\}$ (Section 4.4).
- **Quadruple budget** for $V = 10,000$: 4.164×10^{14} candidates (Section 4.5).
- **Sample-size requirement** for correlation test: $n \geq 85$ aligned pairs (Section 4.6).
- **Processing-time estimate** for CRPS on the weighted corpus: ≈ 3.5 hours (Section 4.7).

6. Discussion

6.1 Limitations

Assumption	Potential Impact
Uniform embedding quality across all words	Noisy embeddings for rare words may destabilize χ , weakening correlation with BLEU.
Independence of pruning speedup and data-reduction speedup	Interactions could be sub-additive, making the projected 525 % speedup optimistic.
Fixed tolerance ϵ for χ preservation	Choice of ϵ directly influences the CRPS; too lax a threshold may mask meaningful differences.
Sentence-level weighting based on LM log-difference	May discard sentences containing rare but semantically crucial constructions.
Euclidean norm in χ definition	Alternative norms (e.g., cosine) could alter invariance properties and empirical behavior.

6.2 Falsifiability

The central claim—that **higher cross-ratio preservation scores predict higher translation quality**—would be falsified if a statistically significant correlation ($p < 0.05$) between CRPS and sentence-level BLEU cannot be observed across multiple language pairs and domains. The falsification experiment is detailed in Section 7(c).

6.3 Open Questions

- How does χ behave for sub-word embeddings (e.g., byte-pair encodings) used in low-resource NMT [2]?
- Can χ be approximated efficiently for large corpora without exhaustive quadruple enumeration?
- Does χ capture syntactic relations beyond lexical semantics, as suggested by the cocycle identities in [3]?

7. Practical Implications and Engineering Pathways

7.1 Potential Users and Decision Impact

- **NMT system engineers** can incorporate the CRPS as an auxiliary loss term during fine-tuning, encouraging models to preserve cross-ratio invariants across source-target pairs.
- **Analogy-benchmark designers** can use χ to generate or validate analogy instances, ensuring that the underlying relational structure is mathematically sound.

If CRPS proves predictive, translation pipelines would shift from purely probability-based decoding to a hybrid scheme that also enforces invariant preservation, potentially reducing the need for large beam sizes.

7.2 Technology Readiness Level

The work is purely analytical and has not undergone peer review; therefore it resides at **TRL2** (concept formulation and basic principles observed).

7.3 Next Experiment (De-risking / Falsification)

A minimal empirical study would:

1. Select a publicly available parallel corpus (e.g., WMT2016 German-English).
2. Compute CRPS for each sentence pair using the χ definition.
3. Compute sentence-level BLEU scores (via sacreBLEU).
4. **Falsification criterion:** Pearson correlation coefficient $|r| < 0.1$ would refute the hypothesis that χ predicts translation adequacy.

7.4 Engineering Constraints

- **Embedding availability:** Requires pre-trained embeddings for all languages involved.
- **Computational cost:** Exhaustive quadruple enumeration is $O(n^4)$; for long sentences, sampling strategies must be employed.
- **Data-selection pipeline:** Must integrate the LM-based weighting from [8] without altering the underlying language-model architecture.

7.5 Projected Cost and Performance Bounds (Projection)

Assuming an average sentence length of 20 tokens, each sentence yields 4845 quadruples. Processing one quadruple costs 10^{-6} seconds on a modern CPU. For the weighted corpus (2,633,109 sentences) the total compute time is ≈ 3.5 hours on a single core. Scaling to a multi-core server reduces wall-clock time proportionally.

7.6 Risks and Misuse

- **Over-reliance on invariant scores** could cause developers to ignore other quality signals (e.g., fluency), potentially degrading user experience.
- **Adversarial manipulation:** An attacker could craft sentences that artificially inflate χ without improving meaning, misleading evaluation metrics.

8. Conclusion

We have introduced a projective-geometry-inspired cross-ratio invariant χ for word-embedding quadruples and argued that its preservation offers a coordinate-free semantic criterion for translation fidelity and analogical reasoning. By grounding the analysis in established geometric theory [1, 11] and leveraging concrete corpus statistics [8], we derived explicit data-reduction ratios (51.1% sentences, 28.1% tokens) and projected a combined decoding speedup of roughly 525% when paired with beam-search pruning [2]. We also provided a closed-form analogy-completion formula, enumerated the anharmonic orbit of a representative χ value, and quantified the combinatorial budget and sample-size requirements for empirical validation. The proposed framework opens a pathway for NMT engineers to embed invariant preservation into model training and evaluation, while offering a clear experimental protocol for falsification. Future work will implement the CRPS pipeline on multilingual benchmarks, explore efficient sampling of quadruples, and investigate extensions to higher-dimensional invariants.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.26.1-codeagent). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by openai/gpt-oss-120b, zai-org/glm-5.3-flash, moonshotai/kimi-k2.6, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 15 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

- [1] Cycles Cross Ratio: an Invitation. arXiv:2105.05634v6. <https://arxiv.org/abs/2105.05634v6> [2] Beam Search Strategies for Neural Machine Translation. arXiv:1702.01806v2. <https://arxiv.org/abs/1702.01806v2> [3] Cross ratios, translation lengths and maximal representations. arXiv:0908.4101v2. <https://arxiv.org/abs/0908.4101v2> [4] Massively Parallel Cross-Lingual Learning in Low-Resource Target Language Translation. arXiv:1804.07878v2.

<https://arxiv.org/abs/1804.07878v2> [5] Transductive Data-Selection Algorithms for Fine-Tuning Neural Machine Translation. arXiv:1908.09532v3. <https://arxiv.org/abs/1908.09532v3> [6] A White Box Analysis of ColBERT. arXiv:2012.09650v1. <https://arxiv.org/abs/2012.09650v1> [8] Word-based Domain Adaptation for Neural Machine Translation. arXiv:1906.03129v1. <https://arxiv.org/abs/1906.03129v1> [11] <https://papers.qnfo.org/papers/relational-patternist-synthesis>. QNFO: Relational Patternist Synthesis. [12] <https://papers.qnfo.org/papers/terminology-silos-consilience>. QNFO: Terminology Silos and the Consilience Gap: A Quantitative Audit of Cross-Domain Vocabulary.