NEURAL LANGUAGE DYNAMICS: SHORT-HORIZON STABILITY AND DISSIPATION OF PERTURBATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

We study neural language dynamics: treating the text rollout as a dynamic system, quantized by a pretrained embedding model one token at a time. We formulate this inference-time process as a discrete input–state–output system in which tokens drive an internal Markov state, read out by the embedding model. Our benchmark spans 15 model configurations, 7 languages, parallel sentence-scale FLORES+, and document-scale FineWeb/FineWeb2 trajectories. Local geometry separates sharply by readout: last-token trajectories retain large turning and approximately constant token-scale speed, whereas pooled trajectories slow with prefix length. A predictive gate accepts a reduced or delay coordinate on 267 of 280 shards. Accepted causal coordinates are moderately predictable and usually contract rapidly under a fitted output-space surrogate; pooled coordinates are far more predictable but evolve near the unit timescale, while Chinese causal trajectories form a slower, more metastable intermediate regime. Paired interventions on English FineWeb across three Qwen3-Base sizes expose the central stability distinction. Under a shared suffix, plausible and random one-token edits contract with half-lives of roughly 1.5–5.5 tokens. Once the branches generate freely, their deviation instead returns to the ordinary sampling-noise scale, so shared-input contraction does not imply synchronization under free generation.

1 INTRODUCTION

Language itself is a dynamical system: as a text unfolds, each new token changes the context from which the next is produced. Embedding models provide a quantitative instrument for studying this evolution: applying a frozen model to successive text prefixes traces a sequence of vectors, which records how far the representation moves, how it turns, and how much history its future retains. A dynamical interpretation nevertheless requires care because the exposed embedding is an observation, not necessarily the computational state; similar embeddings may conceal different histories and futures. The prefix encoder is therefore a latent state rollout followed by a readout, not an autonomous update on embeddings, and its adequacy as a predictive coordinate must be tested. Our central stability question is whether a local trajectory difference is erased by subsequent context or persists as the text continues.

We organize the study around three progressively stronger questions, corresponding to spatial geometry, predictive temporal organization, and dynamical stability. First, what geometry do prefix embeddings trace under natural text replay? We measure local motion, accumulated path geometry, and cross-lingual endpoint agreement without treating the embedding as a dynamical state. Second, does a reduced coordinate derived from the current embedding—or a short delay coordinate built from it—predict the next reduced output? Document-level train–test splits, persistence and shuffled controls, and multi-step forecasts provide an operational gate for predictive sufficiency. Third, where that gate passes, what effective output-space dynamics are supported, and does their apparent contraction survive a direct intervention? We fit low-dimensional dynamic-mode-decomposition (DMD) surrogates, examine their spectra and timescales, test fixed-point candidates against observed support, and coarse-grain recurrent or metastable regions. Paired token interventions then compare two branches under a shared suffix and under free generation, separating contraction under common input from synchronization under free generation. These are properties of a fitted

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

output-space system; they do not assign the same fixed points or Lyapunov spectrum to the hidden Transformer state.

Our benchmark combines parallel FLORES+ sentences in seven languages with document-scale FineWeb and FineWeb2 trajectories, advancing one model-specific token at a time with frozen weights and readout depth. The results separate sharply by readout: last-token Qwen trajectories turn strongly, maintain their speed, and lose multi-step predictive power quickly, whereas pooled and API trajectories slow with context, are easier to forecast, and evolve near the unit timescale. Chinese causal trajectories form a more predictable and metastable intermediate regime, while cross-lingual endpoint alignment is high for multilingual embedders but language- and size-dependent for base models. In the English Qwen3-Base intervention panel, a shared suffix contracts both plausible and random edits within a few tokens, whereas free generation returns their deviation to the ordinary sampling-noise scale.

What this work is not:

- **It is not layer dynamics or training dynamics.** Layer dynamics takes network depth as its time index: the state is the activation that changes from layer to layer, often modeled as a flow in depth. Training dynamics takes the optimization step as its time index: the state is the trainable parameters, and the trajectory is its motion under gradient updates across minibatches and epochs.
- **It is not "language dynamics" in the sense of language evolution.** In linguistics and related fields, "language dynamics" usually means how languages change over historical or population time, e.g., sound, grammaticalization, contact, and cultural transmission.

Our contributions are threefold:

- We formulate inference-time prefix embeddings as observations of a token-driven input–state–output system and characterize their **spatial geometry** through a multilingual, two-scale benchmark that distinguishes path structure from endpoint agreement and shows that readout organizes local motion.
- We characterize their **predictive temporal organization** by testing the adequacy of output-derived reduced or delay coordinates and, on the 267 of 280 shards that pass, identifying effective output-space structure through DMD modes and timescales, data-supported fixed points, and recurrent or metastable regions.
- We characterize their **dynamical stability** through surrogate stability analysis and paired token perturbations, separating fitted output-space dynamics from the response of the original representation: shared inputs contract plausible and random edits within 1.5 to 5.5 tokens, while free generation sustains divergence at the sampling-noise scale.

2 Related Work

Language Representation Dynamics. Sequential representation dynamics have been studied through recurrent-state models, fixed-point and contextual analyses, and Transformer belief-state geometry (Elman, 1990; 1991; Maheswaranathan et al., 2019; Maheswaranathan & Sussillo, 2020; Shai et al., 2024). A complementary line measures trajectories of token, prefix, or semantic representations, relating their curvature and velocity to depth, context, entropy, reasoning structure, cross-lingual concepts, retrieval, reading times, literary signals, latent generation, activation time series, and paraphrase cycles (Hosseini & Fedorenko, 2023; Hosseini et al., 2026; King et al., 2026; Zhou et al., 2026; Toro-Hernández et al., 2026; Gromov & Dang, 2023; Tran Minh Quan, 2026; Barenholtz, 2026; Zhang et al., 2025; Ugail & Howard, 2026; Wang et al., 2025). Temporal and systems-oriented work further uses switching linear models, Koopman-inspired residuals, and stochastic linear assumptions to analyze reasoning states, hallucination, safety, and cross-embedding distinguishability (Carson & Reisizadeh, 2025; Wilson & Akrout, 2026; Akrout et al., 2026; Akrout & Wilson, 2026). We combine these perspectives through natural-text prefix measurements under explicit tokenization and readout protocols, while treating fitted output-space dynamics as distinct from unobserved internal states and attractors of the text-generating process.

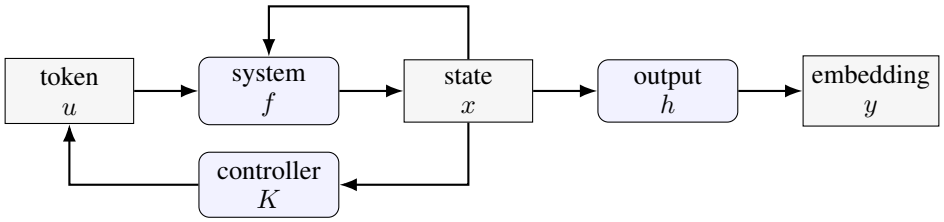

Figure 1: Input–state–output view. At step $t+1$, K samples u_{t+1} from x_t, f maps (x_t, u_{t+1}) to x_{t+1}, and h produces y_{t+1}.

Incremental and Streaming NLP. Incremental NLP studies partial-input interpretation, modular dialogue processing, read–write policies for simultaneous translation, and the behavior of bidirectional encoders under incremental interfaces (Köhn, 2018; Schlangen & Skantze, 2009; Ma et al., 2019; Madureira & Schlangen, 2020). These works motivate our prefix protocol and its incomplete-input caveat; we analyze representation geometry and temporal dependence rather than emission latency or output revision.

Layer and Training Dynamics. Layer dynamics tracks activations through depth, including linguistic information, attention rank collapse, geometric signal propagation, and low-order surrogates for layer interventions (Tenney et al., 2019; Dong et al., 2021; Cowsik et al., 2024; Nettasinghe & Joseph, 2026); training dynamics instead follows parameter updates and phenomena such as the edge of stability (Cohen et al., 2021). Both differ from our token-driven setting, which holds weights and readout depth fixed while advancing the text prefix.

3 Formulation

We view the representation of a text s as a dynamical process that unfolds as progressively longer prefixes become visible. The model-associated tokenizer T maps s to a discrete control sequence. Because tokenizers segment the same surface string differently, both the time index and the vocabulary are model-dependent. We formulate neural language dynamics as a discrete-time input–state–output system on this token stream with a feedback controller. Let

$$u_{1:t} = (u_1, \ldots, u_t) = T(s)_{1:t} \in \mathcal{V}^t \tag{1}$$

denote the token prefix visible through time t, where $u_t \in \mathcal{V}$ is the control input appended at step t and $\mathcal{V}$ is the vocabulary of T.

State update, controller, and output map. The state $x_t \in \mathcal{X}$ is a sufficient statistic of the past for future updates: it contains all information required to process the next token. For a causal Transformer, this may include the key–value cache and any activations required by the readout. For a bidirectional prefix encoder that recomputes every prefix, an exact state may instead include the token history and need not have a fixed-size implementation. For human-generated text, x_t is a latent cognitive state that summarizes the past and guides the next token choice. With frozen model parameters, the one-step state update is

$$x_{t+1} = f(x_t, u_{t+1}), \qquad x_0 = x_\varnothing, \tag{2}$$

where $x_\varnothing$ is the empty-prefix initial state.

A controller (language-producing policy) generates the next input token,

$$u_{t+1} \sim K(\cdot \mid x_t). \tag{3}$$

An autoregressive LM generally has this form: an LM head reads a current last-position activation derived from x_t to parameterize the next-token distribution. While for human-generated text, K is a separate cognitive process that produces the next token. Since u_{t+1} is determined by x_t, the state process is Markovian.

The output is an observable of the underlying state,

$$y_t = h(x_t) \in \mathbb{R}^d, \tag{4}$$

Table 1: Embedding models for language dynamics analysis.

MODEL	TYPE	SIZE	EMBEDDING DIM	MAX INPUT
Perplexity Embed V1	bidirectional	0.6B / 4B	1,024 / 2,560	32,768
Qwen3-Embedding	causal	0.6B / 4B / 8B	1,024 / 2,560 / 4,096	32,768
Qwen3-Base	causal	0.6B / 4B / 8B / 14B	1,024 / 2,560 / 4,096 / 5,120	32,768
BAAI bge-m3	bidirectional	569M	1,024	8,192
OpenAI Text-Embedding 3	undisclosed	Small / Large	1,536 / 3,072	8,192
Voyage-4	bidirectional	Lite / Standard / Large	1,024	32,000

where h is a fixed readout. In general, y_t does not determine x_t; it is a potentially many-to-one observation of the state. The Markovian property of x_t therefore need not be inherited by y_t: states with the same embedding can produce different next embeddings, even under the same next token. Whether an output-derived coordinate is nevertheless predictively adequate is an empirical question addressed in Section 4.3. For a causal last-token readout, the controller may select the next token from the observed embedding; this does not make the embedding sufficient for the subsequent state update. Mean-pooled and dedicated embeddings need not support such output feedback.

Composite embedding map. Rolling out Equation (2) from $x_\varnothing$ along $u_{1:t}$ defines a prefix-to-state map $F : \mathcal{V}^* \to \mathcal{X}$ with $x_t = F(u_{1:t})$. The map reported by embedding model APIs is the composition

$$E \coloneqq h \circ F, \qquad y_t = E(u_{1:t}). \tag{5}$$

Thus E is not the one-step map f; it is f rolled out over a prefix, followed by the output map. For a Transformer, f is generally nonlinear in its continuous state variables because it contains attention, normalization, and nonlinear activations; h may be linear, as in mean pooling, or may itself contain a nonlinear head. In any case, curvature of a single trajectory does not establish nonlinearity of the underlying state update: the external token sequence can bend the trajectory independently of the map's intrinsic nonlinearity. When an implementation exposes only y_t, the quantities in Section 4 are properties of recorded trajectories $\gamma = \{y_t\}$ and, where stated, of paired trajectories under controlled changes to u.

4 Analysis

4.1 Experimental Setup

All analyses use one discrete step per model-specific token, frozen weights, and deterministic inference. For each successive plain-text prefix, open-weight checkpoints are evaluated locally through Hugging Face and proprietary embedding models are queried through OpenRouter, without chat templates, conversational roles, or added instructions. Because embedding dimensions and native scales differ across models, we emphasize scale-robust summaries and within-model contrasts.

Embedding model selection. We select mainstream embedding models based on popularity on OpenRouter, Hugging Face, Ollama, and related communities. Table 1 lists selected models with architecture type, size, embedding dimension, and vendor-documented maximum input length. Gemini Embedding 2 is omitted because its tokenizer is not public, so a prefix cannot be cut on that model's tokens. Two popular open-source models, all-MiniLM-v2 (the L6 and L12 sizes, 256 tokens) and nomic-embed-text-v2-moe (512 tokens), are omitted because their maximum input is too short for document-length rollout. Different models use different types of readouts: mean pooling or CLS for encoders, last-token for causal LMs (`last_token`, `mean_pool`, or `cls` in the implementation). Besides embedding-specialized checkpoints (contrastive or retrieval fine-tuning), we include Qwen3-Base, a decoder-only language-model pretrained checkpoint without embedding fine-tuning, as a control: its last-token readout yields trajectories from a pure next-token pretraining objective rather than from a sentence/document embedding objective.

Dataset selection. The reported analysis uses two public corpora with complementary roles (Table 2). FineWeb (Penedo et al., 2024) and FineWeb2 (Penedo et al., 2025) together supply in-the-

Table 2: Corpora used for language dynamics analysis.

NAME	ORIGINAL ENTRIES	USED ENTRIES	UNIT	CROSS-LANG.
FineWeb	25.9B	2,000	document	—
FineWeb2	∼5B	2,000 / language	document	independent
FLORES+	2,009 / language	2,009 / language	sentence	corresponding

Table 3: Selected languages for language dynamics analysis.

LANGUAGE	FAMILY / BRANCH	MORPHOLOGY
English	Indo-European / Germanic	analytic (residual inflection)
German	Indo-European / Germanic	fusional
French	Indo-European / Romance	fusional
Spanish	Indo-European / Romance	fusional
Modern Standard Arabic	Afro-Asiatic / Semitic	fusional
Mandarin Chinese	Sino-Tibetan / Sinitic	isolating
Japanese	Japonic	agglutinative

wild, document-level web text (filtered Common Crawl): FineWeb covers English, and FineWeb2 covers the remaining language–script pairs. FLORES+ (NLLB Team, 2024) supplies professionally translated, sentence-level parallel text (230 language varieties in version 4.6, maintained by the Open Language Data Initiative), used for controlled same-meaning comparisons across languages and not as training data. We use FLORES+ version 4.6, with the `dev` and `devtest` splits for the seven languages listed below. For the FineWeb and FineWeb2 dataset, the downloaded subsets and the sampling rule are specified in Appendix A.1. Each input is a single plain-text corpus item, tokenized with the associated T; cross-model comparisons therefore use model–tokenizer pairs rather than a shared token clock. If the term "prompt" is used informally, it denotes this entire text rather than a separate instruction or conversation history. Trajectory statistics cover the full recorded rollout, subject to the configured length limit; there is no separate prompt span excluded from these statistics.

Language selection. We select a mainstream, deliberately diverse set of languages so that embedding trajectories can be compared across genealogical families and morphological types under matched protocols (comparable documents on FineWeb2, parallel sentences on FLORES+). Table 3 lists the selected seven languages, one row per language. FLORES+ provides parallel samples for all selected languages; English documents come from FineWeb, while documents in the other six languages come from FineWeb2.

4.2 Observed trajectory geometry

Geometry describes the observed path without assuming that the embedding is a state. By itself, it does not identify linearity, expansion, or an attractor. For $\Delta y_t = y_{t+1} - y_t$, we measure speed $s_t = \|\Delta y_t\|$, turning angle $\theta_t = \arccos\big(\langle \Delta y_t, \Delta y_{t+1}\rangle / (\|\Delta y_t\|\|\Delta y_{t+1}\|)\big)$, path length $L = \sum_{t=1}^{T-1} s_t$, displacement $D = \|y_T - y_1\|$, and tortuosity $\tau = L/D$. To compare unequal embedding dimensions, we emphasize τ, θ_t, and relative speed. The latter assigns each step to one of five equal prefix-position bins and normalizes its median speed by the trajectory-wide median, yielding σ_k for $k \in \{0, \ldots, 4\}$. We compare turning angles with isotropic increments in the same dimension through $\Delta\theta = q_{0.5}(\theta) - q_{0.5}(\theta^{\text{iso}})$; positive values indicate greater turning than the isotropic baseline.

FLORES+ additionally separates path geometry from semantic endpoint agreement. For each language ℓ, we compare the final embeddings of parallel sentences with English using $\cos(y_T^{(\ell)}, y_T^{(\text{en})})$. Euclidean endpoint distances are also recorded, but remain on each model's native scale. Full shard-level summaries are reported in Appendix B.

Local geometry is organized primarily by readout. Every stored geometry shard has positive angle excess, but its magnitude separates model families. Last-token Qwen3 base and embedding models have the largest excess (0.32–0.52 radians across the reported cells), whereas mean-pooled open models and Voyage are much closer to isotropic (0.02–0.17); OpenAI lies between these groups (0.16–0.28). The positional speed profile gives the same separation. For last-token Qwen readouts, σ_k remains close to one from the first to the final fifth of the prefix. For mean-pooled and API embeddings, σ_0 is typically 2–3, while σ_4 falls to 0.55–0.70. Both patterns recur across all seven languages and on sentence-scale FLORES+ as well as document-scale FineWeb.

Global path length depends strongly on the rollout horizon. The corpus-stable local geometry does not imply corpus-stable tortuosity. For last-token Qwen models, median tortuosity is roughly 13–38 on FLORES+ but 4×10^2–1.7×10^3 on FineWeb: displacement grows far more slowly than the accumulated path length. Mean-pooled open models remain much less tortuous, approximately 6–10 on FLORES+ and 17–38 on FineWeb, because their increments shrink with prefix length. Thus tortuosity largely records how long the same local update pattern is applied, whereas angle excess and σ_k expose the update pattern itself.

Cross-lingual endpoints reflect the training objective, not model size alone. Multilingual embedding models and Voyage retain high English–non-English cosine agreement, usually 0.80–0.95 even for Arabic and Japanese. OpenAI embeddings are lower: about 0.75–0.79 on German, French, and Spanish, and 0.49–0.74 on Arabic, Chinese, and Japanese. Qwen3 base models are strongly language- and size-dependent. Their Arabic cosine remains below 0.29 at every tested size; Chinese and Japanese rise from $0.35/0.19$ at 0.6B to $0.91/0.76$ at 8B, but fall again at 14B. The non-monotonicity rules out parameter count alone as an explanation. Together, these results identify the readout and embedding objective as the dominant sources of local geometry, while language and scale mainly modulate endpoint alignment and the length over which that geometry accumulates.

4.3 Predictive sufficiency of reduced coordinates

We next ask whether the observed path admits a low-dimensional predictive coordinate for the corpus-induced transitions. All train–test splits are made by document, preventing adjacent windows from the same trajectory from entering both sides. This is a predictive sufficiency test for an output coordinate, not an attempt to recover the internal state x_t.

Coordinates and predictive gate. PCA is fitted on training documents only, with $z_t = V_r^\top (y_t - \mu)$ and $r \in \{8, 16, 32, 64, 128\}$. We select the smallest rank within 5% of the best held-out predictive error. At each rank we compare an affine current-coordinate predictor $\widehat{z}_{t+1} = Az_t + b$ with persistence, a shuffled-target control, and delay predictors on $[z_t^\top, \ldots, z_{t-\ell+1}^\top]^\top$ for $\ell \in \{1, 2, 4\}$. Writing the selected current or delay coordinate as ξ_t, the empirical question is therefore whether ξ_t adequately predicts the next reduced output z_{t+1}; we do not directly posit a Markov law on the full embedding y_t. Normalized MSE is measured both one step ahead and under autonomous rollout to horizons $\{1, 2, 4, 8\}$. If delay reduces aggregate error by more than 5%, the shortest saturating delay vector, rather than z_t alone, becomes the selected state ξ_t. The gate accepts ξ_t only if it beats persistence at one step and retains NMSE below one at the longest horizon. Projection itself can introduce memory, so passing this gate establishes predictive adequacy only at the selected resolution.

The predictive gate passes 267 of 280 shards. The grid contains 210 FLORES+ shards (separate `dev` and `devtest`) and 70 FineWeb/FineWeb2 shards. Every selected coordinate improves on one-step persistence. All 70 document-scale shards pass the full gate. The 13 failures are confined to Qwen3 base models on sentence-scale FLORES+, in German, French, Spanish, or Arabic; their horizon-8 NMSE is already indistinguishable from one. Consequently, the downstream dynamical interpretation is withheld precisely where the short trajectory contains no useful multi-step forecast.

Predictability and required memory separate the same readout families. Causal base models usually select lag one and rank 16 or 32, but remain only moderately predictable: one-step NMSE is about 0.75, and horizon-8 error returns to approximately one for most European-language and Arabic shards. English remains near 0.95 at horizon eight. Chinese is the systematic exception: its one-step NMSE is about 0.53, its horizon-8 error about 0.87, and 8 of the 12 Chinese base shards

select lag two or four. Qwen3-Embedding models most often select lag two, with occasional lag-one or lag-four choices, and have one-step NMSE near 0.65. Mean-pooled and API models typically select a low-rank, four-lag coordinate. Their one-step NMSE is 0.04–0.09 on FLORES+ and, for the mean-pooled FineWeb runs, often of order 10^{-3}; horizon-8 error remains about 0.20–0.29 on FLORES+ and near 10^{-2} on FineWeb.

These results expose two distinct regimes. Last-token coordinates require little explicit history but have weak multi-step predictability, whereas pooled and API embeddings are highly predictable only after a short delay state absorbs their smoothing-induced memory. FineWeb therefore provides the primary evidence for the effective dynamics below; FLORES+ serves as a shorter-horizon diagnostic.

4.4 Effective dynamics and convergence

Only the 267 accepted coordinates receive a dynamical interpretation. For a rejected shard, the same fit remains a descriptive regression and its spectrum is not interpreted.

Surrogate, fixed point, and convergence criteria. The affine DMD surrogate is $\widehat{\xi}_{t+1} = A\xi_t + b$; for delay coordinates, its lower blocks shift the lagged coordinates exactly. Its modes describe the fitted output-space map, and a stable mode with $0 < |\lambda_j(A)| < 1$ has decay time $-1/\log|\lambda_j(A)|$. The candidate fixed point is $\xi^\star = (I - A)^\dagger b$. We call it data-supported only when it lies within three typical nearest-neighbor distances of a held-out observation. It is called a data-supported surrogate attractor only if it is also locally stable and at least 90% of horizon-64 surrogate rollouts initialized from observed states enter a prescribed relative neighborhood. This terminology applies to the fitted output-space map over sampled support, not to the unobserved internal state or free generation.

Readout separates fast contraction from near-stationarity. Every accepted Qwen base or embedding shard has a locally stable, data-supported fixed-point candidate. For Qwen3 base models, the stronger surrogate-attractor criterion holds on all 43 accepted FLORES+ shards and on 24 of 28 FineWeb shards. The four failures are exactly the FineWeb Chinese shards. There the spectral radius is 0.991–0.995 and only about 0.4–1.9% of horizon-64 rollouts reach the target neighborhood, despite local stability and data support. Away from this exception, base-model radii are substantially smaller (about 0.82 overall and 0.93 in English), with a typical slow decay time of roughly five tokens. Qwen3-Embedding is slower, with characteristic decay times around 11–14 tokens and a less uniform attractor outcome.

In contrast, pooled and API surrogates place their leading modes near the unit circle: radii are typically 0.98–1.01 on FLORES+. The open-weight pooled models evaluated on FineWeb are likewise essentially at one. On stable shards, the corresponding decay estimates range from hundreds to thousands of tokens. Their horizon-64 convergence fraction is therefore zero or negligible, and none satisfies the surrogate-attractor rule. This is not a contradiction with their low predictive error: they predict a slowly moving coordinate accurately.

Autonomous contraction does not describe corpus replay. Even for causal base models whose autonomous surrogate converges, the driven held-out trajectories do not systematically approach $\xi^\star$. Approximately half of observed steps move closer, and the median endpoint distance is about 0.95 of the initial distance. The external token stream continually displaces a state that the autonomous conditional-mean map would contract. Chinese documents sharpen this distinction: they are more predictable, but their near-unit spectral radius removes the short-horizon attracting behavior.

4.5 Recurrence and stability

Fixed-point convergence and regional persistence answer different questions. We therefore coarse-grain each accepted state trajectory into eight regions c_t and, without crossing document boundaries, estimate $P_{ij} = \Pr(c_{t+1} = j \mid c_t = i)$, occupancy, and dwell time. The occupancy-weighted self-transition probability summarizes metastability. These empirical regions need not be invariant and are not relabeled as attractors.

Surrogate Lyapunov spectrum. For completeness, we also apply QR iteration to the constant surrogate Jacobian A. Because the model is affine, this surrogate Lyapunov spectrum restates the DMD spectrum rather than supplying an independent estimate of the internal-state Lyapunov exponents. We report it only for coordinates that pass the predictive gate.

Regional persistence mirrors the spectral timescale. Most non-Chinese causal Qwen shards have metastability around 0.17–0.33: their trajectories move readily among coarse regions. Chinese base shards are markedly stickier, ranging from about 0.30 to 0.84 and reaching 0.69–0.84 on FineWeb. Mean-pooled and API models occupy the opposite regime, with metastability typically 0.86–1.00. Their regions are rarely left within a document. The leading surrogate Lyapunov exponent is negative whenever the fitted surrogate is locally stable, but its magnitude varies with the spectral radius: it is clearly negative for rapidly contracting causal fits and approaches zero for Chinese, pooled, and API fits near the unit circle.

The combined picture distinguishes two mechanisms that can both look persistent in a long trajectory. Causal coordinates mix among regions and would contract quickly without continued token input; their long observed paths are sustained by repeated driving. Pooled and API coordinates instead change slowly and remain in the same coarse region. High self-transition in the latter case indicates a near-unit-timescale coordinate, not convergence to a point attractor.

4.6 Shared-input contraction and free-generation divergence

The preceding analyses are observational. We now intervene on one input token and recompute the trajectory to test the paper's central stability claim: whether contraction inferred for causal models appears under a controlled shared suffix and whether it survives free generation. This validation panel covers English FineWeb and Qwen3-0.6B-, 4B-, and 8B-Base, with 24, 28, and 27 usable documents, respectively, at the 256-token anchor.

Intervention protocol and output metric. We replace u_{t_p} at anchors $t_p \in \{256, 512\}$ and at two positions drawn uniformly from $[256, 512]$. A plausible replacement is drawn from the temperature-one, top-$p = 0.95$ nucleus after removing the original token; a random replacement is uniform over nonspecial vocabulary items. The forced condition then replays the original suffix on both branches. The sampled condition instead generates both suffixes with common random numbers, while independent-seed baseline continuations estimate the ordinary sampling floor. The delayed-release condition forces the original suffix for $\delta \in \{1, 4, 16, 64\}$ steps before allowing both branches to generate. We align branches at $\tau = t - t_p$, compare their common support up to 512 steps, and measure both $\|\delta y_\tau\| = \|y'_\tau - y_\tau\|$ and $1 - \cos(y'_\tau, y_\tau)$. Aggregation first takes the median over branches within a document and then the median and interquartile range across documents.

For a forced pair we additionally report $\lambda^{\mathrm{out}}_{t,k} = k^{-1} \log(\|\delta y_{t+k}\|/\|\delta y_t\|)$. This finite-time rate measures contraction of the full output under one prescribed suffix. It is distinct from both the eigenvalues of the reduced surrogate and a Lyapunov exponent of f.

A shared suffix rapidly erases a one-token edit. Table 4 reports the 256-token anchor; the other positions show the same sign and short timescale. For plausible replacements, the immediate cosine deviation decreases from 0.194 at 0.6B to 0.114 at 4B and 0.036 at 8B. The Euclidean deviation reaches half of its initial value after only 2.0–2.5 tokens, and its late-to-immediate ratio is 0.008–0.013. The corresponding four-step output rates are -0.48, -0.43, and -0.42. Random replacements begin farther away and leave larger residuals, most clearly at 8B, but their four-step rates remain negative (-0.45, -0.38, and -0.35). Late cosine residuals are of order 10^{-5} for plausible edits and 10^{-4} for random edits.

Free generation sustains divergence at the sampling-noise scale. Table 5 reports the direct-release and 64-token delayed-release conditions at the 256-token anchor. In the sampled condition, late Euclidean deviation remains near 100–125, comparable to the upper quartile of ordinary unedited sampling noise, and the decay ratio remains of order one (0.656–1.336) rather than approaching the forced-suffix range. Here divergence means sustained branch separation, not exponential growth or a positive Lyapunov exponent. The delayed-release sweep isolates the transition. For plausible edits at position 256, increasing the shared suffix from $\delta = 1$ to 64 reduces release

Table 4: Shared-suffix contraction after one-token edits on English FineWeb at token 256. Entries are cohort medians.

MODEL	ARM	$1 - \cos(y'_0, y_0)$	DECAY RATIO	HALF-LIFE	$\lambda^{\text{out}}_{0,4}$
Qwen3-0.6B-Base	plausible	0.194	0.008	2.5	-0.48
Qwen3-0.6B-Base	random	0.459	0.017	1.5	-0.45
Qwen3-4B-Base	plausible	0.114	0.013	2.0	-0.43
Qwen3-4B-Base	random	0.514	0.030	2.5	-0.38
Qwen3-8B-Base	plausible	0.036	0.012	2.5	-0.42
Qwen3-8B-Base	random	0.143	0.062	5.5	-0.35

Table 5: Free-generation divergence after direct or 64-token delayed release on the same panel. Entries are cohort-median Euclidean deviations.

MODEL	ARM	FREE LATE	FREE RATIO	RELEASE_{64}	LATE_{64}
Qwen3-0.6B-Base	plausible	101.023	1.095	1.405	100.711
Qwen3-0.6B-Base	random	102.660	0.835	3.891	99.870
Qwen3-4B-Base	plausible	125.156	1.253	2.298	120.388
Qwen3-4B-Base	random	123.916	0.733	10.062	121.080
Qwen3-8B-Base	plausible	101.150	1.336	2.043	100.659
Qwen3-8B-Base	random	104.739	0.656	12.079	101.880

deviation from about 70 to 1.4 at 0.6B, 75 to 2.3 at 4B, and 62 to 2.0 at 8B. After release, however, all three return to the free-generation noise scale.

The intervention therefore supports a narrow but clear conclusion: shared inputs contract both plausible and random one-token perturbations, but free generation sustains divergence at the ordinary sampling-noise scale. This does not show convergence to $\xi^\star$ or contraction under freely diverging token streams. The shared-suffix timescale is comparable to the roughly five-token decay of the causal DMD surrogates, but the agreement is only at the level of timescale: λ^{out} acts in the full embedding under a fixed input, whereas $\log|\lambda(A)|$ belongs to a reduced conditional-mean map. This distinction explains why local forgetting, autonomous surrogate convergence, and persistent corpus-driven trajectories can coexist.

5 Conclusion

Language unfolds dynamically, and embedding models provide a quantitative view by mapping successive text prefixes into vectors. We measured this evolution across 15 model configurations and seven languages, using parallel FLORES+ sentences for controlled cross-lingual comparisons and FineWeb/FineWeb2 documents for longer temporal analysis. We characterized trajectory geometry, tested the predictive adequacy of reduced coordinates, fitted effective output-space dynamics, and directly perturbed individual tokens. The results reveal readout-dependent regimes: last-token representations turn sharply and keep moving, whereas pooled and API representations slow with context. The predictive gate accepts 267 of 280 shards; accepted causal coordinates usually contract rapidly in fitted surrogates, while pooled and API coordinates are easier to predict but remain near the unit timescale. Chinese causal trajectories form a slower, more metastable intermediate regime, and cross-lingual endpoints are more consistently aligned for multilingual embedders than for base models. Most importantly, a shared suffix contracts plausible and random edits within 1.5–5.5 tokens, but free generation restores deviation to the ordinary sampling-noise scale. Shared-input contraction therefore does not imply convergence or synchronization under free generation. Future work should compare subword rollout with linguistic-boundary clocks and study embedders adapted explicitly for streaming input. It should also extend perturbation validation across languages and model families and better disentangle the effects of readout, architecture, and training objective.

AI USE STATEMENT

In this work, we used generative AI tools (LLM-based coding and writing assistants) for implementing parts of the benchmark codebase and running experiments. We also used generative AI tools for searching and retrieving related literature, and drafting portions of the manuscript. We have not used generative AI tools to generate synthetic datasets or prove mathematical claims. We have reviewed all AI-assisted work. The generated code was either manually checked or tested. The retrieved literature was reviewed to confirm the validity and accuracy. The generated text was reviewed and further edited by the authors. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work is about dynamics analysis of well-known publicly-available language models and datasets. It does not involve human subjects, and it does not present any new model or dataset that could be misused.

REPRODUCIBILITY STATEMENT

We have made the following efforts to support reproducibility. The full codebase—including the FLORES+ and FineWeb/FineWeb2 loaders, the Hugging Face and OpenRouter prefix rollouts, and the scripts for post-processing and analysis—is submitted with this paper as Supplementary Material (anonymous during review; public release upon acceptance). Dataset construction, tokenization, and inference settings are specified in Appendices A.1 and A.2, including the pinned corpus commits, the seven-language sample, and the sentences omitted by the Voyage tokenizer round trip. The open-weight rollout, geometry, and surrogate-dynamics procedures are specified in Sections 4.2 to 4.5. More detailed and thorough shard-level results are listed in Appendix B.

REFERENCES

Mohamed Akrout and Dan Wilson. Guarantees on dynamical system distinguishability for LLM token generation. arXiv preprint arXiv:2607.28667, 2026. URL https://arxiv.org/abs/2607.28667.

Mohamed Akrout, Olivera Kotevska, and Dan Wilson. Enforcing LLM safety through DMD-based classification of prompt-response embedding dynamics. arXiv preprint arXiv:2608.19579, 2026. URL https://arxiv.org/abs/2608.19579.

Elan Barenholtz. Trajectory dynamics in language model hidden states predict human processing costs beyond surprisal. arXiv preprint arXiv:2606.05346, 2026. URL https://arxiv.org/abs/2606.05346.

Jack David Carson and Amir Reisizadeh. A statistical physics of language model reasoning. arXiv preprint arXiv:2506.04374, 2025. URL https://arxiv.org/abs/2506.04374.

Jeremy M Cohen, Simran Kaur, Yuanzhi Li, J Zico Kolter, and Ameet Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability. In International Conference on Learning Representations, 2021.

Aditya Cowsik, Tamra Nebabu, Xiao-Liang Qi, and Surya Ganguli. Geometric dynamics of signal propagation predict trainability of transformers. arXiv preprint arXiv:2403.02579, 2024. URL https://arxiv.org/abs/2403.02579.

Yihe Dong, Jean-Baptiste Cordonnier, and Andreas Loukas. Attention is not all you need: pure attention loses rank doubly exponentially with depth. In Proceedings of the 38th International Conference on Machine Learning, volume 139 of Proceedings of Machine Learning Research, pp. 2793–2803, 2021. URL https://proceedings.mlr.press/v139/dong21a.html.

Jeffrey L. Elman. Finding structure in time. Cognitive Science, 14(2):179–211, 1990. doi: 10.1207/s15516709cog1402_1.

Jeffrey L. Elman. Distributed representations, simple recurrent networks, and grammatical structure. Machine Learning, 7:195–225, 1991. doi: 10.1023/A:1022699029236.

Vasilii A. Gromov and Quynh Nhu Dang. Semantic and sentiment trajectories of literary masterpieces. Chaos, Solitons & Fractals, 175:113934, 2023. doi: 10.1016/j.chaos.2023.113934.

Eghbal Hosseini and Evelina Fedorenko. Large language models implicitly learn to straighten neural sentence trajectories to construct a predictive representation of natural language. In Advances in Neural Information Processing Systems, volume 36, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/88dddaf430b5bc38ab8228902bb61821-Abstract-Conference.html.

Eghbal A. Hosseini, Yuxuan Li, Yasaman Bahri, Declan Campbell, and Andrew Kyle Lampinen. Context structure reshapes the representational geometry of language models. arXiv preprint arXiv:2601.22364, 2026. URL https://arxiv.org/abs/2601.22364.

Jack King, Evelina Fedorenko, and Eghbal A. Hosseini. Representational curvature modulates behavioral uncertainty in large language models. arXiv preprint arXiv:2604.23985, 2026. URL https://arxiv.org/abs/2604.23985.

Arne Köhn. Incremental natural language processing: Challenges, strategies, and evaluation. In Proceedings of the 27th International Conference on Computational Linguistics, pp. 2990–3003, 2018. URL https://aclanthology.org/C18-1253/.

Mingbo Ma, Liang Huang, Hao Xiong, Renjie Zheng, Kaibo Liu, Baigong Zheng, Chuanqiang Zhang, Zhongjun He, Hairong Liu, Xing Li, Hua Wu, and Haifeng Wang. STACL: Simultaneous translation with implicit anticipation and controllable latency using prefix-to-prefix framework. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 2019. URL https://aclanthology.org/P19-1289/.

Brielen Madureira and David Schlangen. Incremental processing in the age of non-incremental encoders: An empirical assessment of bidirectional models for incremental NLU. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, pp. 357–374, 2020. URL https://aclanthology.org/2020.emnlp-main.26/.

Niru Maheswaranathan and David Sussillo. How recurrent networks implement contextual processing in sentiment analysis. In Proceedings of the 37th International Conference on Machine Learning, volume 119 of Proceedings of Machine Learning Research, pp. 6608–6619, 2020. URL https://proceedings.mlr.press/v119/maheswaranathan20a.html.

Niru Maheswaranathan, Alex H. Williams, Matthew D. Golub, Surya Ganguli, and David Sussillo. Reverse engineering recurrent networks for sentiment classification reveals line attractor dynamics. In Advances in Neural Information Processing Systems, volume 32, 2019. URL https://arxiv.org/abs/1906.10720.

Buddhika Nettasinghe and Geethu Joseph. As language models scale, low-order linear depth dynamics emerge. arXiv preprint arXiv:2603.12541, 2026. URL https://arxiv.org/abs/2603.12541.

NLLB Team. Scaling neural machine translation to 200 languages. Nature, 630:841–846, 2024.

Guilherme Penedo, Hynek Kydlíček, Loubna Ben allal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. The FineWeb datasets: Decanting the web for the finest text data at scale. arXiv preprint arXiv:2406.17557, 2024.

Guilherme Penedo, Hynek Kydlíček, Vinko Sabolčec, Bettina Messmer, Negar Foroutan, Amir Hossein Kargaran, Colin Raffel, Martin Jaggi, Leandro Von Werra, and Thomas Wolf. FineWeb2: One pipeline to scale them all — adapting pre-training data processing to every language. arXiv preprint arXiv:2506.20920, 2025.

David Schlangen and Gabriel Skantze. A general, abstract model of incremental dialogue processing. In Proceedings of the 12th Conference of the European Chapter of the ACL, pp. 710–718, 2009. URL https://aclanthology.org/E09-1081/.

Adam S Shai, Sarah E Marzen, Lucas Teixeira, Alexander Gietelink Oldenziel, and Paul M Riechers. Transformers represent belief state geometry in their residual stream. In Advances in Neural Information Processing Systems, volume 37, 2024.

Ian Tenney, Dipanjan Das, and Ellie Pavlick. Bert rediscovers the classical nlp pipeline. In Proceedings of the 57th annual meeting of the association for computational linguistics, pp. 4593–4601, 2019.

Felipe D. Toro-Hernández, Jesuino Vieira Filho, and Rodrigo M. Cabral-Carvalho. Characterizing human semantic navigation in concept production as trajectories in embedding space. In International Conference on Learning Representations, 2026. URL https://openreview.net/forum?id=QQVmIR97sf.

Tran Minh Quan. Reading a legal question word by word: Embedding trajectories of 2,144 vietnamese legal headlines. arXiv preprint arXiv:2609.08372, 2026. URL https://arxiv.org/abs/2609.08372.

Hassan Ugail and Newton Howard. Dynamical systems analysis reveals functional regimes in large language models. arXiv preprint arXiv:2601.11622, 2026.

Zhilin Wang, Yafu Li, Jianhao Yan, Yu Cheng, and Yue Zhang. Unveiling attractor cycles in large language models: A dynamical systems view of successive paraphrasing. arXiv preprint arXiv:2502.15208, 2025. URL https://arxiv.org/abs/2502.15208.

Dan Wilson and Mohamed Akrout. Low-cost black-box detection of LLM hallucinations via dynamical system prediction. arXiv preprint arXiv:2605.05134, 2026. URL https://arxiv.org/abs/2605.05134.

Yukun Zhang, Qi Dong, and Mengkang Li. Latent trajectory dynamics in large language models: A manifold evolution framework with empirical validation. arXiv preprint arXiv:2505.20340, 2025. URL https://arxiv.org/abs/2505.20340v3. Version 3, revised May 2026.

Yufa Zhou, Yixiao Wang, Xunjian Yin, Shuyan Zhou, and Anru Zhang. The geometry of reasoning: Flowing logics in representation space. In International Conference on Learning Representations, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/7644353d580a9e027e0069d6480d971b-Abstract-Conference.html.

A Detailed Experimental Setups

A.1 Dataset Construction

The experiments use two public corpora and only the seven languages in Table 3: English, German, French, Spanish, Modern Standard Arabic, Mandarin Chinese, and Japanese. No other FineWeb2 language–script pair and no other FLORES+ variety is downloaded. A download manifest records the pinned dataset revision, the row count of each file, and a SHA-256 checksum, so the same files can be checked after the fact.

FLORES+. The sentence data are version 4.6 of `openlanguagedata/flores_plus`, commit `5fec6c13f9e5a4db2f745d4ec0d7c9721ddc4f06`. For each of the seven languages the loader keeps the official `dev` and `devtest` JSONL files unchanged: 997 and 1,012 sentences. The language codes are `eng_Latn`, `deu_Latn`, `fra_Latn`, `spa_Latn`, `arb_Arab`, `cmn_Hans`, and `jpn_Jpan`. Parallel sentences share the pair `(split, id)` across those files. The release is CC-BY-SA-4.0 and is used for evaluation only.

FineWeb and FineWeb2. English is taken from `HuggingFaceFW/fineweb`, configuration `sample-10BT`, commit `9bb295ddab0e05d785b879661af7260fed5140fc`. German, French, Spanish, Arabic, Chinese, and Japanese are taken from `HuggingFaceFW/fineweb-2`, commit `af9c13333eb981300149d5ca60a8e9d659b276b9`, under the configurations `deu_Latn`, `fra_Latn`, `spa_Latn`, `arb_Arab`, `cmn_Hani`, and `jpn_Jpan`. The Chinese configuration is Han script, which is broader than the simplified Mandarin `cmn_Hans` used for the

same language in FLORES+. Each language is read as a stream in the order of that commit. A document is kept when its `language_score` is at least 0.99 and its GPT-2 token count lies in $[1000, 8000]$. English uses the `token_count` field published with FineWeb. FineWeb2 does not publish that field, so the count is computed with the GPT-2 tokenizer and special tokens left off; a text shorter than 250 characters or longer than 120,000 characters is rejected before that encode. The first 2,000 documents that pass are the sample. The scan stops after 5,000,000 rows if fewer matches have appeared; every language in the stored run reached 2,000. Membership is the prefix of the filtered stream, so a later draw from the same commit and the same filters reproduces the same 2,000 documents.

A.2 Tokenization and Embedding Model Inference

Open-weight models. Every open-weight model in Table 1 is loaded from a local Hugging Face checkpoint and run on a single NVIDIA A100-SXM4 GPU with 80 GB of memory. The loader reads that local snapshot with downloads disabled. Perplexity Embed V1 uses the modeling code shipped inside its checkpoint. Weights and the stored readout are `float32`. TF32 is disabled for CUDA matrix multiplication and for cuDNN, and attention uses the eager implementation, in line with the deterministic inference of Section 4.1. The recorded vector is the last-layer hidden state: the last token for Qwen3-Embedding and Qwen3-Base, and the mean over positions for bge-m3 and Perplexity Embed V1. Causal models advance one token at a time and reuse the key–value cache. Bidirectional models run a separate forward pass on each prefix $u_{1:t}$. Tokenization uses that model's own tokenizer with special tokens enabled, so a BOS or CLS token inserted by the tokenizer is part of u. The input is the plain corpus string. The length cap is the vendor maximum in Table 1: 32,768 tokens for the Qwen3 and Perplexity models, and 8,192 tokens for bge-m3. A longer document is cut at that cap. The FLORES+ sentences sit well below every cap.

Proprietary models. OpenAI Text-Embedding 3 (Small and Large) and Voyage-4 (lite, standard, and large) are queried through the OpenRouter embeddings endpoint. The API runs cover FLORES+ only: the `dev` and `devtest` splits of the seven languages in Table 3, 997 and 1,012 sentences. Requests use the vendor default dimension, 1,536 and 3,072 for the two OpenAI models and 1,024 for all three Voyage models, with the full vector retained. OpenAI Text-Embedding 3 Small and Large share the `cl100k_base` tokenizer. Each OpenAI request sends token-id arrays, so the prefix is exactly the ids produced locally, including a token that ends inside a UTF-8 character. The three Voyage sizes share one tokenizer. Special tokens are left off for both APIs. Prefixes go out in batches of 64, with up to 64 requests in flight.

Voyage sentences left out. The Voyage embedding API accepts only strings. Each trajectory step is a token prefix $u_{1:t}$, so that prefix has to be converted back to a string before the request. The conversion can fail. A token may end inside a UTF-8 character; decoding the prefix then inserts U+FFFD, and encoding the resulting string does not recover the original token ids. If any prefix of a sentence fails this check, including the full sentence, the sentence is omitted before any request. The store keeps a sentence only when every decoded prefix encodes back to the same ids. The three Voyage sizes share the tokenizer, so they miss the same sentences. Table 6 counts them. English is complete on both splits. Across `dev` and `devtest`, the omitted counts are 2 (German), 64 (French), 3 (Spanish), 256 (Arabic), 386 (Chinese), and 1,383 (Japanese), 2,094 sentences in total. Each Voyage model therefore stores 11,969 of the 14,063 sentences. The OpenAI runs contain every sentence, because the token-id request does not pass through this decode step.

B Detailed Experimental Results

The tables below provide a complete shard-level inventory of the measurements most directly supporting Sections 4.2 to 4.6. To keep the appendix readable, they retain the quantities needed to audit the main conclusions rather than every intermediate diagnostic or hyperparameter candidate. The supplementary machine-readable outputs retain per-document or per-branch records, the full rank–lag–horizon grid, fixed-point diagnostics, region statistics, and perturbation curves. FLORES+ surrogate rows are separate for `dev` and `devtest`. FineWeb and FineWeb2 rows are marked `fw`. Geometry rows pool the two FLORES+ splits because they report descriptive medians. A surrogate

Table 6: Voyage sentences omitted from the FLORES+ store. The three Voyage sizes share these counts. Dev has 997 sentences and devtest has 1,012. Share is the omitted count as a percentage of the 2,009 sentences in that language. The total row uses all 14,063 sentences.

LANGUAGE	DEV	DEVTEST	TOTAL	SHARE
English	0	0	0	0%
German	2	0	2	0.10%
French	38	26	64	3.19%
Spanish	2	1	3	0.15%
Arabic	124	132	256	12.74%
Chinese	195	191	386	19.21%
Japanese	671	712	1,383	68.84%
Total	1,032	1,062	2,094	14.89%

quantity is left blank when the predictive gate rejects that shard, in which case Section 4.4 assigns it no dynamical interpretation.

B.1 Observed trajectory geometry

Table 7: Complete shard-level geometry inventory. All numeric entries are medians: Tort. is path length divided by endpoint displacement; Excess is turning angle minus the isotropic median; Path, Disp., and Speed are on the model's native Euclidean scale; Jump is the between-sentence endpoint distance available on FLORES+. Rows marked `flores` pool dev and devtest, while `fw` denotes FineWeb or FineWeb2. Cross-model comparisons should use Tort. and Excess rather than native-scale columns.

Model	Shard	Tort.	Excess	Path	Disp.	Speed	Jump
Qwen3-0.6B-Base	flores/en	17.39	0.388	2221.7	127.0	89.6	53.0
Qwen3-0.6B-Base	flores/de	32.85	0.399	4538.6	137.7	119.3	40.9
Qwen3-0.6B-Base	flores/fr	33.64	0.389	4667.0	133.9	116.1	42.7
Qwen3-0.6B-Base	flores/es	30.63	0.375	4285.2	135.7	108.9	43.2
Qwen3-0.6B-Base	flores/ar	30.56	0.416	4887.2	160.6	131.4	38.0
Qwen3-0.6B-Base	flores/zh	18.10	0.349	1807.3	97.9	74.5	47.1
Qwen3-0.6B-Base	flores/ja	24.44	0.397	3641.6	147.3	99.4	43.0
Qwen3-4B-Base	flores/en	16.14	0.389	2658.9	160.0	107.2	63.0
Qwen3-4B-Base	flores/de	31.32	0.371	5109.9	163.6	131.6	51.6
Qwen3-4B-Base	flores/fr	30.12	0.367	5134.2	164.9	128.0	50.8
Qwen3-4B-Base	flores/es	28.59	0.357	4768.0	164.3	122.0	49.8
Qwen3-4B-Base	flores/ar	35.35	0.386	5890.7	166.7	154.1	47.4
Qwen3-4B-Base	flores/zh	14.69	0.356	2013.8	135.7	83.9	53.4
Qwen3-4B-Base	flores/ja	23.16	0.374	3811.1	166.2	104.0	53.1
Qwen3-8B-Base	flores/en	17.15	0.393	2418.7	137.8	93.2	54.1
Qwen3-8B-Base	flores/de	38.15	0.379	4988.2	131.8	119.9	50.3
Qwen3-8B-Base	flores/fr	35.97	0.430	5187.1	143.0	120.8	47.6
Qwen3-8B-Base	flores/es	34.42	0.463	5018.2	146.0	120.6	46.1
Qwen3-8B-Base	flores/ar	28.67	0.408	4297.2	149.0	112.2	45.3
Qwen3-8B-Base	flores/zh	15.04	0.394	1942.7	127.8	75.9	44.4
Qwen3-8B-Base	flores/ja	26.90	0.383	3375.5	124.8	88.4	57.8
Qwen3-14B-Base	flores/en	23.33	0.486	4982.2	212.9	188.6	110.0
Qwen3-14B-Base	flores/de	32.66	0.399	5704.3	175.5	143.7	62.1
Qwen3-14B-Base	flores/fr	29.38	0.432	6062.0	207.6	152.0	64.1
Qwen3-14B-Base	flores/es	32.64	0.435	5707.6	174.1	150.9	63.9
Qwen3-14B-Base	flores/ar	34.54	0.397	4545.5	131.5	119.4	49.0
Qwen3-14B-Base	flores/zh	13.57	0.411	2589.0	192.3	106.1	64.1

Table 7 continued

Model	Shard	Tort.	Excess	Path	Disp.	Speed	Jump
Qwen3-14B-Base	flores/ja	23.00	0.405	3500.0	151.0	93.0	58.3
Qwen3-Embedding-0.6B	flores/en	19.10	0.360	2501.8	130.7	100.1	98.1
Qwen3-Embedding-0.6B	flores/de	29.25	0.331	4258.7	146.4	110.7	118.4
Qwen3-Embedding-0.6B	flores/fr	29.14	0.331	4228.7	146.5	105.2	116.0
Qwen3-Embedding-0.6B	flores/es	28.04	0.320	4062.3	146.0	104.4	117.6
Qwen3-Embedding-0.6B	flores/ar	32.72	0.355	4614.2	141.1	116.0	116.0
Qwen3-Embedding-0.6B	flores/zh	18.11	0.352	2400.6	133.3	99.9	97.1
Qwen3-Embedding-0.6B	flores/ja	25.68	0.325	3781.5	147.3	108.4	119.8
Qwen3-Embedding-4B	flores/en	18.58	0.392	2923.9	156.9	114.6	130.6
Qwen3-Embedding-4B	flores/de	28.57	0.352	4752.9	166.4	122.1	141.9
Qwen3-Embedding-4B	flores/fr	29.80	0.372	4942.5	165.8	123.1	137.3
Qwen3-Embedding-4B	flores/es	27.64	0.362	4638.8	167.5	119.2	143.0
Qwen3-Embedding-4B	flores/ar	31.78	0.375	5294.8	166.8	129.4	140.9
Qwen3-Embedding-4B	flores/zh	17.75	0.373	2807.6	157.7	114.1	131.6
Qwen3-Embedding-4B	flores/ja	26.62	0.353	4358.4	163.3	122.5	140.9
Qwen3-Embedding-8B	flores/en	12.85	0.404	3268.1	255.1	126.0	123.3
Qwen3-Embedding-8B	flores/de	21.18	0.384	5194.0	250.7	132.5	135.1
Qwen3-Embedding-8B	flores/fr	22.16	0.394	5478.9	252.3	135.0	131.5
Qwen3-Embedding-8B	flores/es	20.48	0.387	5088.2	250.3	130.8	136.4
Qwen3-Embedding-8B	flores/ar	24.38	0.404	5274.4	216.7	130.4	136.7
Qwen3-Embedding-8B	flores/zh	14.37	0.393	2937.4	205.6	120.3	130.9
Qwen3-Embedding-8B	flores/ja	20.77	0.381	4306.3	207.8	122.5	134.4
bge-m3	flores/en	6.32	0.124	180.4	28.6	4.3	17.9
bge-m3	flores/de	6.87	0.124	196.5	28.6	3.9	18.0
bge-m3	flores/fr	7.05	0.120	202.2	28.7	3.5	17.7
bge-m3	flores/es	7.16	0.146	205.8	28.7	4.0	17.8
bge-m3	flores/ar	7.00	0.124	199.7	28.6	4.0	18.2
bge-m3	flores/zh	6.02	0.107	171.2	28.5	4.2	17.6
bge-m3	flores/ja	6.39	0.095	182.1	28.5	3.8	18.3
pplx-embed-v1-0.6b	flores/en	5.80	0.116	62.4	10.5	1.5	5.9
pplx-embed-v1-0.6b	flores/de	7.65	0.126	84.1	10.3	1.2	5.2
pplx-embed-v1-0.6b	flores/fr	7.80	0.114	80.9	10.4	1.1	5.1
pplx-embed-v1-0.6b	flores/es	7.63	0.107	76.3	10.1	1.1	5.1
pplx-embed-v1-0.6b	flores/ar	8.66	0.128	74.0	8.3	1.0	4.9
pplx-embed-v1-0.6b	flores/zh	5.91	0.109	51.1	8.4	1.3	5.0
pplx-embed-v1-0.6b	flores/ja	7.16	0.094	61.5	8.4	1.0	4.8
pplx-embed-v1-4b	flores/en	6.19	0.117	125.0	20.0	4.0	16.0
pplx-embed-v1-4b	flores/de	9.40	0.113	184.8	19.6	3.6	15.9
pplx-embed-v1-4b	flores/fr	9.02	0.107	170.0	18.9	3.2	15.6
pplx-embed-v1-4b	flores/es	8.40	0.098	163.0	19.5	3.0	14.4
pplx-embed-v1-4b	flores/ar	10.10	0.118	174.5	17.4	3.1	14.3
pplx-embed-v1-4b	flores/zh	6.86	0.118	118.3	17.1	3.7	14.6
pplx-embed-v1-4b	flores/ja	8.79	0.108	149.5	17.1	3.1	14.1
text-embedding-3-small	flores/en	6.85	0.189	8.8	1.3	0.3	1.1
text-embedding-3-small	flores/de	9.28	0.182	12.2	1.3	0.2	1.0
text-embedding-3-small	flores/fr	9.79	0.187	13.0	1.3	0.3	1.0
text-embedding-3-small	flores/es	9.64	0.200	12.6	1.3	0.3	1.0
text-embedding-3-small	flores/ar	15.27	0.164	20.4	1.3	0.2	1.2
text-embedding-3-small	flores/zh	10.10	0.194	13.0	1.3	0.2	1.1
text-embedding-3-small	flores/ja	11.46	0.194	14.7	1.3	0.2	1.1
text-embedding-3-large	flores/en	8.09	0.234	10.6	1.3	0.4	1.0
text-embedding-3-large	flores/de	13.17	0.272	17.5	1.3	0.4	1.0
text-embedding-3-large	flores/fr	13.42	0.277	18.2	1.4	0.4	1.0
text-embedding-3-large	flores/es	13.25	0.282	17.5	1.3	0.4	1.0
text-embedding-3-large	flores/ar	25.14	0.263	33.9	1.4	0.4	1.1
text-embedding-3-large	flores/zh	13.68	0.268	18.1	1.3	0.3	1.1
text-embedding-3-large	flores/ja	15.60	0.243	20.7	1.3	0.3	1.1

Table 7 continued

Model	Shard	Tort.	Excess	Path	Disp.	Speed	Jump
voyage-4-lite	flores/en	5.51	0.087	6.5	1.2	0.2	0.9
voyage-4-lite	flores/de	7.79	0.108	9.5	1.2	0.2	1.0
voyage-4-lite	flores/fr	7.49	0.100	9.1	1.2	0.1	0.9
voyage-4-lite	flores/es	7.66	0.096	9.4	1.2	0.2	1.0
voyage-4-lite	flores/ar	8.23	0.126	9.9	1.2	0.2	1.0
voyage-4-lite	flores/zh	5.34	0.093	6.3	1.2	0.2	0.9
voyage-4-lite	flores/ja	6.23	0.088	7.1	1.1	0.2	0.9
voyage-4	flores/en	5.02	0.039	6.1	1.2	0.2	0.9
voyage-4	flores/de	6.93	0.043	8.8	1.3	0.1	0.9
voyage-4	flores/fr	6.68	0.033	8.3	1.3	0.1	0.9
voyage-4	flores/es	6.75	0.031	8.6	1.3	0.1	0.9
voyage-4	flores/ar	7.31	0.055	9.1	1.3	0.1	1.0
voyage-4	flores/zh	5.00	0.051	6.0	1.2	0.2	0.9
voyage-4	flores/ja	5.69	0.035	6.7	1.2	0.2	0.9
voyage-4-large	flores/en	5.83	0.109	7.9	1.3	0.2	1.0
voyage-4-large	flores/de	7.28	0.119	9.9	1.4	0.1	1.0
voyage-4-large	flores/fr	7.19	0.119	9.8	1.4	0.1	1.0
voyage-4-large	flores/es	7.30	0.106	9.9	1.4	0.2	1.0
voyage-4-large	flores/ar	8.21	0.169	11.4	1.4	0.2	1.1
voyage-4-large	flores/zh	5.64	0.126	7.6	1.4	0.2	1.0
voyage-4-large	flores/ja	6.39	0.118	8.6	1.4	0.2	1.0
Qwen3-0.6B-Base	fw/en	1074.40	0.410	146346.5	134.9	89.2	–
Qwen3-0.6B-Base	fw/de	1024.41	0.405	141790.7	139.3	116.0	–
Qwen3-0.6B-Base	fw/fr	1113.40	0.391	156498.3	139.1	113.1	–
Qwen3-0.6B-Base	fw/es	1009.25	0.377	143137.0	140.7	107.3	–
Qwen3-0.6B-Base	fw/ar	543.58	0.417	85853.1	162.5	129.8	–
Qwen3-0.6B-Base	fw/zh	508.15	0.341	61445.6	122.5	71.1	–
Qwen3-0.6B-Base	fw/ja	669.36	0.388	94899.2	145.9	93.7	–
Qwen3-4B-Base	fw/en	1017.01	0.411	176081.7	172.6	104.2	–
Qwen3-4B-Base	fw/de	966.96	0.391	164689.0	170.6	132.2	–
Qwen3-4B-Base	fw/fr	1037.90	0.381	179527.5	171.9	127.7	–
Qwen3-4B-Base	fw/es	944.76	0.371	164141.1	171.8	121.4	–
Qwen3-4B-Base	fw/ar	592.22	0.390	105674.7	174.0	155.8	–
Qwen3-4B-Base	fw/zh	414.26	0.355	70548.5	171.8	80.7	–
Qwen3-4B-Base	fw/ja	582.99	0.380	107056.9	184.0	105.3	–
Qwen3-8B-Base	fw/en	1074.61	0.427	156281.5	143.7	88.9	–
Qwen3-8B-Base	fw/de	1097.01	0.390	155783.8	142.4	119.2	–
Qwen3-8B-Base	fw/fr	1205.46	0.441	177466.7	148.6	119.1	–
Qwen3-8B-Base	fw/es	1104.71	0.474	169182.2	152.1	119.7	–
Qwen3-8B-Base	fw/ar	477.01	0.415	77716.5	163.3	114.1	–
Qwen3-8B-Base	fw/zh	477.90	0.395	67883.7	142.1	72.6	–
Qwen3-8B-Base	fw/ja	623.05	0.392	92381.2	147.7	88.3	–
Qwen3-14B-Base	fw/en	1679.55	0.521	357946.3	213.6	213.5	–
Qwen3-14B-Base	fw/de	900.66	0.402	183658.3	202.8	144.7	–
Qwen3-14B-Base	fw/fr	972.98	0.438	216242.4	223.6	157.8	–
Qwen3-14B-Base	fw/es	980.28	0.443	198840.1	200.4	154.5	–
Qwen3-14B-Base	fw/ar	546.10	0.404	83474.2	143.9	122.6	–
Qwen3-14B-Base	fw/zh	433.94	0.412	84824.3	197.1	94.7	–
Qwen3-14B-Base	fw/ja	539.22	0.403	100443.6	188.8	95.5	–
Qwen3-Embedding-0.6B	fw/en	1154.56	0.386	162167.5	139.7	103.5	–
Qwen3-Embedding-0.6B	fw/de	1034.81	0.372	138382.0	132.2	115.7	–
Qwen3-Embedding-0.6B	fw/fr	1149.20	0.368	148337.7	129.4	109.9	–
Qwen3-Embedding-0.6B	fw/es	1050.54	0.355	139762.5	133.0	109.1	–
Qwen3-Embedding-0.6B	fw/ar	614.67	0.377	82121.5	133.0	118.9	–
Qwen3-Embedding-0.6B	fw/zh	638.74	0.370	83440.1	129.3	104.2	–
Qwen3-Embedding-0.6B	fw/ja	849.38	0.372	107006.8	126.5	114.5	–
Qwen3-Embedding-4B	fw/en	1177.26	0.400	190788.3	160.0	118.9	–

Table 7 continued

Model	Shard	Tort.	Excess	Path	Disp.	Speed	Jump
Qwen3-Embedding-4B	fw/de	993.76	0.380	157240.9	158.7	133.0	–
Qwen3-Embedding-4B	fw/fr	1122.02	0.393	175763.3	157.2	131.7	–
Qwen3-Embedding-4B	fw/es	1015.08	0.382	161490.7	162.4	126.4	–
Qwen3-Embedding-4B	fw/ar	563.50	0.396	93139.3	164.4	134.8	–
Qwen3-Embedding-4B	fw/zh	614.75	0.381	95168.0	154.9	117.4	–
Qwen3-Embedding-4B	fw/ja	807.36	0.384	124580.6	154.6	130.7	–
Qwen3-Embedding-8B	fw/en	851.04	0.415	214234.5	257.0	126.1	–
Qwen3-Embedding-8B	fw/de	714.46	0.400	165302.7	243.4	138.4	–
Qwen3-Embedding-8B	fw/fr	851.46	0.410	194742.5	238.9	141.5	–
Qwen3-Embedding-8B	fw/es	754.24	0.402	174830.5	242.3	135.8	–
Qwen3-Embedding-8B	fw/ar	420.84	0.416	89322.9	215.1	130.9	–
Qwen3-Embedding-8B	fw/zh	476.73	0.402	95702.5	203.0	119.3	–
Qwen3-Embedding-8B	fw/ja	581.55	0.404	117161.7	201.5	124.7	–
bge-m3	fw/en	28.91	0.157	806.6	27.9	0.2	–
bge-m3	fw/de	23.30	0.149	642.7	27.6	0.3	–
bge-m3	fw/fr	37.94	0.143	642.1	17.0	0.2	–
bge-m3	fw/es	23.31	0.142	648.7	28.0	0.2	–
bge-m3	fw/ar	16.57	0.145	454.3	27.6	0.3	–
bge-m3	fw/zh	18.87	0.133	516.2	27.5	0.3	–
bge-m3	fw/ja	20.45	0.148	563.2	27.7	0.3	–
pplx-embed-v1-0.6b	fw/en	31.43	0.046	223.3	4.8	0.0	–
pplx-embed-v1-0.6b	fw/de	25.00	0.040	218.5	8.8	0.0	–
pplx-embed-v1-0.6b	fw/fr	22.76	0.030	200.1	8.9	0.0	–
pplx-embed-v1-0.6b	fw/es	22.09	0.020	191.1	8.7	0.0	–
pplx-embed-v1-0.6b	fw/ar	17.87	0.063	150.4	8.1	0.0	–
pplx-embed-v1-0.6b	fw/zh	20.15	0.048	167.1	8.0	0.0	–
pplx-embed-v1-0.6b	fw/ja	22.06	0.028	183.1	8.1	0.0	–
pplx-embed-v1-4b	fw/en	33.74	0.071	532.2	15.4	0.1	–
pplx-embed-v1-4b	fw/de	32.67	0.058	493.7	15.4	0.1	–
pplx-embed-v1-4b	fw/fr	32.26	0.058	474.3	14.6	0.1	–
pplx-embed-v1-4b	fw/es	29.52	0.045	457.8	15.4	0.1	–
pplx-embed-v1-4b	fw/ar	27.56	0.081	417.1	15.0	0.1	–
pplx-embed-v1-4b	fw/zh	27.93	0.070	430.4	15.6	0.2	–
pplx-embed-v1-4b	fw/ja	31.48	0.051	459.1	14.7	0.1	–

Table 8: Complete positional-speed and parallel-endpoint inventory. $\sigma_0, \ldots, \sigma_4$ are median speeds in five equal bins of normalized prefix position, divided by the trajectory-wide median. Cosine and L2 compare each FLORES+ sentence endpoint with its English parallel endpoint; English and FineWeb rows have no parallel-reference entry. L2 remains on the model's native scale.

Model	Shard	σ_0	σ_1	σ_2	σ_3	σ_4	Cosine	L2
Qwen3-0.6B-Base	flores/en	1.020	0.966	0.972	0.978	1.106	–	–
Qwen3-0.6B-Base	flores/de	0.980	0.992	0.998	1.014	1.027	0.735	96.3
Qwen3-0.6B-Base	flores/fr	0.993	0.987	1.003	1.000	1.013	0.734	93.8
Qwen3-0.6B-Base	flores/es	1.011	0.989	0.973	0.984	1.028	0.747	93.9
Qwen3-0.6B-Base	flores/ar	1.034	0.998	0.995	1.009	0.951	0.175	170.0
Qwen3-0.6B-Base	flores/zh	1.060	0.990	0.985	0.979	1.001	0.354	139.8
Qwen3-0.6B-Base	flores/ja	1.038	0.984	0.985	0.985	1.009	0.189	187.0
Qwen3-4B-Base	flores/en	1.032	0.960	0.971	0.968	1.123	–	–
Qwen3-4B-Base	flores/de	0.987	0.975	0.991	1.006	1.030	0.855	100.2
Qwen3-4B-Base	flores/fr	0.998	0.967	0.988	0.989	1.045	0.877	92.9
Qwen3-4B-Base	flores/es	1.025	0.968	0.965	0.978	1.043	0.875	93.6
Qwen3-4B-Base	flores/ar	1.020	0.993	0.997	1.019	0.964	0.285	246.4

Table 8 continued

Model	Shard	σ_0	σ_1	σ_2	σ_3	σ_4	Cosine	L2
Qwen3-4B-Base	flores/zh	1.092	0.984	0.983	0.973	1.006	0.582	167.4
Qwen3-4B-Base	flores/ja	1.072	0.977	0.979	0.986	1.003	0.352	222.9
Qwen3-8B-Base	flores/en	1.010	0.947	0.967	0.966	1.197	–	–
Qwen3-8B-Base	flores/de	0.956	0.955	0.982	0.978	1.105	0.965	91.6
Qwen3-8B-Base	flores/fr	0.972	0.962	0.979	0.990	1.083	0.970	80.6
Qwen3-8B-Base	flores/es	0.979	0.957	0.962	0.976	1.097	0.971	77.8
Qwen3-8B-Base	flores/ar	0.998	0.985	0.998	1.019	1.001	0.283	304.0
Qwen3-8B-Base	flores/zh	1.096	0.997	0.994	0.957	0.973	0.913	128.1
Qwen3-8B-Base	flores/ja	1.084	0.967	0.966	0.958	1.039	0.764	203.7
Qwen3-14B-Base	flores/en	0.928	0.977	0.992	0.985	1.148	–	–
Qwen3-14B-Base	flores/de	0.956	0.968	0.994	0.989	1.117	0.814	134.0
Qwen3-14B-Base	flores/fr	0.970	0.974	0.990	1.002	1.067	0.820	131.0
Qwen3-14B-Base	flores/es	0.971	0.974	0.981	0.993	1.080	0.837	126.2
Qwen3-14B-Base	flores/ar	0.997	0.992	0.999	1.023	0.987	0.207	238.0
Qwen3-14B-Base	flores/zh	1.011	0.991	0.983	0.964	1.053	0.546	218.0
Qwen3-14B-Base	flores/ja	1.052	0.963	0.968	0.966	1.075	0.259	238.4
Qwen3-Embedding-0.6B	flores/en	1.033	0.997	0.991	0.997	0.989	–	–
Qwen3-Embedding-0.6B	flores/de	1.040	1.001	1.001	1.008	0.940	0.820	65.4
Qwen3-Embedding-0.6B	flores/fr	1.052	0.983	0.997	0.996	0.979	0.790	71.4
Qwen3-Embedding-0.6B	flores/es	1.050	0.994	0.991	0.992	0.969	0.815	66.2
Qwen3-Embedding-0.6B	flores/ar	1.045	1.018	0.991	1.005	0.923	0.734	77.8
Qwen3-Embedding-0.6B	flores/zh	1.080	1.000	0.990	1.004	0.943	0.825	60.9
Qwen3-Embedding-0.6B	flores/ja	1.085	1.032	1.019	1.025	0.803	0.748	78.2
Qwen3-Embedding-4B	flores/en	0.988	0.973	0.993	1.012	1.026	–	–
Qwen3-Embedding-4B	flores/de	1.030	0.983	0.991	1.005	0.993	0.860	71.3
Qwen3-Embedding-4B	flores/fr	0.999	0.969	0.997	1.011	1.017	0.842	76.5
Qwen3-Embedding-4B	flores/es	0.994	0.969	0.989	1.007	1.030	0.844	75.4
Qwen3-Embedding-4B	flores/ar	1.032	0.989	0.990	1.007	0.983	0.803	84.5
Qwen3-Embedding-4B	flores/zh	1.022	0.978	0.996	1.001	1.010	0.839	76.0
Qwen3-Embedding-4B	flores/ja	1.032	0.996	1.009	1.017	0.961	0.796	86.0
Qwen3-Embedding-8B	flores/en	0.950	0.968	1.012	1.031	1.019	–	–
Qwen3-Embedding-8B	flores/de	1.024	0.996	1.006	1.011	0.969	0.860	69.6
Qwen3-Embedding-8B	flores/fr	0.981	0.980	1.009	1.024	1.004	0.851	71.9
Qwen3-Embedding-8B	flores/es	0.980	0.976	1.008	1.026	1.002	0.844	73.2
Qwen3-Embedding-8B	flores/ar	1.026	0.998	0.997	1.006	0.976	0.808	81.2
Qwen3-Embedding-8B	flores/zh	1.003	0.987	1.008	1.013	0.990	0.835	75.2
Qwen3-Embedding-8B	flores/ja	1.025	0.997	1.009	1.016	0.956	0.808	81.0
bge-m3	flores/en	2.905	1.342	0.841	0.609	0.597	–	–
bge-m3	flores/de	3.111	1.376	0.871	0.630	0.540	0.952	8.2
bge-m3	flores/fr	3.157	1.371	0.854	0.616	0.559	0.955	7.9
bge-m3	flores/es	3.033	1.394	0.835	0.589	0.564	0.945	8.7
bge-m3	flores/ar	2.981	1.397	0.876	0.619	0.512	0.916	10.8
bge-m3	flores/zh	3.177	1.445	0.876	0.632	0.552	0.919	10.6
bge-m3	flores/ja	3.164	1.382	0.877	0.646	0.585	0.920	10.6
pplx-embed-v1-0.6b	flores/en	4.102	1.447	0.905	0.639	0.576	–	–
pplx-embed-v1-0.6b	flores/de	4.292	1.452	0.905	0.641	0.479	0.822	3.1
pplx-embed-v1-0.6b	flores/fr	4.571	1.454	0.892	0.607	0.490	0.831	3.0
pplx-embed-v1-0.6b	flores/es	4.540	1.488	0.893	0.608	0.498	0.798	3.2
pplx-embed-v1-0.6b	flores/ar	4.258	1.500	0.864	0.603	0.489	0.692	3.9
pplx-embed-v1-0.6b	flores/zh	4.116	1.482	0.909	0.664	0.524	0.775	3.4
pplx-embed-v1-0.6b	flores/ja	4.348	1.465	0.908	0.674	0.507	0.761	3.5
pplx-embed-v1-4b	flores/en	2.341	1.283	0.922	0.724	0.621	–	–
pplx-embed-v1-4b	flores/de	2.366	1.318	0.937	0.691	0.494	0.877	7.7
pplx-embed-v1-4b	flores/fr	2.456	1.321	0.942	0.674	0.519	0.874	7.7
pplx-embed-v1-4b	flores/es	2.613	1.336	0.928	0.669	0.525	0.859	7.9
pplx-embed-v1-4b	flores/ar	2.615	1.337	0.931	0.663	0.512	0.813	9.1
pplx-embed-v1-4b	flores/zh	2.390	1.297	0.928	0.713	0.553	0.842	8.4

Table 8 continued

Model	Shard	σ_0	σ_1	σ_2	σ_3	σ_4	Cosine	L2
pplx-embed-v1-4b	flores/ja	2.717	1.335	0.918	0.706	0.571	0.826	8.8
text-embedding-3-small	flores/en	2.097	1.248	0.936	0.771	0.655	–	–
text-embedding-3-small	flores/de	2.202	1.274	0.931	0.767	0.635	0.754	0.7
text-embedding-3-small	flores/fr	2.074	1.264	0.936	0.757	0.666	0.787	0.7
text-embedding-3-small	flores/es	2.069	1.275	0.939	0.747	0.666	0.755	0.7
text-embedding-3-small	flores/ar	2.654	1.334	0.897	0.687	0.582	0.489	1.0
text-embedding-3-small	flores/zh	2.287	1.284	0.928	0.747	0.642	0.668	0.8
text-embedding-3-small	flores/ja	2.427	1.284	0.941	0.767	0.608	0.622	0.9
text-embedding-3-large	flores/en	1.758	1.140	0.964	0.860	0.725	–	–
text-embedding-3-large	flores/de	1.609	1.117	0.968	0.881	0.757	0.759	0.7
text-embedding-3-large	flores/fr	1.576	1.125	0.969	0.871	0.751	0.786	0.7
text-embedding-3-large	flores/es	1.543	1.118	0.971	0.877	0.761	0.747	0.7
text-embedding-3-large	flores/ar	1.742	1.240	0.972	0.763	0.562	0.659	0.8
text-embedding-3-large	flores/zh	1.839	1.217	0.967	0.803	0.635	0.742	0.7
text-embedding-3-large	flores/ja	1.908	1.229	0.980	0.806	0.558	0.726	0.7
voyage-4-lite	flores/en	2.999	1.366	0.871	0.651	0.551	–	–
voyage-4-lite	flores/de	2.995	1.363	0.875	0.665	0.514	0.896	0.5
voyage-4-lite	flores/fr	3.065	1.319	0.862	0.649	0.531	0.898	0.5
voyage-4-lite	flores/es	2.970	1.353	0.873	0.647	0.549	0.885	0.5
voyage-4-lite	flores/ar	2.764	1.372	0.868	0.643	0.522	0.831	0.6
voyage-4-lite	flores/zh	2.682	1.387	0.881	0.669	0.550	0.873	0.5
voyage-4-lite	flores/ja	2.671	1.382	0.919	0.681	0.515	0.859	0.5
voyage-4	flores/en	3.345	1.427	0.859	0.636	0.515	–	–
voyage-4	flores/de	3.526	1.397	0.858	0.641	0.474	0.916	0.4
voyage-4	flores/fr	3.605	1.353	0.838	0.626	0.499	0.910	0.4
voyage-4	flores/es	3.379	1.377	0.855	0.626	0.503	0.898	0.5
voyage-4	flores/ar	3.302	1.446	0.833	0.609	0.469	0.861	0.5
voyage-4	flores/zh	2.945	1.423	0.886	0.661	0.524	0.883	0.5
voyage-4	flores/ja	2.930	1.416	0.911	0.666	0.487	0.869	0.5
voyage-4-large	flores/en	3.092	1.378	0.883	0.666	0.632	–	–
voyage-4-large	flores/de	3.354	1.356	0.857	0.680	0.635	0.867	0.5
voyage-4-large	flores/fr	3.274	1.317	0.858	0.675	0.652	0.864	0.5
voyage-4-large	flores/es	3.242	1.333	0.866	0.673	0.639	0.846	0.6
voyage-4-large	flores/ar	3.388	1.382	0.863	0.664	0.631	0.799	0.6
voyage-4-large	flores/zh	2.921	1.416	0.886	0.684	0.636	0.833	0.6
voyage-4-large	flores/ja	2.994	1.391	0.900	0.672	0.643	0.811	0.6
Qwen3-0.6B-Base	fw/en	1.004	1.000	1.000	0.998	0.998	–	–
Qwen3-0.6B-Base	fw/de	0.999	1.013	1.005	1.001	0.982	–	–
Qwen3-0.6B-Base	fw/fr	0.994	1.006	1.007	1.005	0.989	–	–
Qwen3-0.6B-Base	fw/es	1.003	1.007	1.005	0.999	0.987	–	–
Qwen3-0.6B-Base	fw/ar	0.998	1.000	1.002	1.002	0.997	–	–
Qwen3-0.6B-Base	fw/zh	1.021	0.997	0.990	0.991	1.002	–	–
Qwen3-0.6B-Base	fw/ja	1.007	1.000	0.996	0.995	1.001	–	–
Qwen3-4B-Base	fw/en	1.004	0.999	0.999	0.997	1.001	–	–
Qwen3-4B-Base	fw/de	1.002	1.006	1.001	0.998	0.993	–	–
Qwen3-4B-Base	fw/fr	1.000	1.003	1.002	1.000	0.995	–	–
Qwen3-4B-Base	fw/es	1.004	1.002	1.002	0.997	0.995	–	–
Qwen3-4B-Base	fw/ar	1.004	0.998	0.999	1.000	0.999	–	–
Qwen3-4B-Base	fw/zh	1.022	0.994	0.988	0.990	1.006	–	–
Qwen3-4B-Base	fw/ja	1.007	0.997	0.994	0.994	1.008	–	–
Qwen3-8B-Base	fw/en	1.006	0.997	0.997	0.996	1.004	–	–
Qwen3-8B-Base	fw/de	1.001	1.002	0.997	0.997	1.003	–	–
Qwen3-8B-Base	fw/fr	1.005	1.001	0.999	0.997	0.998	–	–
Qwen3-8B-Base	fw/es	1.005	1.000	1.000	0.997	0.998	–	–
Qwen3-8B-Base	fw/ar	1.008	0.996	0.997	0.997	1.001	–	–
Qwen3-8B-Base	fw/zh	1.030	0.991	0.984	0.986	1.010	–	–
Qwen3-8B-Base	fw/ja	1.015	0.991	0.988	0.991	1.017	–	–

Table 8 continued

Model	Shard	σ_0	σ_1	σ_2	σ_3	σ_4	Cosine	L2
Qwen3-14B-Base	fw/en	0.976	1.002	1.008	1.009	1.005	–	–
Qwen3-14B-Base	fw/de	0.993	0.998	0.998	1.000	1.013	–	–
Qwen3-14B-Base	fw/fr	0.986	0.999	1.004	1.006	1.006	–	–
Qwen3-14B-Base	fw/es	0.986	1.001	1.004	1.004	1.006	–	–
Qwen3-14B-Base	fw/ar	1.002	0.995	0.997	0.999	1.006	–	–
Qwen3-14B-Base	fw/zh	1.007	0.990	0.990	0.995	1.019	–	–
Qwen3-14B-Base	fw/ja	1.006	0.988	0.987	0.995	1.026	–	–
Qwen3-Embedding-0.6B	fw/en	0.988	0.999	1.004	1.004	1.004	–	–
Qwen3-Embedding-0.6B	fw/de	0.993	1.007	1.006	1.002	0.992	–	–
Qwen3-Embedding-0.6B	fw/fr	0.995	1.001	1.003	1.002	0.998	–	–
Qwen3-Embedding-0.6B	fw/es	0.994	1.000	1.002	1.002	1.002	–	–
Qwen3-Embedding-0.6B	fw/ar	0.992	0.993	1.001	1.008	1.006	–	–
Qwen3-Embedding-0.6B	fw/zh	1.011	1.001	0.996	0.996	0.996	–	–
Qwen3-Embedding-0.6B	fw/ja	0.991	0.998	1.003	1.004	1.003	–	–
Qwen3-Embedding-4B	fw/en	1.008	1.003	0.999	0.994	0.996	–	–
Qwen3-Embedding-4B	fw/de	0.980	1.001	1.007	1.007	1.006	–	–
Qwen3-Embedding-4B	fw/fr	0.986	1.003	1.006	1.005	1.002	–	–
Qwen3-Embedding-4B	fw/es	0.999	1.000	1.000	1.001	1.000	–	–
Qwen3-Embedding-4B	fw/ar	0.989	0.999	1.002	1.004	1.007	–	–
Qwen3-Embedding-4B	fw/zh	1.016	1.005	0.994	0.989	0.995	–	–
Qwen3-Embedding-4B	fw/ja	0.991	1.001	1.002	1.001	1.005	–	–
Qwen3-Embedding-8B	fw/en	1.041	1.010	0.991	0.981	0.977	–	–
Qwen3-Embedding-8B	fw/de	0.990	1.004	1.006	1.005	0.997	–	–
Qwen3-Embedding-8B	fw/fr	0.990	1.006	1.006	1.003	0.996	–	–
Qwen3-Embedding-8B	fw/es	0.994	1.006	1.004	1.000	0.997	–	–
Qwen3-Embedding-8B	fw/ar	1.001	1.001	1.000	0.999	0.999	–	–
Qwen3-Embedding-8B	fw/zh	1.032	1.010	0.990	0.979	0.982	–	–
Qwen3-Embedding-8B	fw/ja	1.009	1.004	0.996	0.992	0.997	–	–
bge-m3	fw/en	2.424	1.135	0.868	0.728	0.666	–	–
bge-m3	fw/de	2.762	1.130	0.850	0.735	0.671	–	–
bge-m3	fw/fr	1.031	0.994	0.987	1.000	0.989	–	–
bge-m3	fw/es	2.656	1.081	0.847	0.742	0.701	–	–
bge-m3	fw/ar	3.101	1.165	0.858	0.716	0.663	–	–
bge-m3	fw/zh	2.652	1.124	0.860	0.739	0.692	–	–
bge-m3	fw/ja	2.556	1.132	0.861	0.741	0.683	–	–
pplx-embed-v1-0.6b	fw/en	2.123	1.236	0.932	0.750	0.658	–	–
pplx-embed-v1-0.6b	fw/de	5.332	1.435	0.834	0.576	0.452	–	–
pplx-embed-v1-0.6b	fw/fr	5.025	1.391	0.822	0.588	0.478	–	–
pplx-embed-v1-0.6b	fw/es	4.859	1.401	0.820	0.583	0.474	–	–
pplx-embed-v1-0.6b	fw/ar	5.402	1.350	0.823	0.602	0.489	–	–
pplx-embed-v1-0.6b	fw/zh	4.898	1.421	0.817	0.578	0.467	–	–
pplx-embed-v1-0.6b	fw/ja	5.093	1.452	0.811	0.562	0.440	–	–
pplx-embed-v1-4b	fw/en	4.907	1.356	0.837	0.590	0.486	–	–
pplx-embed-v1-4b	fw/de	5.040	1.361	0.821	0.596	0.504	–	–
pplx-embed-v1-4b	fw/fr	4.806	1.361	0.835	0.606	0.508	–	–
pplx-embed-v1-4b	fw/es	4.606	1.365	0.824	0.596	0.493	–	–
pplx-embed-v1-4b	fw/ar	4.990	1.379	0.826	0.591	0.477	–	–
pplx-embed-v1-4b	fw/zh	4.545	1.365	0.807	0.593	0.512	–	–
pplx-embed-v1-4b	fw/ja	4.874	1.399	0.812	0.588	0.481	–	–

B.2 PREDICTIVE SUFFICIENCY OF REDUCED COORDINATES

Table 9: Complete predictive-gate inventory over 280 stored shards. Rank and Lag identify the selected PCA/delay coordinate; Markov records whether that selected state passes the predictive gate. $NMSE_1$ and $NMSE_8$ are held-out autonomous errors at horizons one and eight, and Pers. is the horizon-one persistence error. The omitted rank–lag–horizon candidates are retained in the machine-readable outputs.

Model	Shard	Rank	Lag	Markov	$NMSE_1$	$NMSE_8$	Pers.
Qwen3-0.6B-Base	dev/eng	32	1	yes	0.650	0.930	1.264
Qwen3-0.6B-Base	dev/deu	8	1	no	0.815	1.001	1.595
Qwen3-0.6B-Base	dev/fra	8	1	yes	0.812	1.000	1.578
Qwen3-0.6B-Base	dev/spa	16	1	yes	0.786	0.996	1.610
Qwen3-0.6B-Base	dev/arb	8	1	yes	0.859	1.000	2.017
Qwen3-0.6B-Base	dev/cmn	16	2	yes	0.505	0.872	0.824
Qwen3-0.6B-Base	dev/jpn	16	1	yes	0.708	0.983	1.391
Qwen3-0.6B-Base	devtest/eng	32	1	yes	0.652	0.930	1.274
Qwen3-0.6B-Base	devtest/deu	8	1	no	0.827	1.002	1.650
Qwen3-0.6B-Base	devtest/fra	8	1	no	0.814	1.001	1.584
Qwen3-0.6B-Base	devtest/spa	16	1	yes	0.774	0.998	1.633
Qwen3-0.6B-Base	devtest/arb	8	1	no	0.859	1.000	1.999
Qwen3-0.6B-Base	devtest/cmn	16	2	yes	0.509	0.865	0.852
Qwen3-0.6B-Base	devtest/jpn	16	1	yes	0.720	0.982	1.412
Qwen3-0.6B-Base	fw/en	32	1	yes	0.691	0.921	1.346
Qwen3-0.6B-Base	fw/de	8	1	yes	0.813	0.994	1.540
Qwen3-0.6B-Base	fw/fr	16	1	yes	0.783	0.985	1.494
Qwen3-0.6B-Base	fw/es	16	1	yes	0.778	0.987	1.551
Qwen3-0.6B-Base	fw/ar	8	1	yes	0.831	0.990	1.889
Qwen3-0.6B-Base	fw/zh	8	4	yes	0.264	0.378	0.403
Qwen3-0.6B-Base	fw/ja	16	2	yes	0.602	0.859	1.092
Qwen3-4B-Base	dev/eng	32	1	yes	0.649	0.946	1.263
Qwen3-4B-Base	dev/deu	16	1	yes	0.767	1.000	1.560
Qwen3-4B-Base	dev/fra	16	1	yes	0.805	1.000	1.589
Qwen3-4B-Base	dev/spa	16	1	yes	0.795	0.997	1.636
Qwen3-4B-Base	dev/arb	8	1	no	0.876	1.000	2.004
Qwen3-4B-Base	dev/cmn	32	2	yes	0.543	0.883	0.977
Qwen3-4B-Base	dev/jpn	16	1	yes	0.717	0.987	1.354
Qwen3-4B-Base	devtest/eng	32	1	yes	0.644	0.945	1.275
Qwen3-4B-Base	devtest/deu	16	1	no	0.773	1.001	1.616
Qwen3-4B-Base	devtest/fra	16	1	no	0.793	1.001	1.605
Qwen3-4B-Base	devtest/spa	16	1	no	0.779	1.001	1.655
Qwen3-4B-Base	devtest/arb	8	1	yes	0.858	1.000	1.986
Qwen3-4B-Base	devtest/cmn	32	2	yes	0.543	0.876	1.018
Qwen3-4B-Base	devtest/jpn	32	1	yes	0.684	0.979	1.370
Qwen3-4B-Base	fw/en	64	1	yes	0.654	0.929	1.336
Qwen3-4B-Base	fw/de	32	1	yes	0.755	0.985	1.509
Qwen3-4B-Base	fw/fr	16	1	yes	0.786	0.990	1.564
Qwen3-4B-Base	fw/es	32	1	yes	0.736	0.982	1.569
Qwen3-4B-Base	fw/ar	16	1	yes	0.807	0.995	1.853
Qwen3-4B-Base	fw/zh	8	4	yes	0.280	0.368	0.437
Qwen3-4B-Base	fw/ja	16	2	yes	0.638	0.877	1.123
Qwen3-8B-Base	dev/eng	32	1	yes	0.717	0.960	1.297
Qwen3-8B-Base	dev/deu	16	1	no	0.788	1.001	1.559
Qwen3-8B-Base	dev/fra	32	1	yes	0.776	0.998	1.669
Qwen3-8B-Base	dev/spa	16	1	yes	0.819	1.000	1.772
Qwen3-8B-Base	dev/arb	16	1	no	0.819	1.000	1.831
Qwen3-8B-Base	dev/cmn	32	1	yes	0.682	0.944	1.172
Qwen3-8B-Base	dev/jpn	32	1	yes	0.633	0.989	1.136
Qwen3-8B-Base	devtest/eng	32	1	yes	0.720	0.961	1.336

Table 9 continued

Model	Shard	Rank	Lag	Markov	$NMSE_1$	$NMSE_8$	Pers.
Qwen3-8B-Base	devtest/deu	32	1	yes	0.756	1.000	1.618
Qwen3-8B-Base	devtest/fra	32	1	no	0.767	1.000	1.690
Qwen3-8B-Base	devtest/spa	16	1	no	0.805	1.002	1.793
Qwen3-8B-Base	devtest/arb	16	1	yes	0.813	1.000	1.832
Qwen3-8B-Base	devtest/cmn	32	1	yes	0.703	0.949	1.254
Qwen3-8B-Base	devtest/jpn	32	1	yes	0.648	0.991	1.174
Qwen3-8B-Base	fw/en	64	1	yes	0.701	0.938	1.425
Qwen3-8B-Base	fw/de	32	1	yes	0.754	0.987	1.526
Qwen3-8B-Base	fw/fr	32	1	yes	0.761	0.987	1.654
Qwen3-8B-Base	fw/es	32	1	yes	0.744	0.989	1.757
Qwen3-8B-Base	fw/ar	16	1	yes	0.766	0.990	1.576
Qwen3-8B-Base	fw/zh	32	2	yes	0.379	0.517	0.619
Qwen3-8B-Base	fw/ja	32	2	yes	0.557	0.844	0.932
Qwen3-14B-Base	dev/eng	32	1	yes	0.733	0.969	1.500
Qwen3-14B-Base	dev/deu	32	1	yes	0.759	0.998	1.524
Qwen3-14B-Base	dev/fra	32	1	yes	0.778	0.997	1.566
Qwen3-14B-Base	dev/spa	16	1	yes	0.815	1.000	1.624
Qwen3-14B-Base	dev/arb	8	1	no	0.855	1.001	1.887
Qwen3-14B-Base	dev/cmn	32	1	yes	0.631	0.958	1.068
Qwen3-14B-Base	dev/jpn	32	1	yes	0.700	0.987	1.321
Qwen3-14B-Base	devtest/eng	32	1	yes	0.717	0.976	1.526
Qwen3-14B-Base	devtest/deu	32	1	yes	0.756	0.998	1.596
Qwen3-14B-Base	devtest/fra	32	1	yes	0.782	0.997	1.579
Qwen3-14B-Base	devtest/spa	32	1	yes	0.780	0.996	1.648
Qwen3-14B-Base	devtest/arb	16	1	yes	0.799	0.999	1.852
Qwen3-14B-Base	devtest/cmn	32	1	yes	0.661	0.951	1.136
Qwen3-14B-Base	devtest/jpn	32	1	yes	0.713	0.988	1.370
Qwen3-14B-Base	fw/en	32	1	yes	0.822	0.985	1.760
Qwen3-14B-Base	fw/de	32	1	yes	0.779	0.991	1.542
Qwen3-14B-Base	fw/fr	32	1	yes	0.812	0.992	1.606
Qwen3-14B-Base	fw/es	32	1	yes	0.797	0.991	1.639
Qwen3-14B-Base	fw/ar	16	1	yes	0.778	0.992	1.735
Qwen3-14B-Base	fw/zh	16	2	yes	0.478	0.621	0.770
Qwen3-14B-Base	fw/ja	32	2	yes	0.649	0.886	1.125
Qwen3-Embedding-0.6B	dev/eng	32	4	yes	0.501	0.789	0.944
Qwen3-Embedding-0.6B	dev/deu	16	2	yes	0.698	0.951	1.263
Qwen3-Embedding-0.6B	dev/fra	16	2	yes	0.666	0.949	1.230
Qwen3-Embedding-0.6B	dev/spa	16	2	yes	0.654	0.940	1.213
Qwen3-Embedding-0.6B	dev/arb	16	1	yes	0.778	0.987	1.439
Qwen3-Embedding-0.6B	dev/cmn	32	4	yes	0.513	0.814	0.918
Qwen3-Embedding-0.6B	dev/jpn	16	2	yes	0.633	0.941	1.144
Qwen3-Embedding-0.6B	devtest/eng	16	4	yes	0.522	0.808	1.063
Qwen3-Embedding-0.6B	devtest/deu	16	2	yes	0.713	0.956	1.294
Qwen3-Embedding-0.6B	devtest/fra	16	2	yes	0.672	0.954	1.243
Qwen3-Embedding-0.6B	devtest/spa	16	2	yes	0.652	0.936	1.221
Qwen3-Embedding-0.6B	devtest/arb	16	1	yes	0.777	0.988	1.447
Qwen3-Embedding-0.6B	devtest/cmn	32	4	yes	0.516	0.806	0.938
Qwen3-Embedding-0.6B	devtest/jpn	16	2	yes	0.644	0.948	1.159
Qwen3-Embedding-0.6B	fw/en	32	2	yes	0.573	0.862	0.961
Qwen3-Embedding-0.6B	fw/de	16	2	yes	0.731	0.958	1.229
Qwen3-Embedding-0.6B	fw/fr	16	1	yes	0.710	0.976	1.212
Qwen3-Embedding-0.6B	fw/es	16	2	yes	0.681	0.939	1.191
Qwen3-Embedding-0.6B	fw/ar	16	2	yes	0.733	0.937	1.311
Qwen3-Embedding-0.6B	fw/zh	8	4	yes	0.473	0.645	0.756
Qwen3-Embedding-0.6B	fw/ja	8	2	yes	0.683	0.939	1.166
Qwen3-Embedding-4B	dev/eng	32	2	yes	0.552	0.843	1.059
Qwen3-Embedding-4B	dev/deu	32	2	yes	0.640	0.901	1.141

Table 9 continued

Model	Shard	Rank	Lag	Markov	$NMSE_1$	$NMSE_8$	Pers.
Qwen3-Embedding-4B	dev/fra	32	2	yes	0.675	0.919	1.256
Qwen3-Embedding-4B	dev/spa	32	2	yes	0.639	0.904	1.195
Qwen3-Embedding-4B	dev/arb	32	2	yes	0.731	0.938	1.384
Qwen3-Embedding-4B	dev/cmn	32	2	yes	0.538	0.854	1.032
Qwen3-Embedding-4B	dev/jpn	32	2	yes	0.625	0.910	1.103
Qwen3-Embedding-4B	devtest/eng	32	2	yes	0.566	0.856	1.098
Qwen3-Embedding-4B	devtest/deu	32	2	yes	0.658	0.910	1.181
Qwen3-Embedding-4B	devtest/fra	32	2	yes	0.678	0.917	1.266
Qwen3-Embedding-4B	devtest/spa	32	2	yes	0.644	0.906	1.212
Qwen3-Embedding-4B	devtest/arb	32	2	yes	0.733	0.941	1.392
Qwen3-Embedding-4B	devtest/cmn	32	2	yes	0.539	0.853	1.045
Qwen3-Embedding-4B	devtest/jpn	32	2	yes	0.633	0.914	1.121
Qwen3-Embedding-4B	fw/en	32	2	yes	0.642	0.884	1.155
Qwen3-Embedding-4B	fw/de	16	2	yes	0.696	0.930	1.155
Qwen3-Embedding-4B	fw/fr	16	2	yes	0.729	0.941	1.310
Qwen3-Embedding-4B	fw/es	32	2	yes	0.661	0.901	1.187
Qwen3-Embedding-4B	fw/ar	16	2	yes	0.690	0.908	1.214
Qwen3-Embedding-4B	fw/zh	8	2	yes	0.586	0.796	0.947
Qwen3-Embedding-4B	fw/ja	16	2	yes	0.639	0.898	1.056
Qwen3-Embedding-8B	dev/eng	32	2	yes	0.623	0.897	1.209
Qwen3-Embedding-8B	dev/deu	32	1	yes	0.737	0.973	1.431
Qwen3-Embedding-8B	dev/fra	32	1	yes	0.778	0.982	1.562
Qwen3-Embedding-8B	dev/spa	32	1	yes	0.750	0.971	1.521
Qwen3-Embedding-8B	dev/arb	32	2	yes	0.731	0.944	1.551
Qwen3-Embedding-8B	dev/cmn	32	2	yes	0.584	0.856	1.245
Qwen3-Embedding-8B	dev/jpn	32	2	yes	0.640	0.913	1.323
Qwen3-Embedding-8B	devtest/eng	64	2	yes	0.609	0.886	1.234
Qwen3-Embedding-8B	devtest/deu	32	1	yes	0.752	0.978	1.488
Qwen3-Embedding-8B	devtest/fra	32	1	yes	0.786	0.984	1.577
Qwen3-Embedding-8B	devtest/spa	32	1	yes	0.749	0.976	1.556
Qwen3-Embedding-8B	devtest/arb	64	2	yes	0.714	0.931	1.490
Qwen3-Embedding-8B	devtest/cmn	32	2	yes	0.581	0.855	1.243
Qwen3-Embedding-8B	devtest/jpn	32	2	yes	0.645	0.916	1.333
Qwen3-Embedding-8B	fw/en	32	2	yes	0.692	0.914	1.302
Qwen3-Embedding-8B	fw/de	16	1	yes	0.775	0.985	1.392
Qwen3-Embedding-8B	fw/fr	32	1	yes	0.771	0.983	1.532
Qwen3-Embedding-8B	fw/es	32	2	yes	0.747	0.954	1.478
Qwen3-Embedding-8B	fw/ar	16	2	yes	0.742	0.935	1.383
Qwen3-Embedding-8B	fw/zh	32	2	yes	0.594	0.809	1.031
Qwen3-Embedding-8B	fw/ja	16	2	yes	0.654	0.909	1.124
bge-m3	dev/eng	16	4	yes	0.125	0.383	0.380
bge-m3	dev/deu	8	4	yes	0.105	0.334	0.391
bge-m3	dev/fra	16	4	yes	0.095	0.319	0.305
bge-m3	dev/spa	16	4	yes	0.136	0.389	0.379
bge-m3	dev/arb	8	4	yes	0.093	0.338	0.377
bge-m3	dev/cmn	16	4	yes	0.105	0.364	0.367
bge-m3	dev/jpn	16	4	yes	0.074	0.295	0.310
bge-m3	devtest/eng	16	4	yes	0.109	0.346	0.363
bge-m3	devtest/deu	8	4	yes	0.099	0.300	0.370
bge-m3	devtest/fra	16	4	yes	0.088	0.295	0.290
bge-m3	devtest/spa	32	4	yes	0.115	0.354	0.296
bge-m3	devtest/arb	8	4	yes	0.083	0.293	0.345
bge-m3	devtest/cmn	8	4	yes	0.101	0.319	0.406
bge-m3	devtest/jpn	8	4	yes	0.074	0.280	0.336
bge-m3	fw/en	8	4	yes	0.001	0.004	0.001
bge-m3	fw/de	8	4	yes	0.001	0.005	0.002
bge-m3	fw/fr	16	1	yes	0.005	0.036	0.005

Table 9 continued

Model	Shard	Rank	Lag	Markov	$NMSE_1$	$NMSE_8$	Pers.
bge-m3	fw/es	8	4	yes	0.001	0.005	0.002
bge-m3	fw/ar	8	4	yes	0.001	0.006	0.002
bge-m3	fw/zh	8	4	yes	0.001	0.005	0.002
bge-m3	fw/ja	8	4	yes	0.001	0.004	0.001
pplx-embed-v1-0.6b	dev/eng	8	4	yes	0.132	0.365	0.339
pplx-embed-v1-0.6b	dev/deu	8	4	yes	0.096	0.294	0.287
pplx-embed-v1-0.6b	dev/fra	8	4	yes	0.107	0.333	0.263
pplx-embed-v1-0.6b	dev/spa	8	4	yes	0.098	0.331	0.288
pplx-embed-v1-0.6b	dev/arb	8	4	yes	0.077	0.301	0.212
pplx-embed-v1-0.6b	dev/cmn	8	4	yes	0.098	0.317	0.274
pplx-embed-v1-0.6b	dev/jpn	8	4	yes	0.076	0.277	0.214
pplx-embed-v1-0.6b	devtest/eng	8	4	yes	0.114	0.350	0.350
pplx-embed-v1-0.6b	devtest/deu	8	4	yes	0.098	0.292	0.309
pplx-embed-v1-0.6b	devtest/fra	8	4	yes	0.111	0.327	0.276
pplx-embed-v1-0.6b	devtest/spa	8	4	yes	0.095	0.300	0.303
pplx-embed-v1-0.6b	devtest/arb	8	4	yes	0.077	0.284	0.193
pplx-embed-v1-0.6b	devtest/cmn	8	4	yes	0.108	0.303	0.280
pplx-embed-v1-0.6b	devtest/jpn	8	4	yes	0.079	0.265	0.211
pplx-embed-v1-0.6b	fw/en	16	4	yes	0.009	0.031	0.011
pplx-embed-v1-0.6b	fw/de	16	4	yes	0.007	0.026	0.019
pplx-embed-v1-0.6b	fw/fr	16	4	yes	0.006	0.023	0.014
pplx-embed-v1-0.6b	fw/es	16	4	yes	0.005	0.020	0.016
pplx-embed-v1-0.6b	fw/ar	8	4	yes	0.007	0.029	0.018
pplx-embed-v1-0.6b	fw/zh	16	4	yes	0.007	0.023	0.019
pplx-embed-v1-0.6b	fw/ja	16	4	yes	0.006	0.022	0.013
pplx-embed-v1-4b	dev/eng	8	4	yes	0.058	0.253	0.106
pplx-embed-v1-4b	dev/deu	8	4	yes	0.046	0.204	0.073
pplx-embed-v1-4b	dev/fra	8	4	yes	0.044	0.205	0.066
pplx-embed-v1-4b	dev/spa	8	4	yes	0.045	0.214	0.076
pplx-embed-v1-4b	dev/arb	8	4	yes	0.049	0.220	0.071
pplx-embed-v1-4b	dev/cmn	8	4	yes	0.052	0.226	0.102
pplx-embed-v1-4b	dev/jpn	8	4	yes	0.048	0.207	0.072
pplx-embed-v1-4b	devtest/eng	8	4	yes	0.059	0.250	0.111
pplx-embed-v1-4b	devtest/deu	8	4	yes	0.044	0.196	0.073
pplx-embed-v1-4b	devtest/fra	8	4	yes	0.046	0.206	0.068
pplx-embed-v1-4b	devtest/spa	8	4	yes	0.044	0.208	0.075
pplx-embed-v1-4b	devtest/arb	8	4	yes	0.047	0.217	0.070
pplx-embed-v1-4b	devtest/cmn	8	4	yes	0.049	0.215	0.093
pplx-embed-v1-4b	devtest/jpn	8	4	yes	0.045	0.199	0.067
pplx-embed-v1-4b	fw/en	8	4	yes	0.002	0.011	0.003
pplx-embed-v1-4b	fw/de	8	4	yes	0.004	0.016	0.006
pplx-embed-v1-4b	fw/fr	8	4	yes	0.003	0.013	0.004
pplx-embed-v1-4b	fw/es	8	4	yes	0.002	0.008	0.003
pplx-embed-v1-4b	fw/ar	8	4	yes	0.003	0.016	0.004
pplx-embed-v1-4b	fw/zh	8	4	yes	0.002	0.009	0.003
pplx-embed-v1-4b	fw/ja	8	4	yes	0.003	0.011	0.004
text-embedding-3-small	dev/eng	8	4	yes	0.047	0.198	0.108
text-embedding-3-small	dev/deu	8	4	yes	0.044	0.197	0.091
text-embedding-3-small	dev/fra	8	4	yes	0.049	0.193	0.097
text-embedding-3-small	dev/spa	8	4	yes	0.050	0.211	0.094
text-embedding-3-small	dev/arb	8	4	yes	0.032	0.150	0.045
text-embedding-3-small	dev/cmn	8	4	yes	0.034	0.182	0.053
text-embedding-3-small	dev/jpn	8	4	yes	0.031	0.157	0.047
text-embedding-3-small	devtest/eng	8	4	yes	0.045	0.182	0.106
text-embedding-3-small	devtest/deu	8	4	yes	0.042	0.176	0.089
text-embedding-3-small	devtest/fra	8	4	yes	0.045	0.183	0.090
text-embedding-3-small	devtest/spa	8	4	yes	0.044	0.185	0.083

Table 9 continued

Model	Shard	Rank	Lag	Markov	$NMSE_1$	$NMSE_8$	Pers.
text-embedding-3-small	devtest/arb	8	4	yes	0.032	0.148	0.045
text-embedding-3-small	devtest/cmn	8	4	yes	0.032	0.161	0.051
text-embedding-3-small	devtest/jpn	8	4	yes	0.031	0.149	0.046
text-embedding-3-large	dev/eng	8	4	yes	0.076	0.214	0.145
text-embedding-3-large	dev/deu	32	4	yes	0.150	0.315	0.259
text-embedding-3-large	dev/fra	16	4	yes	0.136	0.295	0.234
text-embedding-3-large	dev/spa	16	4	yes	0.143	0.306	0.245
text-embedding-3-large	dev/arb	16	4	yes	0.124	0.261	0.180
text-embedding-3-large	dev/cmn	8	4	yes	0.065	0.192	0.095
text-embedding-3-large	dev/jpn	16	4	yes	0.066	0.187	0.096
text-embedding-3-large	devtest/eng	8	4	yes	0.078	0.219	0.150
text-embedding-3-large	devtest/deu	32	4	yes	0.147	0.298	0.261
text-embedding-3-large	devtest/fra	16	4	yes	0.140	0.299	0.238
text-embedding-3-large	devtest/spa	16	4	yes	0.152	0.313	0.260
text-embedding-3-large	devtest/arb	8	4	yes	0.131	0.266	0.188
text-embedding-3-large	devtest/cmn	16	4	yes	0.071	0.206	0.104
text-embedding-3-large	devtest/jpn	16	4	yes	0.066	0.189	0.096
voyage-4-lite	dev/eng	8	4	yes	0.052	0.242	0.113
voyage-4-lite	dev/deu	8	4	yes	0.048	0.228	0.082
voyage-4-lite	dev/fra	8	4	yes	0.043	0.198	0.084
voyage-4-lite	dev/spa	8	4	yes	0.049	0.238	0.088
voyage-4-lite	dev/arb	8	4	yes	0.044	0.216	0.068
voyage-4-lite	dev/cmn	8	4	yes	0.044	0.206	0.101
voyage-4-lite	dev/jpn	8	4	yes	0.044	0.213	0.088
voyage-4-lite	devtest/eng	8	4	yes	0.046	0.223	0.109
voyage-4-lite	devtest/deu	8	4	yes	0.045	0.216	0.081
voyage-4-lite	devtest/fra	8	4	yes	0.041	0.206	0.077
voyage-4-lite	devtest/spa	8	4	yes	0.045	0.216	0.080
voyage-4-lite	devtest/arb	8	4	yes	0.046	0.223	0.069
voyage-4-lite	devtest/cmn	8	4	yes	0.049	0.219	0.107
voyage-4-lite	devtest/jpn	8	4	yes	0.040	0.206	0.077
voyage-4	dev/eng	8	4	yes	0.047	0.232	0.098
voyage-4	dev/deu	8	4	yes	0.041	0.208	0.068
voyage-4	dev/fra	8	4	yes	0.038	0.191	0.068
voyage-4	dev/spa	8	4	yes	0.039	0.211	0.066
voyage-4	dev/arb	8	4	yes	0.039	0.196	0.055
voyage-4	dev/cmn	8	4	yes	0.038	0.190	0.080
voyage-4	dev/jpn	8	4	yes	0.038	0.185	0.070
voyage-4	devtest/eng	8	4	yes	0.040	0.208	0.088
voyage-4	devtest/deu	8	4	yes	0.040	0.196	0.070
voyage-4	devtest/fra	8	4	yes	0.036	0.192	0.064
voyage-4	devtest/spa	8	4	yes	0.038	0.200	0.064
voyage-4	devtest/arb	8	4	yes	0.041	0.195	0.058
voyage-4	devtest/cmn	8	4	yes	0.043	0.200	0.086
voyage-4	devtest/jpn	8	4	yes	0.039	0.205	0.069
voyage-4-large	dev/eng	8	4	yes	0.050	0.250	0.176
voyage-4-large	dev/deu	8	4	yes	0.035	0.195	0.092
voyage-4-large	dev/fra	8	4	yes	0.035	0.196	0.087
voyage-4-large	dev/spa	8	4	yes	0.040	0.224	0.086
voyage-4-large	dev/arb	8	4	yes	0.045	0.203	0.134
voyage-4-large	dev/cmn	8	4	yes	0.040	0.198	0.225
voyage-4-large	dev/jpn	8	4	yes	0.042	0.209	0.158
voyage-4-large	devtest/eng	8	4	yes	0.046	0.233	0.161
voyage-4-large	devtest/deu	8	4	yes	0.037	0.198	0.090
voyage-4-large	devtest/fra	8	4	yes	0.035	0.201	0.082
voyage-4-large	devtest/spa	8	4	yes	0.038	0.216	0.087
voyage-4-large	devtest/arb	8	4	yes	0.045	0.206	0.148

Table 9 continued

Model	Shard	Rank	Lag	Markov	$NMSE_1$	$NMSE_8$	Pers.
voyage-4-large	devtest/cmn	8	4	yes	0.046	0.208	0.205
voyage-4-large	devtest/jpn	8	4	yes	0.039	0.204	0.147

B.3 Effective dynamics and convergence

Table 10: Conditional affine-DMD results on all stored shards. Radius is the spectral radius; Test NMSE is held-out one-step error. Supported means that the fixed-point candidate lies within three typical nearest-neighbor distances of observed support. Attractor additionally requires local stability and at least 90% horizon-64 convergence; Conv. gives that convergence fraction. Rejected predictive coordinates are left blank rather than interpreted dynamically.

Model	Shard	Radius	Test NMSE	Stable	Supported	Attractor	Conv.
Qwen3-0.6B-Base	dev/eng	0.927	0.650	yes	yes	yes	1.000
Qwen3-0.6B-Base	dev/deu	–	–	–	–	–	–
Qwen3-0.6B-Base	dev/fra	0.647	0.812	yes	yes	yes	1.000
Qwen3-0.6B-Base	dev/spa	0.715	0.786	yes	yes	yes	1.000
Qwen3-0.6B-Base	dev/arb	0.694	0.859	yes	yes	yes	1.000
Qwen3-0.6B-Base	dev/cmn	0.944	0.505	yes	yes	yes	1.000
Qwen3-0.6B-Base	dev/jpn	0.792	0.708	yes	yes	yes	1.000
Qwen3-0.6B-Base	devtest/eng	0.923	0.652	yes	yes	yes	1.000
Qwen3-0.6B-Base	devtest/deu	–	–	–	–	–	–
Qwen3-0.6B-Base	devtest/fra	–	–	–	–	–	–
Qwen3-0.6B-Base	devtest/spa	0.715	0.774	yes	yes	yes	1.000
Qwen3-0.6B-Base	devtest/arb	–	–	–	–	–	–
Qwen3-0.6B-Base	devtest/cmn	0.942	0.509	yes	yes	yes	1.000
Qwen3-0.6B-Base	devtest/jpn	0.787	0.720	yes	yes	yes	1.000
Qwen3-0.6B-Base	fw/en	0.942	0.691	yes	yes	yes	1.000
Qwen3-0.6B-Base	fw/de	0.735	0.813	yes	yes	yes	1.000
Qwen3-0.6B-Base	fw/fr	0.792	0.783	yes	yes	yes	1.000
Qwen3-0.6B-Base	fw/es	0.771	0.778	yes	yes	yes	1.000
Qwen3-0.6B-Base	fw/ar	0.808	0.831	yes	yes	yes	1.000
Qwen3-0.6B-Base	fw/zh	0.993	0.264	yes	yes	no	0.008
Qwen3-0.6B-Base	fw/ja	0.920	0.602	yes	yes	yes	1.000
Qwen3-4B-Base	dev/eng	0.924	0.649	yes	yes	yes	1.000
Qwen3-4B-Base	dev/deu	0.706	0.767	yes	yes	yes	1.000
Qwen3-4B-Base	dev/fra	0.671	0.805	yes	yes	yes	1.000
Qwen3-4B-Base	dev/spa	0.761	0.795	yes	yes	yes	1.000
Qwen3-4B-Base	dev/arb	–	–	–	–	–	–
Qwen3-4B-Base	dev/cmn	0.957	0.543	yes	yes	yes	1.000
Qwen3-4B-Base	dev/jpn	0.823	0.717	yes	yes	yes	1.000
Qwen3-4B-Base	devtest/eng	0.918	0.644	yes	yes	yes	1.000
Qwen3-4B-Base	devtest/deu	–	–	–	–	–	–
Qwen3-4B-Base	devtest/fra	–	–	–	–	–	–
Qwen3-4B-Base	devtest/spa	–	–	–	–	–	–
Qwen3-4B-Base	devtest/arb	0.691	0.858	yes	yes	yes	1.000
Qwen3-4B-Base	devtest/cmn	0.961	0.543	yes	yes	yes	0.975
Qwen3-4B-Base	devtest/jpn	0.841	0.684	yes	yes	yes	1.000
Qwen3-4B-Base	fw/en	0.957	0.654	yes	yes	yes	1.000
Qwen3-4B-Base	fw/de	0.815	0.755	yes	yes	yes	1.000
Qwen3-4B-Base	fw/fr	0.771	0.786	yes	yes	yes	1.000
Qwen3-4B-Base	fw/es	0.829	0.736	yes	yes	yes	1.000
Qwen3-4B-Base	fw/ar	0.807	0.807	yes	yes	yes	1.000
Qwen3-4B-Base	fw/zh	0.995	0.280	yes	yes	no	0.004

Table 10 continued

Model	Shard	Radius	Test NMSE	Stable	Supported	Attractor	Conv.
Qwen3-4B-Base	fw/ja	0.919	0.638	yes	yes	yes	1.000
Qwen3-8B-Base	dev/eng	0.931	0.717	yes	yes	yes	1.000
Qwen3-8B-Base	dev/deu	–	–	–	–	–	–
Qwen3-8B-Base	dev/fra	0.773	0.776	yes	yes	yes	1.000
Qwen3-8B-Base	dev/spa	0.737	0.819	yes	yes	yes	1.000
Qwen3-8B-Base	dev/arb	–	–	–	–	–	–
Qwen3-8B-Base	dev/cmn	0.925	0.682	yes	yes	yes	1.000
Qwen3-8B-Base	dev/jpn	0.844	0.633	yes	yes	yes	1.000
Qwen3-8B-Base	devtest/eng	0.932	0.720	yes	yes	yes	1.000
Qwen3-8B-Base	devtest/deu	0.750	0.756	yes	yes	yes	1.000
Qwen3-8B-Base	devtest/fra	–	–	–	–	–	–
Qwen3-8B-Base	devtest/spa	–	–	–	–	–	–
Qwen3-8B-Base	devtest/arb	0.696	0.813	yes	yes	yes	1.000
Qwen3-8B-Base	devtest/cmn	0.919	0.703	yes	yes	yes	1.000
Qwen3-8B-Base	devtest/jpn	0.836	0.648	yes	yes	yes	1.000
Qwen3-8B-Base	fw/en	0.959	0.701	yes	yes	yes	1.000
Qwen3-8B-Base	fw/de	0.816	0.754	yes	yes	yes	1.000
Qwen3-8B-Base	fw/fr	0.843	0.761	yes	yes	yes	1.000
Qwen3-8B-Base	fw/es	0.829	0.744	yes	yes	yes	1.000
Qwen3-8B-Base	fw/ar	0.834	0.766	yes	yes	yes	1.000
Qwen3-8B-Base	fw/zh	0.992	0.379	yes	yes	no	0.008
Qwen3-8B-Base	fw/ja	0.940	0.557	yes	yes	yes	1.000
Qwen3-14B-Base	dev/eng	0.907	0.733	yes	yes	yes	1.000
Qwen3-14B-Base	dev/deu	0.763	0.759	yes	yes	yes	1.000
Qwen3-14B-Base	dev/fra	0.761	0.778	yes	yes	yes	1.000
Qwen3-14B-Base	dev/spa	0.695	0.815	yes	yes	yes	1.000
Qwen3-14B-Base	dev/arb	–	–	–	–	–	–
Qwen3-14B-Base	dev/cmn	0.914	0.631	yes	yes	yes	1.000
Qwen3-14B-Base	dev/jpn	0.851	0.700	yes	yes	yes	1.000
Qwen3-14B-Base	devtest/eng	0.898	0.717	yes	yes	yes	1.000
Qwen3-14B-Base	devtest/deu	0.757	0.756	yes	yes	yes	1.000
Qwen3-14B-Base	devtest/fra	0.748	0.782	yes	yes	yes	1.000
Qwen3-14B-Base	devtest/spa	0.751	0.780	yes	yes	yes	1.000
Qwen3-14B-Base	devtest/arb	0.704	0.799	yes	yes	yes	1.000
Qwen3-14B-Base	devtest/cmn	0.922	0.661	yes	yes	yes	1.000
Qwen3-14B-Base	devtest/jpn	0.841	0.713	yes	yes	yes	1.000
Qwen3-14B-Base	fw/en	0.917	0.822	yes	yes	yes	1.000
Qwen3-14B-Base	fw/de	0.793	0.779	yes	yes	yes	1.000
Qwen3-14B-Base	fw/fr	0.829	0.812	yes	yes	yes	1.000
Qwen3-14B-Base	fw/es	0.804	0.797	yes	yes	yes	1.000
Qwen3-14B-Base	fw/ar	0.802	0.778	yes	yes	yes	1.000
Qwen3-14B-Base	fw/zh	0.991	0.478	yes	yes	no	0.019
Qwen3-14B-Base	fw/ja	0.936	0.649	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	dev/eng	0.978	0.501	yes	yes	no	0.348
Qwen3-Embedding-0.6B	dev/deu	0.870	0.698	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	dev/fra	0.878	0.666	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	dev/spa	0.886	0.654	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	dev/arb	0.795	0.778	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	dev/cmn	0.972	0.513	yes	yes	no	0.798
Qwen3-Embedding-0.6B	dev/jpn	0.873	0.633	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	devtest/eng	0.967	0.522	yes	yes	no	0.848
Qwen3-Embedding-0.6B	devtest/deu	0.869	0.713	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	devtest/fra	0.885	0.672	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	devtest/spa	0.889	0.652	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	devtest/arb	0.783	0.777	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	devtest/cmn	0.964	0.516	yes	yes	yes	0.929
Qwen3-Embedding-0.6B	devtest/jpn	0.862	0.644	yes	yes	yes	1.000

Table 10 continued

Model	Shard	Radius	Test NMSE	Stable	Supported	Attractor	Conv.
Qwen3-Embedding-0.6B	fw/en	0.942	0.573	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	fw/de	0.861	0.731	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	fw/fr	0.784	0.710	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	fw/es	0.888	0.681	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	fw/ar	0.909	0.733	yes	yes	yes	1.000
Qwen3-Embedding-0.6B	fw/zh	0.980	0.473	yes	yes	no	0.083
Qwen3-Embedding-0.6B	fw/ja	0.872	0.683	yes	yes	yes	1.000
Qwen3-Embedding-4B	dev/eng	0.966	0.552	yes	yes	yes	0.925
Qwen3-Embedding-4B	dev/deu	0.942	0.640	yes	yes	yes	1.000
Qwen3-Embedding-4B	dev/fra	0.932	0.675	yes	yes	yes	1.000
Qwen3-Embedding-4B	dev/spa	0.940	0.639	yes	yes	yes	1.000
Qwen3-Embedding-4B	dev/arb	0.929	0.731	yes	yes	yes	1.000
Qwen3-Embedding-4B	dev/cmn	0.956	0.538	yes	yes	yes	1.000
Qwen3-Embedding-4B	dev/jpn	0.940	0.625	yes	yes	yes	1.000
Qwen3-Embedding-4B	devtest/eng	0.957	0.566	yes	yes	yes	1.000
Qwen3-Embedding-4B	devtest/deu	0.937	0.658	yes	yes	yes	1.000
Qwen3-Embedding-4B	devtest/fra	0.931	0.678	yes	yes	yes	1.000
Qwen3-Embedding-4B	devtest/spa	0.938	0.644	yes	yes	yes	1.000
Qwen3-Embedding-4B	devtest/arb	0.924	0.733	yes	yes	yes	1.000
Qwen3-Embedding-4B	devtest/cmn	0.956	0.539	yes	yes	yes	1.000
Qwen3-Embedding-4B	devtest/jpn	0.933	0.633	yes	yes	yes	1.000
Qwen3-Embedding-4B	fw/en	0.951	0.642	yes	yes	yes	1.000
Qwen3-Embedding-4B	fw/de	0.900	0.696	yes	yes	yes	1.000
Qwen3-Embedding-4B	fw/fr	0.896	0.729	yes	yes	yes	1.000
Qwen3-Embedding-4B	fw/es	0.926	0.661	yes	yes	yes	1.000
Qwen3-Embedding-4B	fw/ar	0.921	0.690	yes	yes	yes	1.000
Qwen3-Embedding-4B	fw/zh	0.943	0.586	yes	yes	yes	1.000
Qwen3-Embedding-4B	fw/ja	0.901	0.639	yes	yes	yes	1.000
Qwen3-Embedding-8B	dev/eng	0.971	0.623	yes	yes	no	0.860
Qwen3-Embedding-8B	dev/deu	0.857	0.737	yes	yes	yes	1.000
Qwen3-Embedding-8B	dev/fra	0.851	0.778	yes	yes	yes	1.000
Qwen3-Embedding-8B	dev/spa	0.881	0.750	yes	yes	yes	1.000
Qwen3-Embedding-8B	dev/arb	0.921	0.731	yes	yes	yes	1.000
Qwen3-Embedding-8B	dev/cmn	0.958	0.584	yes	yes	yes	0.998
Qwen3-Embedding-8B	dev/jpn	0.932	0.640	yes	yes	yes	1.000
Qwen3-Embedding-8B	devtest/eng	0.973	0.609	yes	yes	no	0.834
Qwen3-Embedding-8B	devtest/deu	0.847	0.752	yes	yes	yes	1.000
Qwen3-Embedding-8B	devtest/fra	0.842	0.786	yes	yes	yes	1.000
Qwen3-Embedding-8B	devtest/spa	0.867	0.749	yes	yes	yes	1.000
Qwen3-Embedding-8B	devtest/arb	0.933	0.714	yes	yes	yes	1.000
Qwen3-Embedding-8B	devtest/cmn	0.965	0.581	yes	yes	yes	0.943
Qwen3-Embedding-8B	devtest/jpn	0.931	0.645	yes	yes	yes	1.000
Qwen3-Embedding-8B	fw/en	0.951	0.692	yes	yes	yes	1.000
Qwen3-Embedding-8B	fw/de	0.792	0.775	yes	yes	yes	1.000
Qwen3-Embedding-8B	fw/fr	0.837	0.771	yes	yes	yes	1.000
Qwen3-Embedding-8B	fw/es	0.919	0.747	yes	yes	yes	1.000
Qwen3-Embedding-8B	fw/ar	0.918	0.742	yes	yes	yes	1.000
Qwen3-Embedding-8B	fw/zh	0.954	0.594	yes	yes	yes	1.000
Qwen3-Embedding-8B	fw/ja	0.894	0.654	yes	yes	yes	1.000
bge-m3	dev/eng	0.995	0.125	yes	no	no	0.000
bge-m3	dev/deu	0.995	0.105	yes	yes	no	0.000
bge-m3	dev/fra	0.997	0.095	yes	no	no	0.000
bge-m3	dev/spa	1.000	0.136	yes	no	no	0.000
bge-m3	dev/arb	0.992	0.093	yes	yes	no	0.000
bge-m3	dev/cmn	0.997	0.105	yes	yes	no	0.000
bge-m3	dev/jpn	0.994	0.074	yes	yes	no	0.000
bge-m3	devtest/eng	0.995	0.109	yes	yes	no	0.000

Table 10 continued

Model	Shard	Radius	Test NMSE	Stable	Supported	Attractor	Conv.
bge-m3	devtest/deu	0.992	0.099	yes	yes	no	0.000
bge-m3	devtest/fra	0.996	0.088	yes	yes	no	0.000
bge-m3	devtest/spa	0.999	0.115	yes	no	no	0.000
bge-m3	devtest/arb	0.989	0.083	yes	yes	no	0.000
bge-m3	devtest/cmn	0.996	0.101	yes	yes	no	0.000
bge-m3	devtest/jpn	0.991	0.074	yes	yes	no	0.000
bge-m3	fw/en	1.000	0.001	yes	yes	no	0.000
bge-m3	fw/de	1.000	0.001	yes	yes	no	0.000
bge-m3	fw/fr	0.999	0.005	yes	yes	no	0.000
bge-m3	fw/es	1.000	0.001	yes	yes	no	0.000
bge-m3	fw/ar	1.000	0.001	no	no	no	0.000
bge-m3	fw/zh	1.000	0.001	yes	yes	no	0.000
bge-m3	fw/ja	1.000	0.001	yes	yes	no	0.000
pplx-embed-v1-0.6b	dev/eng	0.981	0.132	yes	yes	no	0.049
pplx-embed-v1-0.6b	dev/deu	0.988	0.096	yes	yes	no	0.021
pplx-embed-v1-0.6b	dev/fra	0.980	0.107	yes	yes	no	0.093
pplx-embed-v1-0.6b	dev/spa	0.985	0.098	yes	yes	no	0.039
pplx-embed-v1-0.6b	dev/arb	0.987	0.077	yes	yes	no	0.021
pplx-embed-v1-0.6b	dev/cmn	0.976	0.098	yes	yes	no	0.051
pplx-embed-v1-0.6b	dev/jpn	0.989	0.076	yes	yes	no	0.003
pplx-embed-v1-0.6b	devtest/eng	0.987	0.114	yes	yes	no	0.003
pplx-embed-v1-0.6b	devtest/deu	0.987	0.098	yes	yes	no	0.015
pplx-embed-v1-0.6b	devtest/fra	0.985	0.111	yes	yes	no	0.032
pplx-embed-v1-0.6b	devtest/spa	0.988	0.095	yes	yes	no	0.020
pplx-embed-v1-0.6b	devtest/arb	0.988	0.077	yes	yes	no	0.014
pplx-embed-v1-0.6b	devtest/cmn	0.984	0.108	yes	yes	no	0.021
pplx-embed-v1-0.6b	devtest/jpn	0.987	0.079	yes	yes	no	0.006
pplx-embed-v1-0.6b	fw/en	0.999	0.009	yes	no	no	0.000
pplx-embed-v1-0.6b	fw/de	0.999	0.007	yes	no	no	0.000
pplx-embed-v1-0.6b	fw/fr	0.999	0.006	yes	no	no	0.000
pplx-embed-v1-0.6b	fw/es	0.999	0.005	yes	no	no	0.000
pplx-embed-v1-0.6b	fw/ar	0.999	0.007	yes	no	no	0.000
pplx-embed-v1-0.6b	fw/zh	0.999	0.007	yes	no	no	0.000
pplx-embed-v1-0.6b	fw/ja	0.999	0.006	yes	no	no	0.000
pplx-embed-v1-4b	dev/eng	1.008	0.058	no	no	no	0.000
pplx-embed-v1-4b	dev/deu	1.001	0.046	no	no	no	0.000
pplx-embed-v1-4b	dev/fra	1.001	0.044	no	no	no	0.000
pplx-embed-v1-4b	dev/spa	1.002	0.045	no	no	no	0.000
pplx-embed-v1-4b	dev/arb	1.000	0.049	yes	no	no	0.000
pplx-embed-v1-4b	dev/cmn	1.001	0.052	no	no	no	0.000
pplx-embed-v1-4b	dev/jpn	0.999	0.048	yes	no	no	0.000
pplx-embed-v1-4b	devtest/eng	1.005	0.059	no	no	no	0.000
pplx-embed-v1-4b	devtest/deu	0.998	0.044	yes	no	no	0.000
pplx-embed-v1-4b	devtest/fra	0.998	0.046	yes	no	no	0.000
pplx-embed-v1-4b	devtest/spa	0.998	0.044	yes	no	no	0.000
pplx-embed-v1-4b	devtest/arb	0.997	0.047	yes	yes	no	0.000
pplx-embed-v1-4b	devtest/cmn	0.999	0.049	yes	no	no	0.000
pplx-embed-v1-4b	devtest/jpn	0.997	0.045	yes	yes	no	0.000
pplx-embed-v1-4b	fw/en	1.000	0.002	yes	no	no	0.000
pplx-embed-v1-4b	fw/de	1.000	0.004	yes	no	no	0.000
pplx-embed-v1-4b	fw/fr	1.000	0.003	yes	no	no	0.000
pplx-embed-v1-4b	fw/es	1.000	0.002	yes	no	no	0.000
pplx-embed-v1-4b	fw/ar	0.999	0.003	yes	no	no	0.000
pplx-embed-v1-4b	fw/zh	1.000	0.002	yes	no	no	0.000
pplx-embed-v1-4b	fw/ja	1.000	0.003	yes	no	no	0.000
text-embedding-3-small	dev/eng	1.001	0.047	no	no	no	0.000
text-embedding-3-small	dev/deu	1.001	0.044	no	no	no	0.000

Table 10 continued

Model	Shard	Radius	Test NMSE	Stable	Supported	Attractor	Conv.
text-embedding-3-small	dev/fra	0.998	0.049	yes	no	no	0.000
text-embedding-3-small	dev/spa	1.000	0.050	yes	no	no	0.000
text-embedding-3-small	dev/arb	1.000	0.032	yes	no	no	0.000
text-embedding-3-small	dev/cmn	1.001	0.034	no	no	no	0.000
text-embedding-3-small	dev/jpn	0.998	0.031	yes	yes	no	0.000
text-embedding-3-small	devtest/eng	0.998	0.045	yes	no	no	0.000
text-embedding-3-small	devtest/deu	0.999	0.042	yes	no	no	0.000
text-embedding-3-small	devtest/fra	0.999	0.045	yes	no	no	0.000
text-embedding-3-small	devtest/spa	0.999	0.044	yes	no	no	0.000
text-embedding-3-small	devtest/arb	1.000	0.032	no	no	no	0.000
text-embedding-3-small	devtest/cmn	1.001	0.032	no	no	no	0.000
text-embedding-3-small	devtest/jpn	1.000	0.031	no	no	no	0.000
text-embedding-3-large	dev/eng	1.002	0.076	no	no	no	0.000
text-embedding-3-large	dev/deu	1.004	0.150	no	no	no	0.000
text-embedding-3-large	dev/fra	1.003	0.136	no	yes	no	0.000
text-embedding-3-large	dev/spa	1.002	0.143	no	yes	no	0.000
text-embedding-3-large	dev/arb	1.001	0.124	no	no	no	0.000
text-embedding-3-large	dev/cmn	1.001	0.065	no	no	no	0.000
text-embedding-3-large	dev/jpn	1.002	0.066	no	no	no	0.000
text-embedding-3-large	devtest/eng	0.999	0.078	yes	no	no	0.000
text-embedding-3-large	devtest/deu	1.004	0.147	no	no	no	0.000
text-embedding-3-large	devtest/fra	1.001	0.140	no	no	no	0.000
text-embedding-3-large	devtest/spa	1.000	0.152	no	yes	no	0.000
text-embedding-3-large	devtest/arb	1.000	0.131	yes	no	no	0.000
text-embedding-3-large	devtest/cmn	1.000	0.071	no	yes	no	0.000
text-embedding-3-large	devtest/jpn	1.000	0.066	no	no	no	0.000
voyage-4-lite	dev/eng	1.000	0.052	no	no	no	0.000
voyage-4-lite	dev/deu	1.001	0.048	no	no	no	0.000
voyage-4-lite	dev/fra	1.001	0.043	no	no	no	0.000
voyage-4-lite	dev/spa	1.000	0.049	yes	no	no	0.000
voyage-4-lite	dev/arb	1.000	0.044	yes	no	no	0.000
voyage-4-lite	dev/cmn	1.000	0.044	yes	no	no	0.000
voyage-4-lite	dev/jpn	1.000	0.044	yes	no	no	0.000
voyage-4-lite	devtest/eng	0.997	0.046	yes	no	no	0.000
voyage-4-lite	devtest/deu	0.996	0.045	yes	no	no	0.000
voyage-4-lite	devtest/fra	0.998	0.041	yes	no	no	0.000
voyage-4-lite	devtest/spa	0.997	0.045	yes	no	no	0.000
voyage-4-lite	devtest/arb	0.998	0.046	yes	no	no	0.000
voyage-4-lite	devtest/cmn	0.997	0.049	yes	no	no	0.000
voyage-4-lite	devtest/jpn	1.004	0.040	no	no	no	0.000
voyage-4	dev/eng	1.000	0.047	yes	no	no	0.000
voyage-4	dev/deu	1.001	0.041	no	no	no	0.000
voyage-4	dev/fra	1.002	0.038	no	no	no	0.000
voyage-4	dev/spa	0.999	0.039	yes	no	no	0.000
voyage-4	dev/arb	1.001	0.039	no	no	no	0.000
voyage-4	dev/cmn	0.997	0.038	yes	no	no	0.000
voyage-4	dev/jpn	1.002	0.038	no	no	no	0.000
voyage-4	devtest/eng	0.996	0.040	yes	no	no	0.000
voyage-4	devtest/deu	0.998	0.040	yes	no	no	0.000
voyage-4	devtest/fra	1.001	0.036	no	no	no	0.000
voyage-4	devtest/spa	0.999	0.038	yes	no	no	0.000
voyage-4	devtest/arb	1.000	0.041	yes	no	no	0.000
voyage-4	devtest/cmn	0.998	0.043	yes	no	no	0.000
voyage-4	devtest/jpn	1.004	0.039	no	no	no	0.000
voyage-4-large	dev/eng	0.998	0.050	yes	no	no	0.000
voyage-4-large	dev/deu	0.999	0.035	yes	no	no	0.000
voyage-4-large	dev/fra	1.001	0.035	no	no	no	0.000

Table 10 continued

Model	Shard	Radius	Test NMSE	Stable	Supported	Attractor	Conv.
voyage-4-large	dev/spa	0.999	0.040	yes	no	no	0.000
voyage-4-large	dev/arb	0.998	0.045	yes	no	no	0.000
voyage-4-large	dev/cmn	1.001	0.040	no	no	no	0.000
voyage-4-large	dev/jpn	0.999	0.042	yes	no	no	0.000
voyage-4-large	devtest/eng	0.999	0.046	yes	no	no	0.000
voyage-4-large	devtest/deu	0.997	0.037	yes	no	no	0.000
voyage-4-large	devtest/fra	0.999	0.035	yes	no	no	0.000
voyage-4-large	devtest/spa	0.998	0.038	yes	no	no	0.000
voyage-4-large	devtest/arb	0.995	0.045	yes	yes	no	0.000
voyage-4-large	devtest/cmn	1.000	0.046	no	no	no	0.000
voyage-4-large	devtest/jpn	1.004	0.039	no	no	no	0.000

B.4 Recurrence and stability

Table 11: Conditional recurrence and stability inventory. Metastable is the occupancy-weighted self-transition probability of the eight-region partition. Lyapunov is the largest exponent of the affine surrogate's constant Jacobian and therefore summarizes the same local timescale as its spectrum; it is not an exponent of the model's internal state update. Rejected shards are left blank.

Model	Shard	Metastable	Lyapunov
Qwen3-0.6B-Base	dev/eng	0.268	-0.076
Qwen3-0.6B-Base	dev/deu	–	–
Qwen3-0.6B-Base	dev/fra	0.197	-0.435
Qwen3-0.6B-Base	dev/spa	0.191	-0.335
Qwen3-0.6B-Base	dev/arb	0.168	-0.366
Qwen3-0.6B-Base	dev/cmn	0.596	-0.058
Qwen3-0.6B-Base	dev/jpn	0.240	-0.233
Qwen3-0.6B-Base	devtest/eng	0.244	-0.081
Qwen3-0.6B-Base	devtest/deu	–	–
Qwen3-0.6B-Base	devtest/fra	–	–
Qwen3-0.6B-Base	devtest/spa	0.199	-0.336
Qwen3-0.6B-Base	devtest/arb	–	–
Qwen3-0.6B-Base	devtest/cmn	0.577	-0.060
Qwen3-0.6B-Base	devtest/jpn	0.237	-0.240
Qwen3-0.6B-Base	fw/en	0.193	-0.060
Qwen3-0.6B-Base	fw/de	0.228	-0.308
Qwen3-0.6B-Base	fw/fr	0.199	-0.233
Qwen3-0.6B-Base	fw/es	0.235	-0.260
Qwen3-0.6B-Base	fw/ar	0.201	-0.214
Qwen3-0.6B-Base	fw/zh	0.844	-0.007
Qwen3-0.6B-Base	fw/ja	0.298	-0.084
Qwen3-4B-Base	dev/eng	0.268	-0.079
Qwen3-4B-Base	dev/deu	0.219	-0.348
Qwen3-4B-Base	dev/fra	0.199	-0.399
Qwen3-4B-Base	dev/spa	0.186	-0.273
Qwen3-4B-Base	dev/arb	–	–
Qwen3-4B-Base	dev/cmn	0.427	-0.044
Qwen3-4B-Base	dev/jpn	0.235	-0.195
Qwen3-4B-Base	devtest/eng	0.292	-0.086
Qwen3-4B-Base	devtest/deu	–	–
Qwen3-4B-Base	devtest/fra	–	–
Qwen3-4B-Base	devtest/spa	–	–
Qwen3-4B-Base	devtest/arb	0.167	-0.370

Table 11 continued

Model	Shard	Metastable	Lyapunov
Qwen3-4B-Base	devtest/cmn	0.415	-0.040
Qwen3-4B-Base	devtest/jpn	0.240	-0.173
Qwen3-4B-Base	fw/en	0.222	-0.044
Qwen3-4B-Base	fw/de	0.227	-0.204
Qwen3-4B-Base	fw/fr	0.210	-0.260
Qwen3-4B-Base	fw/es	0.219	-0.188
Qwen3-4B-Base	fw/ar	0.174	-0.215
Qwen3-4B-Base	fw/zh	0.827	-0.005
Qwen3-4B-Base	fw/ja	0.275	-0.085
Qwen3-8B-Base	dev/eng	0.258	-0.071
Qwen3-8B-Base	dev/deu	–	–
Qwen3-8B-Base	dev/fra	0.167	-0.257
Qwen3-8B-Base	dev/spa	0.147	-0.306
Qwen3-8B-Base	dev/arb	–	–
Qwen3-8B-Base	dev/cmn	0.340	-0.078
Qwen3-8B-Base	dev/jpn	0.287	-0.169
Qwen3-8B-Base	devtest/eng	0.271	-0.070
Qwen3-8B-Base	devtest/deu	0.219	-0.288
Qwen3-8B-Base	devtest/fra	–	–
Qwen3-8B-Base	devtest/spa	–	–
Qwen3-8B-Base	devtest/arb	0.251	-0.363
Qwen3-8B-Base	devtest/cmn	0.326	-0.084
Qwen3-8B-Base	devtest/jpn	0.238	-0.179
Qwen3-8B-Base	fw/en	0.271	-0.041
Qwen3-8B-Base	fw/de	0.233	-0.204
Qwen3-8B-Base	fw/fr	0.175	-0.171
Qwen3-8B-Base	fw/es	0.164	-0.188
Qwen3-8B-Base	fw/ar	0.201	-0.181
Qwen3-8B-Base	fw/zh	0.688	-0.008
Qwen3-8B-Base	fw/ja	0.331	-0.062
Qwen3-14B-Base	dev/eng	0.244	-0.098
Qwen3-14B-Base	dev/deu	0.259	-0.270
Qwen3-14B-Base	dev/fra	0.212	-0.273
Qwen3-14B-Base	dev/spa	0.220	-0.364
Qwen3-14B-Base	dev/arb	–	–
Qwen3-14B-Base	dev/cmn	0.375	-0.090
Qwen3-14B-Base	dev/jpn	0.240	-0.162
Qwen3-14B-Base	devtest/eng	0.254	-0.108
Qwen3-14B-Base	devtest/deu	0.270	-0.278
Qwen3-14B-Base	devtest/fra	0.222	-0.291
Qwen3-14B-Base	devtest/spa	0.222	-0.286
Qwen3-14B-Base	devtest/arb	0.188	-0.351
Qwen3-14B-Base	devtest/cmn	0.304	-0.081
Qwen3-14B-Base	devtest/jpn	0.239	-0.173
Qwen3-14B-Base	fw/en	0.189	-0.087
Qwen3-14B-Base	fw/de	0.259	-0.232
Qwen3-14B-Base	fw/fr	0.202	-0.188
Qwen3-14B-Base	fw/es	0.219	-0.218
Qwen3-14B-Base	fw/ar	0.268	-0.220
Qwen3-14B-Base	fw/zh	0.735	-0.009
Qwen3-14B-Base	fw/ja	0.289	-0.066
Qwen3-Embedding-0.6B	dev/eng	0.606	-0.023
Qwen3-Embedding-0.6B	dev/deu	0.285	-0.139
Qwen3-Embedding-0.6B	dev/fra	0.289	-0.131
Qwen3-Embedding-0.6B	dev/spa	0.263	-0.121
Qwen3-Embedding-0.6B	dev/arb	0.262	-0.229
Qwen3-Embedding-0.6B	dev/cmn	0.813	-0.028

Table 11 continued

Model	Shard	Metastable	Lyapunov
Qwen3-Embedding-0.6B	dev/jpn	0.310	-0.136
Qwen3-Embedding-0.6B	devtest/eng	0.603	-0.033
Qwen3-Embedding-0.6B	devtest/deu	0.298	-0.141
Qwen3-Embedding-0.6B	devtest/fra	0.287	-0.122
Qwen3-Embedding-0.6B	devtest/spa	0.301	-0.117
Qwen3-Embedding-0.6B	devtest/arb	0.260	-0.245
Qwen3-Embedding-0.6B	devtest/cmn	0.812	-0.037
Qwen3-Embedding-0.6B	devtest/jpn	0.304	-0.148
Qwen3-Embedding-0.6B	fw/en	0.316	-0.060
Qwen3-Embedding-0.6B	fw/de	0.294	-0.149
Qwen3-Embedding-0.6B	fw/fr	0.262	-0.243
Qwen3-Embedding-0.6B	fw/es	0.260	-0.119
Qwen3-Embedding-0.6B	fw/ar	0.297	-0.095
Qwen3-Embedding-0.6B	fw/zh	0.814	-0.020
Qwen3-Embedding-0.6B	fw/ja	0.318	-0.137
Qwen3-Embedding-4B	dev/eng	0.286	-0.034
Qwen3-Embedding-4B	dev/deu	0.285	-0.060
Qwen3-Embedding-4B	dev/fra	0.229	-0.070
Qwen3-Embedding-4B	dev/spa	0.237	-0.062
Qwen3-Embedding-4B	dev/arb	0.220	-0.074
Qwen3-Embedding-4B	dev/cmn	0.403	-0.045
Qwen3-Embedding-4B	dev/jpn	0.383	-0.062
Qwen3-Embedding-4B	devtest/eng	0.420	-0.044
Qwen3-Embedding-4B	devtest/deu	0.264	-0.065
Qwen3-Embedding-4B	devtest/fra	0.227	-0.072
Qwen3-Embedding-4B	devtest/spa	0.255	-0.064
Qwen3-Embedding-4B	devtest/arb	0.214	-0.079
Qwen3-Embedding-4B	devtest/cmn	0.472	-0.045
Qwen3-Embedding-4B	devtest/jpn	0.455	-0.069
Qwen3-Embedding-4B	fw/en	0.261	-0.050
Qwen3-Embedding-4B	fw/de	0.324	-0.106
Qwen3-Embedding-4B	fw/fr	0.235	-0.110
Qwen3-Embedding-4B	fw/es	0.250	-0.077
Qwen3-Embedding-4B	fw/ar	0.255	-0.082
Qwen3-Embedding-4B	fw/zh	0.545	-0.059
Qwen3-Embedding-4B	fw/ja	0.357	-0.105
Qwen3-Embedding-8B	dev/eng	0.239	-0.029
Qwen3-Embedding-8B	dev/deu	0.249	-0.154
Qwen3-Embedding-8B	dev/fra	0.178	-0.162
Qwen3-Embedding-8B	dev/spa	0.198	-0.127
Qwen3-Embedding-8B	dev/arb	0.229	-0.083
Qwen3-Embedding-8B	dev/cmn	0.392	-0.043
Qwen3-Embedding-8B	dev/jpn	0.260	-0.071
Qwen3-Embedding-8B	devtest/eng	0.269	-0.028
Qwen3-Embedding-8B	devtest/deu	0.233	-0.166
Qwen3-Embedding-8B	devtest/fra	0.193	-0.172
Qwen3-Embedding-8B	devtest/spa	0.178	-0.143
Qwen3-Embedding-8B	devtest/arb	0.186	-0.070
Qwen3-Embedding-8B	devtest/cmn	0.396	-0.036
Qwen3-Embedding-8B	devtest/jpn	0.257	-0.071
Qwen3-Embedding-8B	fw/en	0.292	-0.050
Qwen3-Embedding-8B	fw/de	0.258	-0.234
Qwen3-Embedding-8B	fw/fr	0.192	-0.178
Qwen3-Embedding-8B	fw/es	0.186	-0.084
Qwen3-Embedding-8B	fw/ar	0.198	-0.085
Qwen3-Embedding-8B	fw/zh	0.476	-0.047
Qwen3-Embedding-8B	fw/ja	0.319	-0.113

Table 11 continued

Model	Shard	Metastable	Lyapunov
bge-m3	dev/eng	0.862	-0.005
bge-m3	dev/deu	0.878	-0.005
bge-m3	dev/fra	0.892	-0.004
bge-m3	dev/spa	0.876	-0.001
bge-m3	dev/arb	0.881	-0.008
bge-m3	dev/cmn	0.865	-0.004
bge-m3	dev/jpn	0.886	-0.006
bge-m3	devtest/eng	0.888	-0.006
bge-m3	devtest/deu	0.901	-0.008
bge-m3	devtest/fra	0.899	-0.004
bge-m3	devtest/spa	0.888	-0.002
bge-m3	devtest/arb	0.885	-0.012
bge-m3	devtest/cmn	0.863	-0.004
bge-m3	devtest/jpn	0.906	-0.010
bge-m3	fw/en	0.997	-0.000
bge-m3	fw/de	0.996	-0.000
bge-m3	fw/fr	0.993	-0.001
bge-m3	fw/es	0.996	-0.000
bge-m3	fw/ar	0.996	-0.000
bge-m3	fw/zh	0.996	-0.000
bge-m3	fw/ja	0.995	-0.000
pplx-embed-v1-0.6b	dev/eng	0.884	-0.020
pplx-embed-v1-0.6b	dev/deu	0.899	-0.012
pplx-embed-v1-0.6b	dev/fra	0.901	-0.020
pplx-embed-v1-0.6b	dev/spa	0.896	-0.015
pplx-embed-v1-0.6b	dev/arb	0.906	-0.013
pplx-embed-v1-0.6b	dev/cmn	0.902	-0.025
pplx-embed-v1-0.6b	dev/jpn	0.917	-0.011
pplx-embed-v1-0.6b	devtest/eng	0.883	-0.013
pplx-embed-v1-0.6b	devtest/deu	0.930	-0.014
pplx-embed-v1-0.6b	devtest/fra	0.894	-0.015
pplx-embed-v1-0.6b	devtest/spa	0.909	-0.012
pplx-embed-v1-0.6b	devtest/arb	0.922	-0.013
pplx-embed-v1-0.6b	devtest/cmn	0.902	-0.016
pplx-embed-v1-0.6b	devtest/jpn	0.914	-0.013
pplx-embed-v1-0.6b	fw/en	0.996	-0.001
pplx-embed-v1-0.6b	fw/de	0.995	-0.001
pplx-embed-v1-0.6b	fw/fr	0.996	-0.001
pplx-embed-v1-0.6b	fw/es	0.996	-0.001
pplx-embed-v1-0.6b	fw/ar	0.996	-0.001
pplx-embed-v1-0.6b	fw/zh	0.995	-0.001
pplx-embed-v1-0.6b	fw/ja	0.995	-0.001
pplx-embed-v1-4b	dev/eng	0.932	0.008
pplx-embed-v1-4b	dev/deu	0.950	0.001
pplx-embed-v1-4b	dev/fra	0.947	0.001
pplx-embed-v1-4b	dev/spa	0.948	0.002
pplx-embed-v1-4b	dev/arb	0.947	-0.000
pplx-embed-v1-4b	dev/cmn	0.932	0.001
pplx-embed-v1-4b	dev/jpn	0.947	-0.001
pplx-embed-v1-4b	devtest/eng	0.930	0.005
pplx-embed-v1-4b	devtest/deu	0.946	-0.002
pplx-embed-v1-4b	devtest/fra	0.949	-0.002
pplx-embed-v1-4b	devtest/spa	0.951	-0.002
pplx-embed-v1-4b	devtest/arb	0.946	-0.003
pplx-embed-v1-4b	devtest/cmn	0.930	-0.001
pplx-embed-v1-4b	devtest/jpn	0.945	-0.004
pplx-embed-v1-4b	fw/en	0.997	-0.000

Table 11 continued

Model	Shard	Metastable	Lyapunov
pplx-embed-v1-4b	fw/de	0.997	-0.000
pplx-embed-v1-4b	fw/fr	0.998	-0.000
pplx-embed-v1-4b	fw/es	0.998	-0.000
pplx-embed-v1-4b	fw/ar	0.996	-0.001
pplx-embed-v1-4b	fw/zh	0.996	-0.000
pplx-embed-v1-4b	fw/ja	0.996	-0.000
text-embedding-3-small	dev/eng	0.929	-0.001
text-embedding-3-small	dev/deu	0.941	0.001
text-embedding-3-small	dev/fra	0.940	-0.002
text-embedding-3-small	dev/spa	0.940	-0.001
text-embedding-3-small	dev/arb	0.961	-0.000
text-embedding-3-small	dev/cmn	0.948	0.001
text-embedding-3-small	dev/jpn	0.955	-0.002
text-embedding-3-small	devtest/eng	0.937	-0.001
text-embedding-3-small	devtest/deu	0.941	-0.000
text-embedding-3-small	devtest/fra	0.943	-0.001
text-embedding-3-small	devtest/spa	0.944	-0.002
text-embedding-3-small	devtest/arb	0.961	0.000
text-embedding-3-small	devtest/cmn	0.952	0.000
text-embedding-3-small	devtest/jpn	0.958	-0.000
text-embedding-3-large	dev/eng	0.915	0.002
text-embedding-3-large	dev/deu	0.955	0.004
text-embedding-3-large	dev/fra	0.943	0.003
text-embedding-3-large	dev/spa	0.943	0.002
text-embedding-3-large	dev/arb	0.950	0.001
text-embedding-3-large	dev/cmn	0.945	0.001
text-embedding-3-large	dev/jpn	0.955	0.002
text-embedding-3-large	devtest/eng	0.923	-0.001
text-embedding-3-large	devtest/deu	0.951	0.004
text-embedding-3-large	devtest/fra	0.942	0.001
text-embedding-3-large	devtest/spa	0.941	0.000
text-embedding-3-large	devtest/arb	0.934	-0.001
text-embedding-3-large	devtest/cmn	0.951	0.000
text-embedding-3-large	devtest/jpn	0.956	0.000
voyage-4-lite	dev/eng	0.920	0.000
voyage-4-lite	dev/deu	0.943	0.001
voyage-4-lite	dev/fra	0.941	0.001
voyage-4-lite	dev/spa	0.941	-0.000
voyage-4-lite	dev/arb	0.949	-0.000
voyage-4-lite	dev/cmn	0.922	-0.000
voyage-4-lite	dev/jpn	0.954	-0.000
voyage-4-lite	devtest/eng	0.925	-0.003
voyage-4-lite	devtest/deu	0.947	-0.004
voyage-4-lite	devtest/fra	0.949	-0.002
voyage-4-lite	devtest/spa	0.945	-0.003
voyage-4-lite	devtest/arb	0.949	-0.001
voyage-4-lite	devtest/cmn	0.927	-0.004
voyage-4-lite	devtest/jpn	0.947	0.004
voyage-4	dev/eng	0.929	-0.000
voyage-4	dev/deu	0.945	0.001
voyage-4	dev/fra	0.950	0.002
voyage-4	dev/spa	0.939	-0.001
voyage-4	dev/arb	0.949	0.001
voyage-4	dev/cmn	0.940	-0.002
voyage-4	dev/jpn	0.951	0.001
voyage-4	devtest/eng	0.931	-0.004
voyage-4	devtest/deu	0.950	-0.002

Table 11 continued

Model	Shard	Metastable	Lyapunov
voyage-4	devtest/fra	0.949	0.001
voyage-4	devtest/spa	0.949	-0.001
voyage-4	devtest/arb	0.952	-0.001
voyage-4	devtest/cmn	0.936	-0.003
voyage-4	devtest/jpn	0.953	0.004
voyage-4-large	dev/eng	0.921	-0.002
voyage-4-large	dev/deu	0.944	-0.000
voyage-4-large	dev/fra	0.944	0.001
voyage-4-large	dev/spa	0.940	-0.001
voyage-4-large	dev/arb	0.945	-0.002
voyage-4-large	dev/cmn	0.906	0.001
voyage-4-large	dev/jpn	0.939	-0.001
voyage-4-large	devtest/eng	0.922	-0.002
voyage-4-large	devtest/deu	0.945	-0.003
voyage-4-large	devtest/fra	0.946	-0.001
voyage-4-large	devtest/spa	0.941	-0.003
voyage-4-large	devtest/arb	0.934	-0.005
voyage-4-large	devtest/cmn	0.910	-0.001
voyage-4-large	devtest/jpn	0.946	0.004

B.5 Paired token perturbations

The table lists the complete Euclidean summary grid over the three causal checkpoints, four edit-position slots, both replacement arms, and all delayed-release values used in the validation panel. Scale-free cosine gaps, sustained half-lives, and finite-time output rates for the principal 256-token anchor are reported in Table 4; full aligned curves and support counts remain in the machine-readable outputs.

Table 12: Complete Euclidean summary grid for the English FineWeb perturbation validation slice. Immediate is the deviation at the edited token; Late is the median over the final quarter of common support; Decay is Late/Immediate; Release is the deviation at the first freely generated step and is defined only for the delayed-release sweep. Values first take the within-document branch median and then the cohort median. The 0.6B, 4B, and 8B cohorts contain 24, 28, and 27 usable documents, respectively, at anchor 256.

Model	Run	Group	Immediate	Late	Decay	Release
Qwen3-0.6B-Base	forced	plausible/anchor-256	93.190	0.615	0.008	–
Qwen3-0.6B-Base	forced	random/anchor-256	119.657	2.147	0.017	–
Qwen3-0.6B-Base	forced	plausible/anchor-512	82.250	0.581	0.008	–
Qwen3-0.6B-Base	forced	random/anchor-512	126.366	2.076	0.016	–
Qwen3-0.6B-Base	forced	plausible/uniform-0	87.053	0.647	0.008	–
Qwen3-0.6B-Base	forced	random/uniform-0	127.758	2.018	0.017	–
Qwen3-0.6B-Base	forced	plausible/uniform-1	90.357	0.611	0.009	–
Qwen3-0.6B-Base	forced	random/uniform-1	116.094	2.292	0.018	–
Qwen3-0.6B-Base	sampled	plausible/anchor-256	93.190	101.023	1.095	–
Qwen3-0.6B-Base	sampled	random/anchor-256	119.657	102.660	0.835	–
Qwen3-0.6B-Base	sampled	plausible/anchor-512	82.250	101.336	1.210	–
Qwen3-0.6B-Base	sampled	random/anchor-512	126.366	102.310	0.857	–
Qwen3-0.6B-Base	sampled	plausible/uniform-0	87.053	101.743	1.205	–
Qwen3-0.6B-Base	sampled	random/uniform-0	127.758	103.381	0.848	–
Qwen3-0.6B-Base	sampled	plausible/uniform-1	90.357	102.630	1.161	–
Qwen3-0.6B-Base	sampled	random/uniform-1	116.094	104.448	0.910	–
Qwen3-0.6B-Base	sweep	plausible/anchor-256/delta-1	98.909	100.842	1.173	69.691

Table 12 continued

Model	Run	Group	Immediate	Late	Decay	Release
Qwen3-0.6B-Base	sweep	random/anchor-256/delta-1	123.480	106.419	0.871	52.132
Qwen3-0.6B-Base	sweep	plausible/anchor-256/delta-4	98.909	101.257	1.118	13.797
Qwen3-0.6B-Base	sweep	random/anchor-256/delta-4	123.480	102.494	0.851	24.512
Qwen3-0.6B-Base	sweep	plausible/anchor-256/delta-16	98.909	105.165	1.107	5.073
Qwen3-0.6B-Base	sweep	random/anchor-256/delta-16	123.480	102.753	0.857	9.025
Qwen3-0.6B-Base	sweep	plausible/anchor-256/delta-64	98.909	100.711	1.012	1.405
Qwen3-0.6B-Base	sweep	random/anchor-256/delta-64	123.480	99.870	0.785	3.891
Qwen3-0.6B-Base	sweep	plausible/anchor-512/delta-1	74.557	100.816	1.340	62.420
Qwen3-0.6B-Base	sweep	random/anchor-512/delta-1	130.917	100.488	0.786	62.649
Qwen3-0.6B-Base	sweep	plausible/anchor-512/delta-4	74.557	103.771	1.376	11.004
Qwen3-0.6B-Base	sweep	random/anchor-512/delta-4	130.917	102.951	0.797	14.843
Qwen3-0.6B-Base	sweep	plausible/anchor-512/delta-16	74.557	99.690	1.276	4.200
Qwen3-0.6B-Base	sweep	random/anchor-512/delta-16	130.917	100.756	0.788	11.170
Qwen3-0.6B-Base	sweep	plausible/anchor-512/delta-64	74.557	99.368	1.157	1.060
Qwen3-0.6B-Base	sweep	random/anchor-512/delta-64	130.917	100.166	0.739	4.290
Qwen3-4B-Base	forced	plausible/anchor-256	96.475	1.038	0.013	–
Qwen3-4B-Base	forced	random/anchor-256	176.218	5.603	0.030	–
Qwen3-4B-Base	forced	plausible/anchor-512	101.738	0.775	0.009	–
Qwen3-4B-Base	forced	random/anchor-512	169.391	5.724	0.032	–
Qwen3-4B-Base	forced	plausible/uniform-0	108.024	1.093	0.010	–
Qwen3-4B-Base	forced	random/uniform-0	158.835	5.983	0.039	–
Qwen3-4B-Base	forced	plausible/uniform-1	95.944	1.072	0.012	–
Qwen3-4B-Base	forced	random/uniform-1	167.075	5.230	0.033	–
Qwen3-4B-Base	sampled	plausible/anchor-256	96.475	125.156	1.253	–
Qwen3-4B-Base	sampled	random/anchor-256	176.218	123.916	0.733	–
Qwen3-4B-Base	sampled	plausible/anchor-512	101.738	120.416	1.119	–
Qwen3-4B-Base	sampled	random/anchor-512	169.391	123.453	0.792	–
Qwen3-4B-Base	sampled	plausible/uniform-0	108.024	123.354	1.191	–
Qwen3-4B-Base	sampled	random/uniform-0	158.835	121.878	0.795	–
Qwen3-4B-Base	sampled	plausible/uniform-1	95.944	124.548	1.353	–
Qwen3-4B-Base	sampled	random/uniform-1	167.075	123.019	0.690	–
Qwen3-4B-Base	sweep	plausible/anchor-256/delta-1	94.501	124.215	1.285	75.450
Qwen3-4B-Base	sweep	random/anchor-256/delta-1	184.183	126.275	0.711	79.959
Qwen3-4B-Base	sweep	plausible/anchor-256/delta-4	94.501	121.057	1.366	16.438
Qwen3-4B-Base	sweep	random/anchor-256/delta-4	184.869	124.754	0.704	36.843
Qwen3-4B-Base	sweep	plausible/anchor-256/delta-16	94.501	120.228	1.228	4.955
Qwen3-4B-Base	sweep	random/anchor-256/delta-16	184.869	121.436	0.687	16.900
Qwen3-4B-Base	sweep	plausible/anchor-256/delta-64	94.501	120.388	1.190	2.298
Qwen3-4B-Base	sweep	random/anchor-256/delta-64	184.869	121.080	0.672	10.062
Qwen3-4B-Base	sweep	plausible/anchor-512/delta-1	95.311	123.098	1.252	60.242
Qwen3-4B-Base	sweep	random/anchor-512/delta-1	158.059	119.039	0.712	72.529
Qwen3-4B-Base	sweep	plausible/anchor-512/delta-4	95.311	116.827	1.193	12.381
Qwen3-4B-Base	sweep	random/anchor-512/delta-4	158.059	118.687	0.725	23.036
Qwen3-4B-Base	sweep	plausible/anchor-512/delta-16	95.311	118.855	1.226	4.162
Qwen3-4B-Base	sweep	random/anchor-512/delta-16	158.059	121.398	0.753	17.577
Qwen3-4B-Base	sweep	plausible/anchor-512/delta-64	95.311	121.193	1.228	1.315
Qwen3-4B-Base	sweep	random/anchor-512/delta-64	158.059	119.332	0.682	8.149
Qwen3-8B-Base	forced	plausible/anchor-256	81.853	1.015	0.012	–
Qwen3-8B-Base	forced	random/anchor-256	163.554	12.347	0.062	–
Qwen3-8B-Base	forced	plausible/anchor-512	91.041	1.018	0.014	–
Qwen3-8B-Base	forced	random/anchor-512	172.230	10.243	0.058	–
Qwen3-8B-Base	forced	plausible/uniform-0	72.846	0.812	0.013	–
Qwen3-8B-Base	forced	random/uniform-0	154.453	9.787	0.054	–
Qwen3-8B-Base	forced	plausible/uniform-1	85.400	0.945	0.013	–
Qwen3-8B-Base	forced	random/uniform-1	152.905	9.118	0.062	–
Qwen3-8B-Base	sampled	plausible/anchor-256	81.356	101.150	1.336	–
Qwen3-8B-Base	sampled	random/anchor-256	166.135	104.739	0.656	–

Table 12 continued

Model	Run	Group	Immediate	Late	Decay	Release
Qwen3-8B-Base	sampled	plausible/anchor-512	91.041	101.186	1.220	–
Qwen3-8B-Base	sampled	random/anchor-512	172.230	102.965	0.713	–
Qwen3-8B-Base	sampled	plausible/uniform-0	72.846	100.022	1.373	–
Qwen3-8B-Base	sampled	random/uniform-0	154.453	105.964	0.738	–
Qwen3-8B-Base	sampled	plausible/uniform-1	85.400	102.772	1.217	–
Qwen3-8B-Base	sampled	random/uniform-1	152.905	107.976	0.767	–
Qwen3-8B-Base	sweep	plausible/anchor-256/delta-1	88.580	105.392	1.239	62.077
Qwen3-8B-Base	sweep	random/anchor-256/delta-1	170.609	107.152	0.675	100.680
Qwen3-8B-Base	sweep	plausible/anchor-256/delta-4	88.580	103.514	1.155	17.583
Qwen3-8B-Base	sweep	random/anchor-256/delta-4	170.609	104.585	0.686	38.539
Qwen3-8B-Base	sweep	plausible/anchor-256/delta-16	88.580	99.115	1.126	3.291
Qwen3-8B-Base	sweep	random/anchor-256/delta-16	170.609	106.035	0.654	26.097
Qwen3-8B-Base	sweep	plausible/anchor-256/delta-64	88.580	100.659	1.018	2.043
Qwen3-8B-Base	sweep	random/anchor-256/delta-64	170.609	101.880	0.665	12.079
Qwen3-8B-Base	sweep	plausible/anchor-512/delta-1	85.510	101.249	1.164	52.765
Qwen3-8B-Base	sweep	random/anchor-512/delta-1	164.171	103.507	0.684	83.105
Qwen3-8B-Base	sweep	plausible/anchor-512/delta-4	85.510	103.470	1.162	14.022
Qwen3-8B-Base	sweep	random/anchor-512/delta-4	164.171	104.899	0.636	52.448
Qwen3-8B-Base	sweep	plausible/anchor-512/delta-16	85.510	100.450	1.230	3.795
Qwen3-8B-Base	sweep	random/anchor-512/delta-16	164.171	106.694	0.666	34.870
Qwen3-8B-Base	sweep	plausible/anchor-512/delta-64	85.510	100.117	1.051	1.988
Qwen3-8B-Base	sweep	random/anchor-512/delta-64	164.171	103.182	0.624	15.097