Mapping the Blind Spots: Cultural Asymmetries in LLM Geographic Representations

Anonymous ACL submission

Abstract

Large Language Models (LLMs) encode rich world knowledge, yet how they represent geographic space internally remains poorly understood. We move beyond factual recall of location names and ask whether LLMs can perform geospatial reasoning, specifically, navigating between pairs of locations. We find that `Llama-3.1-70B` and `Qwen-2.5-72B` exhibit strong navigational ability in both localization and directionality, but this ability is unevenly distributed across regions. Further, the relationship between internal representations and geodesic distance across cultural regions reveals a consistent asymmetry where European regions are encoded with up to twice the geometric fidelity of African-Islamic and West & South Asian ones. This distortion is not limited to geometry; it also propagates to decoding confidence and downstream navigation accuracy. To test whether these representations are causally implicated in generation, we apply norm-matched activation addition and show that a single directional intervention at late residual stream layers can steer model outputs toward a target cultural region. Our findings indicate that LLMs learn a structured representation of physical space, but one that is culturally stratified. Reproducible code and dataset are released for review at https://anonymous.4open.science/r/geo-cultural-asymmetric-llm-DD60.

1 Introduction

LLMs are often called "world models." That is, they encode structural knowledge about the physical world in their parameters (Ying et al., 2025). Notably, recent work has shown that geographic structuralism, through latitudes and longitudes, can be recovered from a model's activation space (Gurnee and Tegmark, 2024), as if the models implicitly embed a map of the world in their representations. However, the ability to *recover* this map does not necessarily guarantee that the model encodes various cultural regions with equal fidelity, or that its outputs are accurate. In this paper, we ask: Do LLMs encode geodesic structure uniformly across regions, or do they maintain a distorted "mental map," in which some areas are represented with much higher faithfulness than others? And, if so, is that map causally implicated in generation?

Our work is motivated by recent findings of *behavioral* distortion at the output-level. Zhao et al. (2026) show that modern LLMs perceive geographical locations through a *culturally uneven baseline*. That is, under neutral prompts, an LLM's outputs *lie* systematically closer to European and North American reference frames than to non-Western ones. Further, Manvi et al. (2024) identify output-level disparities for geographic predictions, and Dunn et al. (2024) finds that pretrained language models represent some geographic populations much better than others, with downstream performance varying by region. While prior works identify cultural asymmetry, their analysis is limited to output-level divergence.

In this paper, we look beyond these textual behavioral disparities and ask whether the model's internal representations are themselves geometrically distorted, and whether such distortions causally drive the behaviors we observe.

We investigate two open-weight models, `Llama-3.1-70B` (Grattafiori et al., 2024) and `Qwen-2.5-72B` (Yang et al., 2024); and, we focus on the eight cultural regions defined in the World Values Survey (Inglehart and Welzel, 2005; Haerpfer et al., 2020). We perform three complementary analyses:

Representation (§2.1): We measure how accurately geographic structure is encoded within each region via representational similarity analysis between internal representation and geographic-distance matrices (Kriegeskorte, 2008).

Figure 1: **Geographic fidelity in LLM-generated paths is asymmetric across cultural regions.** (*left*) The quality of a generated path from city X to city Y correlates with the cultural region: paths within well-represented regions (Madrid→Warsaw) stay coherent and near-direct, while paths within poorly represented regions are chaotic (Casablanca→Cairo) and inefficient (Mumbai→Bangkok). (*right*) We provide causal evidence that LLMs encode cultural regions as directions in their activations. These representations can be used to steer generation: adding the Protestant-Europe (PEU) direction reroutes a Madrid→Warsaw query through the Protestant-Europe cluster (olive region), with additional unnecessary waypoints inside the PEU cluster, whereas for the Casablanca→Cairo route, it chooses Rome to fit in a PEU city.

Behavior (§2.2): We define a country-level navigation task, where we generate a sequence of country waypoints between two countries and measure per-region path quality and decoding confidence.

Causal role (§2.3): We apply a norm-matched additive intervention along a target region's direction (Turner et al., 2023); this selectively re-routes generated paths into the steered region, providing causal evidence that the model uses these regional representations during generation.

2 Methodology

Preliminaries. Large language models (LLMs) are functions $f_{\mathcal{M}}$ that map input prompts to a next-token distribution. To do so, input passes through a stack of intermediate computations, exposing residual-stream activations $h^{(\ell)} \in \mathbb{R}^d$ at every layer $\ell \in \{1, \ldots, L\}$. In our case, we are interested in $f_{\mathcal{M}}$'s geographical representations and its performance on a navigation task across different cultural regions. Let $\mathcal{C} = \{c_1, \ldots, c_N\}$ be a set of N geographical locations, where c_i has geographic coordinates (φ_i for latitude and λ_i for longitude). We define $\mathcal{C}_r \subset \mathcal{C}$ as the subset of locations belonging to region r. A summary of the notation used throughout the paper is provided in Appendix A.

We ask whether $f_{\mathcal{M}}$ exhibits performance discrepancies across the $\mathcal{C}_r$, examining $f_{\mathcal{M}}$ from two complementary angles: *behaviorally*, how it solves navigational tasks, and *representationally*, what it encodes internally.

Behavior. We probe $f_{\mathcal{M}}$'s geographic behavior through a **country-level navigation task**: given a start country c_s and a destination c_d, the LLM is prompted to generate a path $p = [c_s, w_1, w_2, \ldots, w_k, c_d]$, whose intermediate **waypoints** w_j are countries that connect c_s and c_d (§2.2). However, as Harding (2023) argues, behavioral evidence alone is not a reliable proxy for what an LLM "knows."

Representation. We therefore also inspect $f_{\mathcal{M}}$ from the inside, examining the residual-stream activations $h_i^{(\ell)}$ associated with a geographic entity c_i. Section 2.1 describes representational similarity analysis (RSA) to quantify how faithfully LLMs encode geographic structure.

From representation to behavior. Internally encoding a concept is not evidence that the LLM *uses* it when generating outputs: a feature can be linearly decodable yet causally irrelevant. Budding (2025) argues that attributing a concept to an LLM requires demonstrating a direct causal effect of its representation on behavior, established through targeted interventions rather than decoding alone. We therefore complement RSA and the navigation task with a **causal intervention** (§2.3).

2.1 Representational Similarity Analysis

Latitude and longitude are an arbitrary parameterization of the Earth's surface. Thus, there is no reason to expect a language model to have internalized this notion. Whereas linear probing (Gurnee and Tegmark, 2024) tests whether absolute geographic

coordinates (latitude and longitude) are decodable, RSA (Kriegeskorte, 2008) quantifies whether the relative geometry of internal representations matches that of geographic space. That is, we define the geographic distance matrix $D_g \in \mathbb{R}^{N\times N}$ via the haversine distance. Further, for each layer ℓ, we record the hidden representation $\mathbf{h}_i^{(\ell)}$ for a city c_i and form the activation dissimilarity matrix $D_a^{(\ell)} \in \mathbb{R}^{N\times N}$ using pairwise cosine dissimilarity (1 - cosine similarity).

Global alignment. Global RSA score ($\rho^{(\ell)}$), or the Spearman rank correlation between the vectorized upper triangles of D_g and $D_a^{(\ell)}$, can faithfully locations are embedded into internal geometry. Since Spearman correlation is invariant to monotone rescaling, $\rho^{(\ell)}$ does not depend on the scale of either matrix. That is, $\rho^{(\ell)} \approx 1$ indicates perfectly aligned pairwise rankings between D_g and $D_a^{(\ell)}$, while 0 indicates no monotone relationship. Comparing $\rho^{(\ell)}$ across the regions $\mathcal{C}_r$ can expose global asymmetries in representational quality.

Entity alignment. To assess representational quality of each location c_i, we can look at the Spearman correlation $\rho_i^{(\ell)}$ between the i-th rows of $D_a^{(\ell)}$ and D_g (local RSA score). A high $\rho_i^{(\ell)}$ indicates c_i's relative distances to all other cities are well-preserved in activation space and a low value indicates poor representation. In Appendix C, we provide additional analysis of the relative geometries.

2.2 Behavior and prompting

Navigation tasks are a common proxy for the geospatial knowledge in LLMs (Truong et al., 2025; Vafa et al., 2024). Given a start c_s and a destination c_d, the LLM is prompted to generate a path p of adjacent *waypoints* that connect the two points (see Appendix G).

If the model has a good representation of the relevant region, the emitted waypoints should closely follow the great-circle[1] *corridor*[2] between c_s and c_d with minimal deviation from it.

Evaluation metric. For each path p, we obtain four behavioral metrics m: mean per-token log-probability on the path (*confidence*), mean per-token entropy (*entropy*), and two waypoint-bearing variants (see formal definition in §3.3). We control for path difficulty via the great-circle distance $d_{\mathrm{SD}}(p)$ between corridor and emitted waypoints. The path-level partial Spearman's correlation is

$$\rho\left(\mathrm{RSA}\left(p\right),\, m\left(p\right) \mid d_{\mathrm{SD}}\left(p\right)\right). \tag{1}$$

In Equation (1), each path p inherits an RSA score based on its start c_s or destination c_d; a cross-region path therefore contributes to both regions.

2.3 Intervention Analysis

We perform steering experiment, to determine whether an LLM has an encoded concept for a given region r, and whether this concept affects the behavior.

Direction for a region. For a target region r, we define the unit direction most representative of the region r at layer ℓ as [3] :

$$\hat{\boldsymbol{v}}_r = \frac{\boldsymbol{\mu}_{\mathrm{in}}^r - \boldsymbol{\mu}_{\mathrm{out}}^r}{\left\|\boldsymbol{\mu}_{\mathrm{in}}^r - \boldsymbol{\mu}_{\mathrm{out}}^r\right\|}, \tag{2}$$

where $\boldsymbol{\mu}_{\mathrm{in}}^r$ and $\boldsymbol{\mu}_{\mathrm{out}}^r$ are the mean residual stream activations for the locations that are in the region and those not in the region, respectively. This is computed over the $2,400$ cities in our dataset (see §3.2 and Appendix K).

We generate the path using a stepwise multi-turn process. At each step t, the model receives the start location c_s, and the previously generated waypoints $w_1, w_2, ..., w_{t-1}$. An LLM predicts the next waypoint w_t via

$$\arg\max_{w}\; f_{\mathcal{M}}\big(w \mid c_s, c_d, w_1, \ldots, w_{t-1}\big). \tag{3}$$

Norm-matched additive intervention. Following the contrastive activation addition family (Turner et al., 2023; Rimsky et al., 2024), at each generation step we perturb the residual stream at the last token of layer ℓ by:

$$\mathbf{h}_{\mathrm{last}}^{(\ell)} \leftarrow \mathbf{h}_{\mathrm{last}}^{(\ell)} + \alpha \cdot \hat{\boldsymbol{v}}_r \cdot ||\mathbf{h}_{\mathrm{last}}^{(\ell)}||, \tag{4}$$

where $\alpha \in \mathbb{R}$ controls the magnitude and sign of the intervention (Wurgaft et al., 2026).

We use this intervention to test whether $\hat{\boldsymbol{v}}_r$ is causally relevant to the LLM's geographic outputs. If the regional vector $\hat{\boldsymbol{v}}_r$ encodes a concept

[1] Throughout, *great-circle* denotes the haversine-shortest arc between two points on the unit sphere (R=6371 km). All bearing and distance quantities are computed on this geodesic.

[2] A *corridor* is the set of countries whose centroid lies within a tube around the $c_s \to c_d$ geodesic: along-track $\in [0, d_{SD}]$, cross-track $\leq \max(300\,\mathrm{km},\ 0.15 \cdot d_{SD})$.

[3] We omit the layer index ℓ from $\hat{\boldsymbol{v}}_r$, $\boldsymbol{\mu}_{\mathrm{in}}^r$, and $\boldsymbol{\mu}_{\mathrm{out}}^r$ for readability; all three quantities are layer-dependent.

the LLM actively uses, then a positive α should bias generation toward region r (higher first-token mass on in-region waypoints and paths that divert through the region) while a negative α should bias generation away from it. A null result on either side would suggest that $\hat{\boldsymbol{v}}_r$, while linearly decodable from the residual stream, does not causally connect to the behavior.

Specificity. A causal effect of $\hat{\boldsymbol{v}}_r$ is only meaningful if it is *specific* to a region r. That is, steering with $\hat{\boldsymbol{v}}_r$ should shift target-region more than steering with a direction associated with another region r' would. Formally,

$$\hat{m}(\hat{\boldsymbol{v}}_r, \alpha) = \frac{1}{|p|} \sum_{w \in p} f_{M,r}(w \mid \hat{\boldsymbol{v}}_r, \alpha) \tag{5}$$

In Equation (5), $\hat{m}(\cdot, \cdot)$ is a per-path average of the first-token probability mass on region-r countries when the residual is steered by direction $\hat{\boldsymbol{v}}_r$ at scale α. The per-direction increase over the unperturbed baseline is

$$\Delta_{\text{target}}(\hat{\boldsymbol{v}}_r, \alpha) = \hat{m}(\hat{\boldsymbol{v}}_r, \alpha) - \hat{m}(\hat{\boldsymbol{v}}_r, 0) \tag{6}$$

Therefore, the *differential-increase* specificity contrast against a wrong-region control direction $\hat{\boldsymbol{v}}_{r'}$ is

$$\Delta_{\text{spec}}(\alpha; r, r') = \Delta_{\text{target}}(\hat{\boldsymbol{v}}_r, \alpha) - \Delta_{\text{target}}(\hat{\boldsymbol{v}}_{r'}, \alpha). \tag{7}$$

A positive Δ_{spec} indicates that $\hat{\boldsymbol{v}}_r$ shifts target-region mass beyond what the matched-norm control; values near zero indicate that $\hat{\boldsymbol{v}}_r$ is not isolating region r from r'. We report confidence intervals via path-paired bootstrap with 10,000 resamples.

3 Experimental Setup

3.1 Data

We look at 8 cultural regions r defined by Inglehart and Welzel (2005) using the World Value Survey Haerpfer et al. (2020). The full 8 regions are used for the representation analyses (§2.1); for layout consistency with the 7-region navigation bank, main-text per-region figures restrict to the 7 regions traversed by paths and Latin America is reported in Appendix C.

Navigational benchmark. We use 225 country-level paths. Each path p is a sequence of five neighboring countries connecting a start country c_s to a destination country c_d along a roughly great-circle adjacency corridor. For instance, given Slovenia as c_s and Albania as c_d, the path is *Slovenia* $\rightarrow$ *Croatia* $\rightarrow$ *Montenegro* $\rightarrow$ *Serbia* $\rightarrow$ *Albania*. The three countries in the middle are waypoints w_j. The composition of the dataset is even across European (75), African (75), and Asian (75) start–destination pairs.[4]

At inference time, we prompt an LLM with a fixed one-shot template with the example "India $\rightarrow$ Nepal $\rightarrow$ Bhutan $\rightarrow$ China" and the schema `[Start]` $\rightarrow$ `[Waypoint]` $\rightarrow$ . . . $\rightarrow$ `[Destination]`. We collect the waypoints using greedy decoding. After each turn, we parse the next waypoint from the emitted text, append it to the running prompt, and re-prompt; the path terminates when the model emits the destination, an empty waypoint, or reaches a length cap.

We record, per step: the emitted token, its per-token log-probability and entropy, and the country resolved from the emitted text. The navigational benchmark is held fixed across two LLMs the verbatim prompt template and termination cap are in Appendix G.

RSA. For each cultural region r, we draw a fixed set $\mathcal{C}_r$ of 300 cities (8 regions, 2,400 cities total). City lists are drawn from the GeoNames database via `GeonamesCache`; selection rule, licensing, and full city listings are in Appendix E and Appendix K.

3.2 Models and Activations

We study open-source LLMs spanning two model families: `Llama-3.1-70B` (Grattafiori et al., 2024) and `Qwen-2.5-72B` (Yang et al., 2024). Further, we refer to these by the family names only. We report two base models at the homologous late residual layer $L_{\text{final}}-4$: `Llama` at L_{76} and `Qwen` at L_{76}. The layer choice $L_{\text{final}}-4$ is motivated by prior work (Wurgaft et al., 2026) that places concept-relevant directions near but not at the unembedding layer.

Activations. We probe each model with the prompt "`Where is {city}?`"and read the residual-stream activation at the final token of the prompt, $\mathbf{h}_i^{(\ell)} \in \mathbb{R}^d$, for each $c_i \in C_r$ at every layer ℓ,

[4] Latin America (South + Central America) and Oceania are excluded from the path bank because countries in these regions have too few land-adjacent neighbors to construct multi-hop great-circle corridors comparable to those in the 225-path bank (e.g., island states in Oceania have zero land neighbors; Latin American countries form a thin land chain).

which is used to compute RSA and $\hat{\boldsymbol{v}}_r$. We extract activations using the `nnsight` library and the NDIF remote-tracing service (Fiotto-Kaufman et al., 2025).

3.3 Metric

For each region r and each path that passes through it we report PRECISION_r (fraction of emitted countries in r that lie on the great-circle corridor), RECALL_r (fraction of reference countries in r that the model emits), region-presence recall $\text{REC}_r^{\text{pres}}$ (the model emits at least one country in r when the reference does), emit-confidence$_r$ (mean log-probability on emitted countries in r), and nav-pass$_{45,r}$, the strict progress-guarded fraction of intermediate hops in r within $45°$ of the great-circle direction: each counted hop must (i) lie at least 25 km from the start, (ii) advance the along-track distance, (iii) not overshoot the destination, and (iv) reduce the remaining distance to destination. *nav-pass*$^{\text{all}}_{45,r}$ uses the same metric with unresolved emissions counted as failures.

4 Results

We report four findings that progress from behavioral observation to mechanistic interpretation. We begin by establishing the behavioral asymmetry: path-completion quality and decoding confidence differ systematically across WVS cultural regions (§4.1). Second, this asymmetry has a representational counterpart: residual-stream activations preserve great-circle distances less faithfully for non-Western regions (§4.2). Third, representational fidelity predicts generation confidence at the path level, even after controlling for geodesic difficulty (§4.3). Finally, the link is causally addressable: perturbing the residual stream along the contrastive Protestant-Europe direction shifts decoding behavior toward that region (§4.4).

4.1 Behavioral asymmetry

Table 1 and Table 6 show a consistent pattern where regions like Protestant, Orthodox, and Catholic Europe occupy strong precision and recall space with the least-negative decoding confidence, while African-Islamic and West & South Asia trail across the same lag on both axes. While the English-speaking region is an outlier, it shows how Path-completion quality and decoding confidence differ systematically across cultural regions for both models.

Additionally, we see that African-Islamic has the highest region-presence recall on the path bank, where the model can reliably detect when the path passes through Africa. This differentiation in navigational ability motivates the reasoning that models' abilities are culturally distorted, thus we analyze internally.

4.2 Region representational asymmetry

The behavioral asymmetry documented above has a representational counterpart. Gurnee and Tegmark (2024) show that absolute geographic coordinates are linearly decodable from LLM activations; we replicate this finding on both models, obtaining joint lat/lon $R^2 = 0.831$ for `Llama` (best probing layer L_{62}) and $R^2 = 0.759$ for `Qwen` (best probing layer L_{78}). Crucially, this decoding accuracy is not uniform across regions (see Appendix B). However, absolute position captures only one aspect of geometry, whereas navigation depends on pairwise relative distance. Whether two cities can plausibly be successive waypoints is a relational property of their activations, and not an intrinsic property of either city in isolation. Thus, we measure the pair-wise representational similarity (RSA) and test whether the European-vs.-non-Western ordering from §4.1 persists when we evaluate how faithfully residual-stream activations preserve great-circle distance within each region.

As seen in Fig. 2, Protestant, Catholic, and Orthodox Europe rank highest on the RSA spider plot while African-Islamic and West & South Asia rank lowest. We observe the same European vs. non-Western ordering that drove navigation behavior in §4.1 re-emerge in relative-distance fidelity.

4.3 Linking representation to behavior

Given that the model has established a regional ordering, we investigate whether this ordering also appears when the model is tasked with navigation. To test this directly, we conduct a path-level analysis on the partial Spearman $\rho(\text{RSA}(p), m(p) \mid d_{\text{SD}}(p))$ (Eq. 1). Since cross-region paths are attributed to both endpoints, our path dataset expands to N=294 from 225.

Table 2 reports the results where we see that the higher RSA correlates with higher confidence and navigation quality, lower entropy, and with seven of eight entries reaching $p < 0.05$ on `Llama` and the weakest `Qwen` entry at $p = .052$. On applying the Holm adjustment to account for the multiple comparisons we make, we see that only the

Region	n_{paths}	PRECISION$_r$	RECALL$_r$	REC$_r^{\text{pres}}$	emit-conf.$_r$	nav-pass$_{45,r}$
Protestant Europe	38	0.86	0.86	0.66	-0.57	0.72
Catholic Europe	63	0.94	0.89	0.76	-0.63	0.69
Orthodox Europe	52	0.92	0.92	0.73	-0.70	0.60
English-Speaking	9	**1.00**	0.67	0.22	**-0.43**	**1.00**
Confucian	28	0.75	**1.00**	0.43	-0.79	0.60
African-Islamic	**135**	0.90	0.89	**0.94**	-0.71	0.64
West & South Asia	46	0.71	**1.00**	0.74	-0.72	0.55

Table 1: **Navigational Metrics across eight WVS regions for `Llama`.** We see that per-region navigation quality and decoding confidence diverges sharply across the eight WVS regions on the country-level navigation task. n_{paths} is the count of paths whose path passes through that WVS region. The `Qwen` and small models' behavioral results are reported in Appendix D.

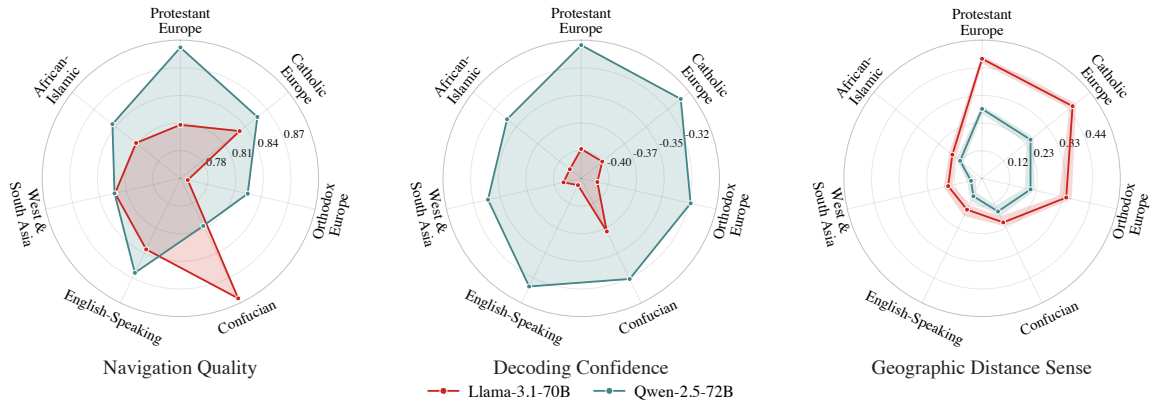

Figure 2: **Navigation Behavior and Representational Fidelity Across Cultural Regions.** We plot three radar plots showing per-region navigation quality (nav-pass$_{45,r}$, left), decoding confidence (middle), and geographic RSA (right) for `Llama` and `Qwen`. We observe a consistent ordering across all three metrics: European regions lead while African-Islamic and West & South Asia trail.

Metric	Llama-3.1-70B L76	Qwen-2.5-72B L76
log-prob	**+0.247 (0.009)**	**+0.214 (0.047)**
entropy	**-0.233 (0.001)**	**-0.156 (0.042)**
nav-pass$_{45}$	**+0.139 (0.022)**	**+0.066 (0.043)**
nav-pass$_{45}^{all}$	**+0.177 (0.039)**	+0.091 (0.052)

Table 2: **Path-level Correlation between RSA and Behavioral Metrics.** We compute the partial Spearman rho between per-path RSA and behavioral metrics, controlling for path difficulty d_{SD}, for all paths (n=294). We observe that representational fidelity consistently and significantly predicts decoding confidence and waypoint accuracy across both models

confidence and entropy metrics remain statistically significant, while navigational metrics lose significance. This shows that navigating through well-represented regions inflates the model's confidence without checking for navigational accuracy.

4.4 Causal steering

As we have seen, models possess asymmetry not only at the behavioral level but also at the representational level. In this section, we now analyze whether the model's representationally distorted internal maps are causally responsible for their asymmetry in navigational behavior, too. We select Protestant Europe as the steering target because it exhibits the highest representational fidelity across both models (Fig. 2, right), making it the direction most likely to yield a clean causal signal and the strongest test of regional specificity against nearby European controls. We derive the contrastive Protestant-EU direction (§2.3) and apply this intervention in a later layer (ℓ=ℓ_{final}−4), following previous work (Rimsky et al., 2024; Ali et al., 2025; Lawson et al., 2025).

We vary the steering strength α across multiples of $\{-2, -1, -0.5, 0, +0.5, +1, +2\}$. We measure five behavioral metrics: **Target Mass** measures the probability the model assigns to the target region at the first token, before emitting a waypoint, capturing implicit regional bias before generation. **Target Share** measures how often the emitted waypoint itself falls within the target region, capturing the effect on the final output. **Confidence** measures the model's certainty over its emissions, com-

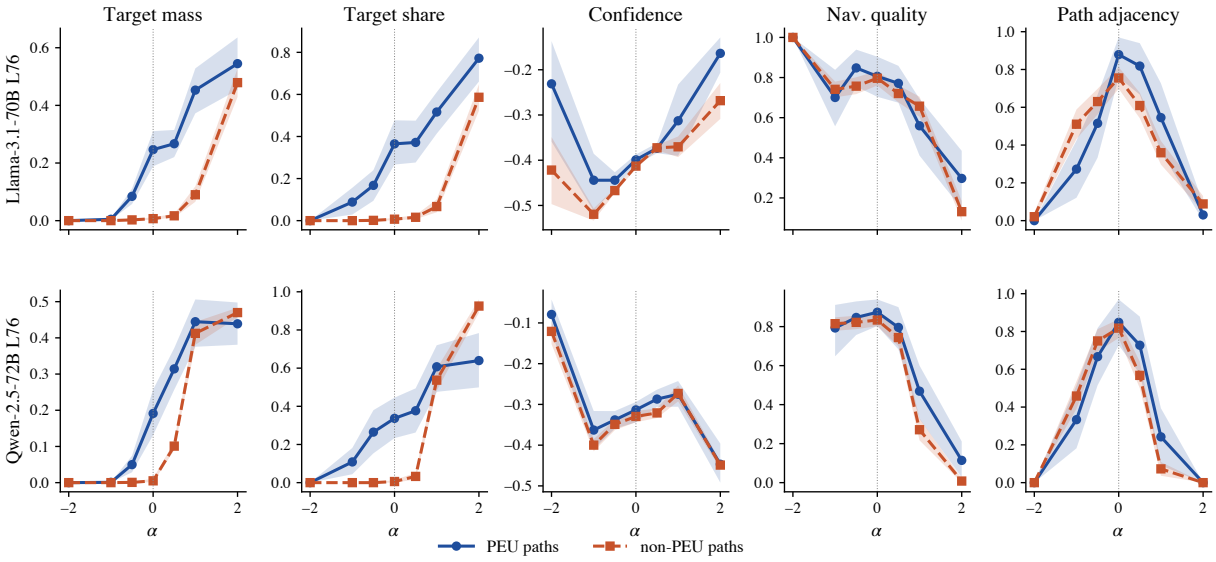

Figure 3: **Steering Curves For Protestant Europe (PEU) and Non-Protestant Regions Europe (non-PEU) paths.** We intervene and steer the residual stream at layer $\ell = 76$ and steer navigation generation at every waypoint toward Protestant Europe across $\alpha \in \{-2, -1, -0.5, 0, +0.5, +1, +2\}$, evaluating on both PEU and non-PEU paths for `Llama` and `Qwen`. We plot five behavioral metrics: Target mass, Target share, Confidence, Navigational quality (nav-pass$_{45,r}$), and Path adjacency. We observe that target mass and share increase selectively on PEU paths, while path validity degrades at large $|\alpha|$ across both subsets.

			Catholic ctrl		English ctrl	
Cell	α	Δ_{target}	cos	Δ_{spec}	cos	Δ_{spec}
`Llama`	$+1$	$+0.101$ [$+.077, +.127$]	$+0.72$	-0.005 [$-.037, +.028$] ×	-0.16	$+0.069$ [$+.042, +.098$] ✓
`Qwen`	$+0.5$	$+0.100$ [$+.087, +.113$]	$+0.37$	$+0.050$ [$+.036, +.064$] ✓	$+0.03$	$+0.090$ [$+.078, +.104$] ✓

Table 3: **Steering Specificity Against Near and Far Control Directions.** We report Δ_{target} and Δ_{spec} against two control directions: Catholic Europe (near-neighbor) and English-Speaking (distant-neighbor), at the comparable α for `Llama` & `Qwen`. cos is the cosine similarity between the target and control directions in model's activation space.

puted as mean per-token log-probability. **Navigation Quality** (nav-pass$_{45,r}$) measures directional consistency, specifically the fraction of hops whose bearing falls within $45°$ of the great-circle corridor between source and destination. Finally, **Path Adjacency** measures geographic validity, whether each emitted waypoint is physically land-adjacent to the previous one.

As shown in Fig. 3, increasing α monotonically increases the target probability mass and the fraction of the emitted countries, confirming that the steering vector successfully biases the model to align with or away from the target region. In contrast, the stricter path-adjacency metric traces an inverted U: moderate steering strengths $[-0.5 \leq \alpha \leq +0.5]$ improve the model's ability to output adjacent paths, while stronger steering begins to degrade it. This suggests that overly aggressive steering pushes the model toward target-region outputs at the expense of maintaining realistic or internally consistent transition structure. The confidence curve corroborates this: at moderate α, confidence is lower but stable, whereas at large $|\alpha|$ it becomes highly variable. We hypothesize that this increased confidence could be a function of model hallucination, which is confirmed via manual inspection.

Is the Steering Effect Consistent Across Model Families? Additionally, we compare the steering capabilities across both models. To enable a fair comparison, we devise a validity-matched setting (Appendix I), where we choose the comparable α for each model where the path-adjacency score remains above 0.24. This corresponds to $\alpha = +1$ for `Llama` and $\alpha = +0.5$ for `Qwen`. In this setting, as seen in Table 3, both models show similar effects ($\Delta_{\text{target}} \approx +0.10$), indicating that the results follow a consistent pattern across models. We see that across metrics, stronger steering increases target-

region mass but reduces confidence and navigation quality, showing a clear tradeoff between effect strength and output reliability (Fig. 3).

Does Steering Generalize Across Representational Distance? Having established that steering interventions bias navigational decisions toward a target region (Protestant Europe), we now evaluate whether this effect is selective to nearby directions or whether it also extends to more distant control directions.

A nearby control direction is defined as one with high cosine similarity to the target region, while a distant control direction has low cosine similarity. As seen in Table 3, the Catholic-Europe direction is relatively close to the target in `Llama` ($\cos = 0.72$) but less so in `Qwen` ($\cos = 0.37$), while English-Speaking is distant from the target in both models ($\cos = -0.16$ on `Llama`, $\cos = 0.03$ on `Qwen`).

On comparing the Δ_{spec} on the distant control group, we see that it is positive in both models. In the nearby control group, this effect weakens substantially: it collapses in `Llama` and is reduced in `Qwen`. This pattern suggests that the Protestant and Catholic directions are not fully separable in `Llama`'s activation space, but instead share a broader European–Christian component.

We hypothesize that this behavior is a consequence of the model's representational geometry. As seen in Fig. 2, for Protestant, Catholic, and Orthodox European regions, the model's internal activations correspond very closely to the physical geodesic metrics. Because these regions are close together, their direction vectors share a large common component. As a result, steering along either the Protestant or Catholic direction mostly moves the representation along this shared "European" axis, rather than cleanly separating the regions.

In smaller models, we observe similar patterns where "European-Christian" regions are less separable and their control effect is less profound. However, in `Qwen`, the Catholic and Protestant directions are less collinear, making them more separable in representation space and reducing interference between them.

Testing against random vector baseline. To further validate our results, we test whether the observed effect can be produced by chance using 10 random directions with the same magnitude as the steering vector. Across all tested values of α, none of the random directions reproduce the `Llama` effect on Δ_{target}. At $\alpha = +1$, the average effect is near zero (mean 0.005 with a small confidence interval). This suggests the result is specific to the learned direction rather than a random vector.

5 Conclusion

Our results demonstrate a consistent picture where we see that LLMs encode a distorted internal map. We investigate this cultural asymmetry across three axes: *navigational behavior*, *representational geometry*, and *causal intervention*. This individually independent analysis exhibits that bias is a structural property of how models internalize geographic knowledge. For example, regions such as Protestant, Orthodox, and Catholic Europe are not just represented better, but this representational advantage shapes how confidently and accurately the model navigates within them.

Furthermore, our work moves beyond the output-level behavioral understanding of model geography to a more mechanistic understanding of how models internally represent spatialness. We extend prior work documenting output-level disparities across cultural regions and show that these disparities correlate with representational distortion of internal maps, which is causally active when the model performs navigational generation. Interestingly, we see that a single additive direction, which is derived purely from a city-based probe and has no access to the navigational task, when intervened in a late residual stream, can reliably steer the model's path. This shows that the cultural-regional structure is encoded as a coherent geometric direction in activation space.

Crucially, our findings also reveal that the model's confidence is more likely associated with its representational fidelity rather than navigational accuracy. For example, a model may generate waypoints in a well-represented region, such as Protestant Europe, with high confidence, even when those waypoints are navigationally wrong, leaving the users with no surface signal to detect a confident wrong answer from a correct one.

Limitations. While exhaustive, our analysis does face certain limitations. First, we focus on only two model families across sizes ranging from 7B to 72B. Secondly, we stratify the world into eight WVS cultural regions derived from the Inglehart and Welzel (2005) cultural map. While this is a strong framework for organizing cultural variation, it has a few limitations, primarily that it may not capture all the innate intricacies and diversity of

cultural identities (Norris et al., 2001). Additionally, there are multiple taxonomies (Huntington and Jervis, 1997) where often culture is viewed through a multi-dimensional lens and cannot be reduced to a single framework (Hofstede, 2011). Finally, our causal evidence is targeted, where we show steering can control region-token probability mass, but this does not translate to improved geographic reasoning. Additionally, for `Llama-3.1-70B`, Protestant and Catholic European directions are geometrically similar and not cleanly separable, limiting the specificity of our intervention. Future work should explore held-out country evaluations, multilingual prompts for diversity, and finer-grained regional directions to address these limitations.

Ethics Considerations. Our work analyzes how models represent different locations stratified by a cultural map. While we observe cultural asymmetry in the models' representation, we hope to position these distortions as a diagnostic roadmap for equitable LLM development. We acknowledge that the use of steering vectors introduces directional bias that may over-correct or under-represent target regions. A potential risk of this work is that the same activation-steering methods used here for diagnosis could be repurposed to manipulate geographic or cultural associations in model outputs. We use steering solely as an analysis tool. Artifact licensing, intended-use compatibility, Personally Identifying Info (PII) screening, and full domain/coverage documentation are reported in Appendix E. We used LLM-assisted tools (Claude, ChatGPT, and Gemini) to help generate code, draft portions of the manuscript, formalize language, and correct grammar. All outputs were verified by the authors. No LLM was used to generate or alter experimental results.

References

Sheikh Abdur Raheem Ali, Justin Xu, Ivory Yang, Jasmine Xinze Li, Ayse Arslan, and Clark Benham. 2025. Scaling laws for activation steering with llama 2 models and refusal mechanisms. *arXiv preprint*.

Céline Budding. 2025. What do large language models know? tacit knowledge as a potential causal-explanatory structure. *arXiv preprint*.

Jonathan Dunn, Benjamin Adams, and Harish Tayyar Madabushi. 2024. Pre-trained language models represent some geographic populations better than others. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 12966–12976, Torino, Italia. ELRA and ICCL.

Jaden Fried Fiotto-Kaufman, Alexander Russell Loftus, Eric Todd, Jannik Brinkmann, Koyena Pal, Dmitrii Troitskii, Michael Ripa, Adam Belfki, Can Rager, Caden Juang, Aaron Mueller, Samuel Marks, Arnab Sen Sharma, Francesca Lucchetti, Nikhil Prakash, Carla E. Brodley, Arjun Guha, Jonathan Bell, Byron C Wallace, and David Bau. 2025. NNsight and NDIF: Democratizing access to foundation model internals. In *The Thirteenth International Conference on Learning Representations*.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*. Hugging Face checkpoint: meta-llama/Meta-Llama-3.1-8B and meta-llama/Llama-3.1-70B.

Wes Gurnee and Max Tegmark. 2024. Language models represent space and time. In *International Conference on Learning Representations (ICLR)*.

Christian Haerpfer, Ronald Inglehart, Alejandro Moreno, Christian Welzel, Kseniya Kizilova, Jaime Diez-Medrano, Marta Lagos, Pippa Norris, Eduard Ponarin, and Bi Puranen. 2020. World values survey wave 7 (2017-2020) cross-national data-set.

Jacqueline Harding. 2023. Operationalising representation in natural language processing. University of Chicago Press.

Geert Hofstede. 2011. Dimensionalizing cultures: The hofstede model in context. *Online readings in psychology and culture*, 2(1):8.

Samuel P Huntington and Robert Jervis. 1997. The clash of civilizations and the remaking of world order. *Finance and Development-English Edition*, 34(2):51–51.

Ronald Inglehart and Christian Welzel. 2005. *Modernization, Cultural Change, and Democracy*. Cambridge University Press.

Nikolaus Kriegeskorte. 2008. Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*.

Tim Lawson, Lucy Farnik, Conor Houghton, and Laurence Aitchison. 2025. Residual stream analysis with multi-layer SAEs. In *The Thirteenth International Conference on Learning Representations*.

Rohin Manvi, Samar Khanna, Gengchen Mai, Marshall Burke, David B. Lobell, and Stefano Ermon. 2024. GeoLLM: Extracting geospatial knowledge from large language models. In *The Twelfth International Conference on Learning Representations*.

Pippa Norris, WL Bennett, and RM Entman. 2001. Communication, society and politics.

Nina Rimsky, Nick Gabrieli, Julian Schulz, Meg Tong, Evan Hubinger, and Alexander Turner. 2024. Steering llama 2 via contrastive activation addition. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15504–15522, Bangkok, Thailand. Association for Computational Linguistics.

Thinh Hung Truong, Jey Han Lau, and Jianzhong Qi. 2025. Understanding the geospatial reasoning capabilities of llms: A trajectory recovery perspective. *arXiv preprint*.

Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Juan J. Vazquez, Ulisse Mini, and Monte MacDiarmid. 2023. Steering language models with activation engineering. *arXiv preprint*.

Keyon Vafa, Justin Y. Chen, Ashesh Rambachan, Sendhil Mullainathan, and Jon Kleinberg. 2024. Evaluating the world model implicit in a generative model. The Thirty-eighth Annual Conference on Neural Information Processing Systems. Curran Associates Inc.

Daniel Wurgaft, Can Rager, Matthew Kowal, Vasudev Shyam, Sheridan Feucht, Usha Bhalla, Tal Haklay, Eric Bigelow, Raphael Sarfati, Thomas McGrath, and 1 others. 2026. Manifold steering reveals the shared geometry of neural network representation and behavior. *arXiv preprint*.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and 23 others. 2024. Qwen2.5 technical report. *arXiv preprint*. Hugging Face checkpoint: Qwen/Qwen2.5-7B and Qwen/Qwen2.5-72B.

Lance Ying, Katherine M Collins, Prafull Sharma, Cedric Colas, Kaiya Ivy Zhao, Adrian Weller, Zenna Tavares, Phillip Isola, Samuel J Gershman, Jacob D Andreas, and 1 others. 2025. Assessing adaptive world models in machines with novel games. *arXiv preprint arXiv:2507.12821*.

Rong Zhao, Wanqi Liu, Zhizhou Sha, Nanxi Su, and Yecheng Zhang. 2026. Large language models perceive cities through a culturally uneven baseline. *arXiv preprint*.

A Notations

We provide an overview of notation used across our manuscript in Table 4.

Symbol	Description	Notes / Shapes
$\mathcal{C}$	Set of N geographic locations (e.g., cities)	$\mathcal{C} = \{c_1, \ldots, c_N\}$
r	Cultural region	
$\mathcal{C}_r$	Set of cities in the region r	
p	Path from start c_s to destination c_d, includes k waypoints	
$\|p\|$	Length of a path $p = \{c_s, w_1, \ldots, w_k, c_d\}$	
c_s	Starting geographical location	
c_d	Destination geographical locations	
w_i	Waypoints between c_s and c_d	
ℓ	Index of a layer	
	Representations Similarity Analysis	
φ_i	Latitude of a geographical location (e.g., city c_i)	$c_i \in \mathcal{C}$
λ_i	Longitude of a geographical location (e.g., city c_i)	$c_i \in \mathcal{C}$
D_g	Geographic distance matrix (based on φ_i, λ_i)	$\mathbb{R}^{N \times N}$
$\mathbf{h}_i^{(\ell)}$	Hidden representation of a location c_i at a layer ℓ	$\mathbf{h}_i^{(\ell)} \in \mathbb{R}^d$
$D_a^{(\ell)}$	Activation dissimilarity matrix (at a layer ℓ)	$\mathbb{R}^{N \times N}$
$\rho^{(\ell)}$	Global RSA score (i.e., aggregated) (at a layer ℓ)	
$\rho_i^{(\ell)}$	Local RSA score (e.g. per-city) at a layer ℓ	
	Behavioral Evaluation (§2.2)	
$m(\cdot)$	Metric function	
$d_{SD}(p)$	Great-circle distance (controls for path difficulty)	
$\rho(\cdot \mid \cdot)$	Partial Spearman Correlation	
	Additive Interventions (§2.3	
t	Step index in stepwise multi-turn generation	
w_t	Waypoint emitted by the LLM at step t	Eq. 3
$P_\theta(\cdot)$	Next-waypoint distribution under model parameters θ	
$\boldsymbol{\mu}_{\text{in}}^r, \boldsymbol{\mu}_{\text{out}}^r$	Mean residual-stream activations over locations inside / outside region r at layer ℓ	$\mathbb{R}^d$; over the 2400-city probe
$\hat{\boldsymbol{v}}_r$	Region direction at layer ℓ (mean-difference, unit-normalised)	$\hat{\boldsymbol{v}}_r \in \mathbb{R}^d$, $\|\hat{\boldsymbol{v}}_r\|_2 = 1$ (Eq. 2)
$\mathbf{h}_{\text{last}}^{(t)}$	Last-token residual at the intervention layer $L = L_{\text{final}} - 4$, step t	$\mathbb{R}^d$
α	Intervention strength (norm-matched)	$\alpha \in \{-2, -1, -0.5, 0, +0.5, +1, +2\}$ (Eq. 4)
$M(\hat{\boldsymbol{v}}, \alpha)$	Per-path mean first-token mass on the target region under $(\hat{\boldsymbol{v}}, \alpha)$	averaged over the 225-path bank
$\Delta_{\text{target}}(\hat{\boldsymbol{v}}, \alpha)$	Per-direction increase $M(\hat{\boldsymbol{v}}, \alpha) - M(\hat{\boldsymbol{v}}, 0)$	
$\Delta_{\text{spec}}(\alpha; r, r')$	Differential-increase specificity contrast vs control $\hat{\boldsymbol{v}}_{r'}$	Eq. 7

Table 4: Notations used throughout the paper. Symbols are grouped by category.

B Per-model geographic probing: prediction maps and layer-wise R^2

This appendix reports two diagnostics for each of the two large-model configurations probed in this study: a prediction map at the model's best probing layer and the layer-wise probe R^2.

Probe. The probe at each layer is a ridge regression that jointly predicts $(\mathrm{lat}, \mathrm{lon})$ in raw degrees from standardized last-token activations (no target standardization). Activations are extracted from the single-template prompt `Where is {city}?`; the prompt does not contain the country name. Targets are split $80/20$ stratified by country (test $n = 480$ cities). The split is city-level within country: for every WVS country, 80% of its cities are in train and 20% in test, so train and test share country labels. The probe can therefore exploit country-specific structure learned from train cities of the same country at test time; we make no cross-country generalisation claim. The split seed is fixed (`seed=42`). Activations are standardised by a `StandardScaler` fitted on the train split only, then applied unchanged to the test split. Ridge α is selected on the train split by 3-fold cross-validation over the grid $\alpha \in \mathrm{logspace}(0, 4, 7) \approx \{1, 4.64, 21.54, 100, 464.16, 2154.43, 10000\}$. The joint test R^2 is $1 - \mathrm{SS}_{\mathrm{res}}/\mathrm{SS}_{\mathrm{tot}}$ summed over the two raw target dimensions, so the joint score is variance-weighted toward longitude (wider range). Per-dimension R^2_{lat} and R^2_{lon} are reported alongside.

Best-layer selection. "Best probing layer" means the layer that maximises joint test R^2 on the held-out split. We follow the replication setup of Gurnee and Tegmark (2024) and do not hold out a separate validation split, so the reported peak R^2 values are post-selected on the test split. We report the full per-layer curve for both 70B+ models so the peak can be read against the surrounding plateau.

Marker encoding in the prediction maps. "Country-level RSA" is the Spearman correlation between the upper-triangular entries of the cosine-distance representational dissimilarity matrix (RDM) of activations of cities in a country and the upper-triangular entries of the haversine-distance RDM of those cities' coordinates, computed at the model's best probing layer over all 100 cities of that country in the WVS dataset. Marker color encodes WVS cultural region, all markers are filled circles, and marker size is proportional to country-level RSA up to a fixed upper clip of $\rho = 0.25$ (so all RSA values ≥ 0.25 render at the maximum size; nonpositive RSA renders at the minimum size). "Network depth" for layer ℓ in a model with valid layers $0 \ldots L_{\max}$ is reported as $100 \cdot \ell / L_{\max}$.

`Llama-3.1-70B` reaches a test-selected peak joint $R^2 = 0.831$ at 78.5% network depth.

`Qwen-2.5-72B` reaches a test-selected peak joint $R^2 = 0.759$ at 98.7% network depth. The lat–lon gap at the peak (0.070) is larger than `Llama-3.1-70B`'s (0.052), and the joint R^2 does not cross 0.7 until 43.0% network depth, against 12.7% for the similar-parameter-count `Llama-3.1-70B` (the two models differ in architecture, tokenizer, and training data; this is a descriptive contrast rather than a controlled comparison).

Summary across the two 70B+ models. Post-selected peak joint R^2 ranges from 0.759 (`Qwen-2.5-72B`) to 0.831 (`Llama-3.1-70B`); the post-selected peak layer falls at 78.5% depth or later on both models. The per-layer curves show that the plateau is broad in the Llama model and narrow (rising into the final layers) in the Qwen model. Latitude R^2 exceeds longitude R^2 at the peak in both models, with gaps of 0.052 (`Llama-3.1-70B`) and 0.070 (`Qwen-2.5-72B`). We draw no further claims from these figures beyond the per-model captions; they are reproduced here as the quantitative anchor for the lat/lon decodability claim in §4.2.

C Additional Details for the Representation Similarity Analysis

In Table 5, we provide detailed results of the representation similarity analysis, including the per correlation and the rank-based evaluation. Additionally, we provide results for smaller models such as `Llama-3.1-8b` (Grattafiori et al., 2024) and `Qwen-2.5-7B` (Yang et al., 2024). The full table covers all 8 WVS cultural regions, including Latin America. Latin America is reported here in the appendix but omitted from the main-text per-region figures for layout consistency with the 7-region navigation benchmark (Latin America and Oceania are excluded from the 225-path bank for the geographic-adjacency reason in the footnote of §3.1). The RSA evidence in this table thus remains the canonical 8-region representation report.

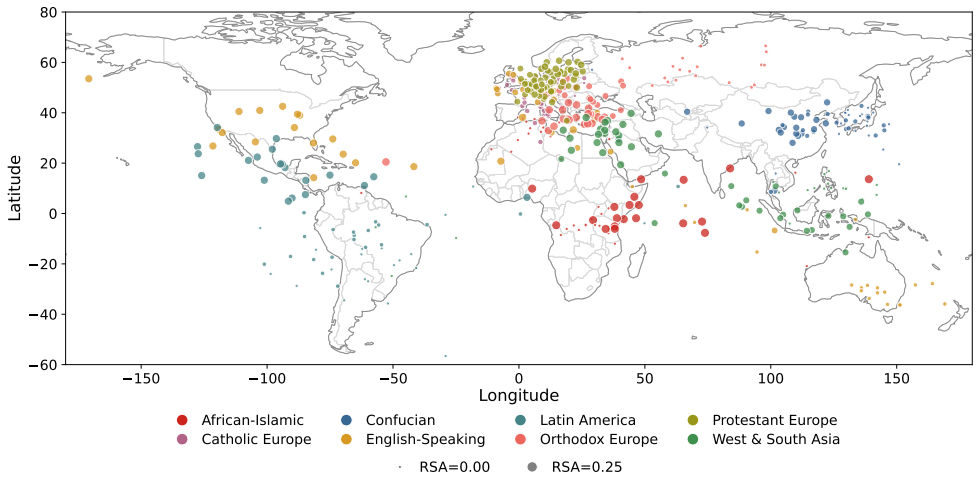

Figure 4: **Predicted city locations from `Llama-3.1-70B` activations.** Best probing layer is L_{62} (valid layers L_0–L_{79}; 78.5% depth; joint $R^2 = 0.831$, lat R^2 exceeds lon R^2 by 0.052). Marker color encodes WVS cultural region and marker size encodes country-level RSA.

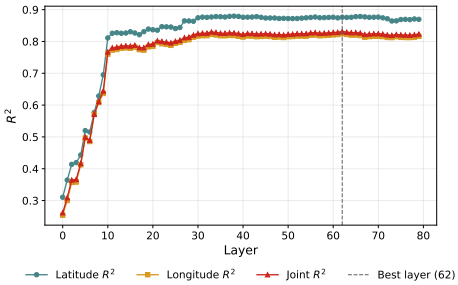

Figure 5: **Layer-wise probe R^2 for `Llama-3.1-70B`.** Joint, latitude, and longitude R^2 across all 80 valid layers (L_0–L_{79}). Joint R^2 first crosses 0.7 at L_{10} (12.7% depth), peaks at L_{62} ($R^2 = 0.831$, 78.5% depth), and falls to 0.823 at the final layer L_{79}. At the peak, lat R^2 exceeds lon R^2 by 0.052.

D `Qwen-2.5-72B` and Small-model behavioral cross-scale validation

Table 6 supplements the `Llama-3.1-70B` headline of Table 1 with three companion cells: the second large-model cell `Qwen-2.5-72B`, and the two small-model cells `Llama-3.1-8B` and `Qwen-2.5-7B`. All four panels report behavior under no intervention (α=0, hook disabled), so the only thing that varies across panels is the model.

`Qwen-2.5-72B` (large-model companion). Qwen-2.5-72B reproduces the same per-region behavioral asymmetry seen on Llama-3.1-70B: European regions (Catholic, Protestant, Orthodox) and the high-n_{paths} African-Islamic corridor lead nav-pass$_{45}$ and decoding confidence, while West & South Asia trails. The column-max pattern in Table 6 matches the Llama-3.1-70B panel of Table 1 cell-for-cell on the bolded categories. Magnitudes differ in absolute terms (Qwen's emit-conf. values are systematically less negative, reflecting higher per-token log-prob calibration on the same corridors), but the per-region *ordering* that underlies the four-pillar narrative is preserved across model families at the headline scale.

Small-model cells (`Llama-3.1-8B`, `Qwen-2.5-7B`). The two small-model cells reproduce the per-region behavioral asymmetry of §4.1: European regions lead nav-pass$_{45}$ and decoding confidence while African-Islamic and West & South Asia trail. The per-cell ranking is noisier at 7–8B (English-Speaking ranks high on confidence but low on nav-pass$_{45}$, reflecting small n_{paths}), consistent with the capacity-limited link reading below.

E Artifacts, licensing, and intended use

City corpus (RSA probe + CAA bank). The 2,400 cities (300 per WVS cultural region $\times$ 8 regions) used to compute residual-stream activations (§3.2) and the CAA direction banks (§2.3) are toponyms drawn from the GeoNames geographical database (`https://www.geonames.org`) via the Python library `GeonamesCache` (`https://pypi.org/project/geonamescache`). For each WVS country we retain the top-population cities reported by GeoNames until 100 cities per country are accumulated; per-region budgets are filled by uniform sampling across countries within each region up to the 300-city cap (full country→city listings in Appendix K). GeoNames does not sup-

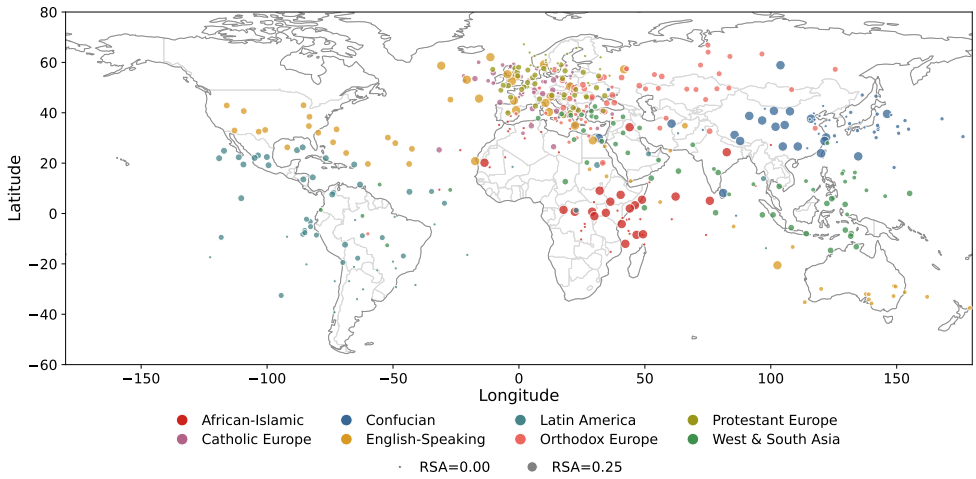

Figure 6: **Predicted city locations from `Qwen-2.5-72B` activations.** Best probing layer is L_{78} (valid layers L_0–L_{79}; 98.7% depth; joint $R^2 = 0.759$, lat $R^2 = 0.819$, lon $R^2 = 0.749$). Peak lon R^2 is 0.075 below `Llama-3.1-70B`.

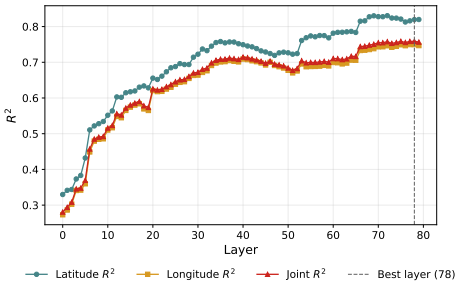

Figure 7: **Layer-wise probe R^2 for `Qwen-2.5-72B`.** Joint R^2 first crosses 0.7 at L_{34} (43.0% depth), crosses 0.75 at L_{69} (87.3% depth), and peaks at L_{78} ($R^2 = 0.759$, 98.7% depth); the final layer L_{79} reports $R^2 = 0.757$. At the peak, lat R^2 exceeds lon R^2 by 0.070.

ply a formal academic citation; we acknowledge it here by URL and license.

Navigation benchmark. The 225 country-level paths (§3.1) were constructed from country-neighbor relations supplied by `GeonamesCache`, partitioned into 75 European, 75 African, and 75 Asian start–destination pairs at country granularity. The benchmark is held fixed across all experiments and all models; there is no train/test split because no learning is performed over paths.

Licenses. GeoNames data is released under CC BY 4.0. `GeonamesCache` is released under the MIT License. The WVS country→region mapping is sourced from Inglehart and Welzel (2005); Haerpfer et al. (2020) and used as descriptive metadata under WVS terms of use. Pretrained models: `meta-llama/Llama-3.1-{8B,70B}` under the Llama 3.1 Community License; `Qwen/Qwen2.5-7B` under Apache-2.0; `Qwen/Qwen2.5-72B` under the Qwen License Agreement (https://huggingface.co/Qwen/Qwen2.5-72B/blob/main/LICENSE). Our released code and derived data (city lists, path bank, CAA directions, intervention results) are released under the MIT License for research use; derivatives of research-accessed data are not intended for non-research deployment.

Intended-use compatibility. GeoNames CC BY 4.0 explicitly permits research re-use with attribution; we redistribute only the small derived city-name and country-adjacency subsets used in our experiments, not the full GeoNames database, and provide the URL above as required attribution. Use of the WVS regional taxonomy is consistent with its descriptive purpose. Use of the four pretrained base models for representational and steering analysis is non-commercial research consistent with each model's license.

PII and offensive content. The artifacts consist of public toponyms (city names) and their WGS-84 coordinates. There is no information that names or uniquely identifies individual people, and no user-generated content. We manually spot-checked the 2,400-city sample for obviously offensive strings and found none. The `Where is {city}?` prompts and the navigation prompt (Appendix G) are likewise toponym-only.

Documentation and coverage. Domain: world cities aggregated into 8 WVS cultural regions (African-Islamic, Catholic Europe, Confucian, English-Speaking, Latin America, Orthodox Europe, Protestant Europe, West & South Asia); full

country→city listing in Appendix K. *Latin America is fully reported in the appendix RSA table (Appendix C) but omitted from the main-text per-region figures for layout consistency with the* 225-*path navigation benchmark, which restricts to* 7 *regions: Latin America and Oceania are excluded from the path bank for the geographic-adjacency reason given in the footnote of §3.1.* Languages: all prompts and city names in English (city names transliterated to ASCII as supplied by GeoNames). Splits: RSA and CAA use the full 2,400-city corpus (no held-out split, no learning); the linear probe uses an $80/20$ city-level stratified-by-country split with n_{test}=480 and `seed=42` (Appendix B); the 225-path benchmark is held fixed and analysis-only. Total experiment trials: $\sim$35,000 across two model families (225 paths $\times$ $\{4{-}7\}$ α values $\times$ $\{1, 3, 5\}$ controls per cell). Full compute, hyperparameter, and package details in Appendix F.

F Compute, hyperparameters, and software stack

Models and parameter counts. We study four open-weight base models: `meta-llama/Llama-3.1-8B` ($\sim$ 8.0 B parameters, 32 residual-stream layers, hidden width 4096), `meta-llama/Llama-3.1-70B` ($\sim$ 70.6 B parameters, 80 layers, hidden width 8192), `Qwen/Qwen2.5-7B` ($\sim$7.6 B parameters, 28 layers, hidden width 3584), and `Qwen/Qwen2.5-72B` ($\sim$ 72.7 B parameters, 80 layers, hidden width 8192). All four are used in their pretrained BASE configuration with no fine-tuning, no instruction tuning, and no LoRA adapters. Inference is fp16.

Hardware. Llama-3.1-8B and Qwen-2.5-7B run on a single NVIDIA L40S GPU (48 GB). `Llama-3.1-70B` and `Qwen-2.5-72B` are sharded across an A100 80 GB $\times$ 2 configuration with `device_map="auto"` from `accelerate`.

Total computational budget. No model training is performed; all compute is forward-pass inference plus hook arithmetic. Approximate GPU time: $\sim$ 14 GPU-hours on L40S for the two small-model cells and $\sim$16 GPU-hours on A100 80 GB $\times$ 2 for the two large- model cells. Total trials: $\sim$35,000 (225 paths $\times$ $\{4{-}7\}$ α values $\times$ $\{1, 3, 5\}$ controls per cell, across 4 model cells), plus a 2,400-city probe forward pass per model for the RSA and CAA banks.

Hyperparameter search. We perform *no* hyperparameter search over the steering direction, intervention layer, or path bank: the layer $L = L_{\text{final}} - 4$ (= L_{28} on the 8B model, L_{24} on the 7B model, L_{76} on the two 70B+ models) is pre-registered following Wurgaft et al. (2026); the α grid $\{-2, -1, -0.5, 0, +0.5, +1, +2\}$ is pre-registered before unblinding the sweep; the 225-path bank is held fixed and analysis-only. The only learned hyperparameter is the Ridge λ for the linear lat/lon probe, selected on the train split by 3-fold cross-validation over the grid $\lambda \in \text{logspace}(0, 4, 7) \approx \{1, 4.64, 21.54, 100, 464.16, 2154.43, 10000\}$ (Appendix B). Probe random seed `seed=42`; split is $80/20$ city-level stratified by country with $n_{\text{test}} = 480$. Decoding is greedy (`do_sample=False`, `max_new_tokens=30`) with a 25-step path cap.

Descriptive statistics (reporting protocol). All reported numbers are point estimates from a single greedy decoding run per (model, path, α, direction) cell — not max-of-N, not mean over re-runs. Greedy decoding is deterministic given the model + prompt + intervention, so re-running yields bit-identical outputs and stochasticity is absent from the decoding stage. Variability is reported via *path-paired bootstrap* with $n_{\text{boot}} = 10{,}000$ resamples and seed 42, yielding 95% percentile confidence intervals on all Δ_{target}, Δ_{spec}, prob-mass, target-share, nav-pass$_{45}$, confidence, and path-adjacency means (§2.3, §4.4). Cross-region significance uses region-clustered permutation tests with 10,000 permutations. Multiplicity is controlled by the Holm-Bonferroni adjustment over the explicit family of 2 large- model comparisons in the headline Δ_{spec} table (§4.4). RSA bootstrap CIs in Appendix C use 1000 resamples of the 300×300 pairwise dissimilarity tables. All confidence intervals reported in tables and figures are 95% unless noted otherwise.

Software stack and existing packages. PyTorch $\geq$ 2.1 (residual-stream hooks via `register_forward_hook`). `transformers` 4.46.3 and `accelerate` (model loading, `device_map="auto"` sharding). `huggingface-hub` (model weight download). `scikit-learn` (`Ridge` regression and `StandardScaler` for the lat/lon probe, fitted on the train split with default settings except the Ridge λ grid above). `scipy.stats.spearmanr` for the RSA Spearman correlations on the upper triangle of the 300×300 activation and great-circle

dissimilarity tables. `numpy`, `pandas`, `pyarrow` for data handling. `GeonamesCache` (MIT-licensed Python wrapper around the GeoNames database) for city names, populations, and country-neighbor relations (Appendix E). `pycountry` (ISO 3166 country-name resolution for path emissions, used with a 46-alias project extension shipped with the released code). No fine-tuning libraries, parameter-efficient adapters, RLHF tooling, or LLM judges are used.

Reproducibility. Bootstrap seed 42, $n_{\text{boot}} = 10{,}000$; permutation seed 42, 10,000 permutations; probe split seed 42. The 2,400-city RSA corpus, the 225-path navigation bank, and the WVS country→region table are held fixed across experiments and ship with the released code (MIT license, Appendix E).

G Navigation prompt template

The verbatim multi-turn template fed to both models at step t=0 (§3.1):

```
You are a Navigator. You are given a
start location and destination. You
must plan the shortest path from the
start location to the destination.
Output a sequential path following this
exact schema.
You should use the format: [Start
Location]-> [Waypoint] -> . . . ->
[Destination]. Do not output anything
else. ONLY THE FINAL OUTPUT.
*EXAMPLES**:
INPUT:
    - Start location: India
    - End location: China
OUTPUT: India -> Nepal -> Bhutan ->
China.

*INPUT**:
    Start location: {loc}
    Destination: {dest}.
OUTPUT: [{loc}] -> __
```

For $t \geq 1$ the prompt is the previous prompt with the parsed waypoint and the literal string `->` appended. Decoding is greedy (`do_sample=False`, `max_new_tokens=30`) and a path terminates when the parsed waypoint equals the destination, an empty waypoint is emitted, or 25 steps are reached.

H Hook implementation

The norm-matched CAA-add hook in Eq. (4) is implemented as a vanilla `register_forward_hook` on the residual-stream output of layer L. The relevant lines (paraphrased from `src/geo_atlas/steering/hooks.py`):

```
def hook_fn(module, inputs, output):
  h = (output[0]
       if isinstance(output, tuple)
       else output)
  last = h[:, -1:, :]
  norm_h = last.float().norm(dim=-1,
                             keepdim=True)
  delta = alpha * w_hat * norm_h.to(h.dtype)
  h[:, -1:, :] = last + delta
  pb = (last * w_hat).sum(-1)
  pa = (h[:, -1:, :] * w_hat).sum(-1)
  # G0 invariant: (pa - pb) / norm_h == alpha
```

The hook is attached *every step* of multi-turn generation (one waypoint per `generate()` call). `w_hat` is the unit-norm CAA direction loaded from `data/processed/caa_dirs/<short_key>_caa.npz`, keyed by `"<region>__layer_<L>"`.

I Causal-steering supplementary tables

Llama-3.1-8B L_{28} steering reference. Table 7 is the validity reference cell for the validity-matched intervention strength. Raw α is not comparable across families: `Qwen` residual norm $\|h_{\text{last}}\|$ is roughly 7× that of `Llama`, so the same nominal α displaces a much larger absolute residual. We define the *validity-matched intervention strength* as the largest α at which the strict path-adjacency validity remains at or above a fixed reference floor of ≈ 0.24. `Llama` satisfies this at α=+1; `Qwen` drops below the floor at α=+1, so we report it at α=+0.5 (validity 0.62).

Small-model steering summary. Table 8 reports the Δ_{target} and distance-dependent Δ_{spec} for Llama-3.1-8B L_{28} and Qwen-2.5-7B L_{24} at α=+1. Both small-model cells preserve specificity against both Catholic (cos=+0.81 near-neighbor) and English-Speaking (cos=−0.15 far) controls, in contrast to the 70B+ main cells (§4.4). Within-family magnitude transfers across Llama scales: L8B L28 $\Delta_{\text{target}} = +0.097$ $[+0.079, +0.116]$ vs. L70B L76 $\Delta_{\text{target}} = +0.101$ $[+0.077, +0.127]$.

Llama-3.1-8B near-neighbor adversarial. On Llama-3.1-8B L_{28}, the same-α paired contrast between Protestant CAA-add and Catholic CAA-add at $\alpha = +1$ (i.e. $M(\hat{w}_{\text{Protestant}}, +1) - M(\hat{w}_{\text{Catholic}}, +1)$ per path) yields $\Delta = +0.094$ $[+0.076, +0.113]$ over 225 paths, with 224/225 paths positive (sign $p \leq .0002$). Note this is the same-α contrast, distinct from the differential-increase Δ_{spec} of Eq. 7; the two agree in sign at $\alpha = +1$ here because the Catholic-direction α=0 baseline equals the Protestant-direction baseline by construction (the hook is the identity at α=0). De-

spite the very high cosine $\cos = +0.81$ between the two directions, Protestant CAA contains a target-specific component.

Mid-stack layer-sensitivity probe. Llama-3.1-8B L_{24} ($= L_{\text{final}} - 8$) at $\alpha = +1$ yields $\Delta_{\text{target}} = +0.192$ $[+0.167, +0.219]$, roughly $2\times$ the L_{28} ($L_{\text{final}} - 4$) effect. The result deviates from the pre-registered homologous-layer rule and is reported only as an exploratory mid-stack probe.

Random-direction null. We computed a matched-norm random-direction null with 10 seeds (5 from v18, 5 added in v19), each evaluated at $\alpha \in \{-1, -0.5, +0.5, +1\}$ on the 225-path bank. Zero of ten seeds reach the observed `Llama` $\Delta_{\text{target}} = +0.101$ at any positive α. The paired random-null mean is $+0.005$ $[-0.012, +0.022]$.

J Additional Information on Representational Similarity Across Layers

We report how well activation D_a^ℓ and geographic D_g geometry align across LLMs depths. This is broken down by cultural region.

The unit of analysis here is per-city entity RSA aggregated over each cultural region r: for each city, we compute the Spearman correlation $\rho_i^{(\ell)}$ between its row in the activation RDM and the corresponding row in the haversine-distance geographic RDM (see §2). We compute this for each of the eight regions, excluding the diagonal entries of the matrices.

Each WVS cultural region in our dataset contains $n = 300$ cities. All per-region, per-layer aggregates are based on the same sample size, and the per-city Spearman row length is 299 for every city.

In 8, we the mean entity RSA over cities within each region as a function of layer. We include 95% dispersion bands obtained by resampling cities within the region with replacement at the city level (1,000 resamples). Because the per-city rows of an RDM are not independent, these bands are descriptive of city-level dispersion rather than a strict inferential CI on the population mean; we make no inferential claims about peak *layer location*, only point-estimate observations. Reported regional peaks are also confounded by the geographic spread, country composition, and city sampling within each WVS region; we make no causal interpretation of between-region peak differences. The dashed horizontal line at $\rho = 0$ is the zero-correlation reference, not a calibrated null distribution. Curves use the WVS-region color palette to identify regions; Higher RSA at a layer means a larger Spearman correlation between activation-distance rows and geographic-distance rows for cities within that region. "Network depth" uses $100 \cdot \ell / L_{\text{max}}$ over valid layers.

In `Llama-3.1-70B` (Fig. 8), the region with the highest value is Protestant Europe ($\rho = 0.502$), while the region with the lowest value is English-Speaking ($\rho = 0.161$). Six of eight regions peak in layers L78–L79 ($\geq 98.7\%$ depth); the two exceptions are Catholic Europe at L63 (79.7% depth) and English-Speaking at L59 (74.7% depth).

For `Qwen-2.5-72B` (Fig. 9), Protestant Europe is again the highest-peaking region ($\rho = 0.348$), while English-Speaking is the lowest-peaking region ($\rho = 0.096$). Five regions peak at L78 (98.7% depth), whereas African–Islamic, West & South Asia, and English-Speaking peak in the first half of the network, at 43.0%, 1.3%, and 41.8% depth respectively.

Summary across 70B/72B models. The cross-model statements below compare raw regional RSA peaks without normalizing for layer count, region composition, post-training, or architecture family; we report them as within-figure observations on the two captioned models, not as causal claims about model design. Across both Llama-3.1-70B and Qwen-2.5-72B (Figs. 8–9), the highest point-estimate regional peak is Protestant Europe's, and the lowest is the English-speaking region. The peak RSA values are consistently higher in Llama-3.1-70B than in Qwen-2.5-72B across the reported regions: for example, Protestant Europe is 0.154 higher in Llama-3.1-70B (0.502 vs. 0.348), Catholic Europe is 0.185 higher (0.471 vs. 0.286), Orthodox Europe is 0.187 higher (0.429 vs. 0.242), and English-Speaking is 0.065 higher (0.161 vs. 0.096).

The depth pattern also differs between the two large models. Llama-3.1-70B has six of eight regions peaking in the final two valid layers, L78–L79 ($\geq 98.7\%$ depth), with Catholic Europe and English-Speaking peaking earlier. Qwen-2.5-72B also has five regions peaking near the end of the network at L78 (98.7% depth), but three regions peak before 50% depth: African–Islamic at L34 (43.0%), West & South Asia at L1 (1.3%), and English-Speaking at L33 (41.8%). Thus, among

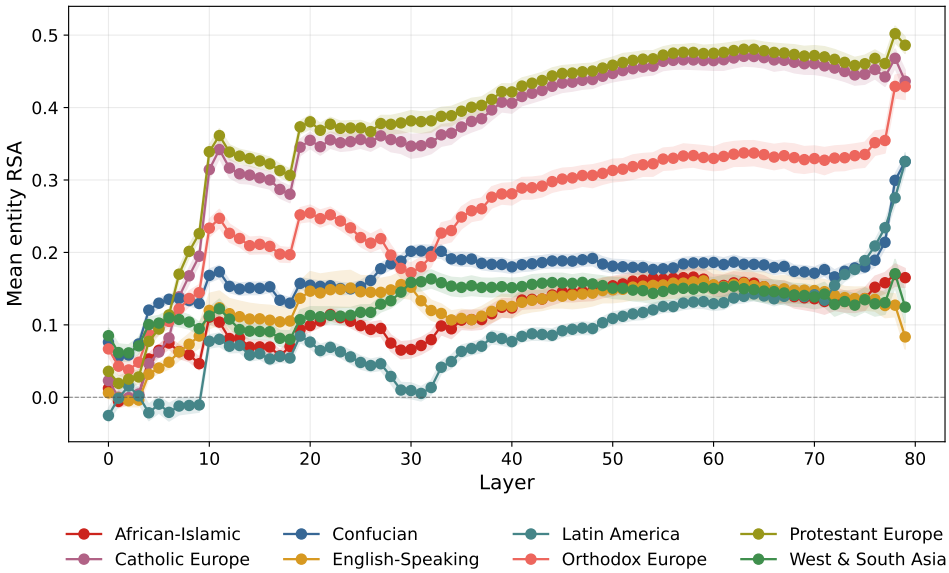

Figure 8: **Region-aggregated entity RSA across layers for Llama-3.1-70B.** Shaded bands are descriptive resampled 95% dispersion bands (see opening paragraph). Peaks: Protestant Europe $\rho = 0.502$ (L78), Catholic Europe $\rho = 0.471$ (L63), Orthodox Europe $\rho = 0.429$ (L78), Latin America $\rho = 0.326$ (L79), Confucian $\rho = 0.325$ (L79), West & South Asia $\rho = 0.171$ (L78), African–Islamic $\rho = 0.170$ (L78), English-Speaking $\rho = 0.161$ (L59).

the 70B/72B models, both show late-layer regional peaks for most regions, but Qwen-2.5-72B exhibits substantially earlier peaks for three lower-RSA regions.

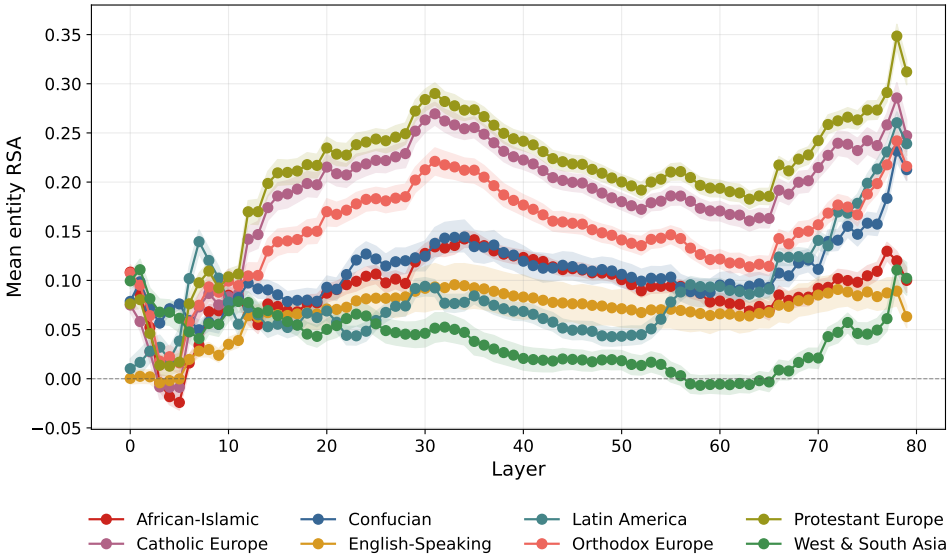

Figure 9: **Region-aggregated entity RSA across layers for Qwen-2.5-72B.** Shaded bands are descriptive resampled 95% dispersion bands (see opening paragraph). Peaks: Protestant Europe $\rho = 0.348$ (L78), Catholic Europe $\rho = 0.286$ (L78), Latin America $\rho = 0.260$ (L78), Orthodox Europe $\rho = 0.242$ (L78), Confucian $\rho = 0.232$ (L78), African–Islamic $\rho = 0.142$ (L34, 43.0% depth), West & South Asia $\rho = 0.111$ (L1, 1.3% depth), English-Speaking $\rho = 0.096$ (L33, 41.8% depth). Five regions peak at L78 (98.7% depth); three peak before 50% depth.

Cultural Region	Model	Rank			Spearman ρ		
		Mean	$CI_{0.025}$	$CI_{0.975}$	Mean	$CI_{0.025}$	$CI_{0.975}$
African-Islamic	Llama-3.1-70B	6.368	6.039	6.737	0.159	0.144	0.173
African-Islamic	Llama-3.1-8B	5.500	4.367	6.633	0.181	0.150	0.193
African-Islamic	Qwen-2.5-72B	5.224	4.961	5.513	0.121	0.107	0.136
African-Islamic	Qwen-2.5-7B	6.192	5.808	6.885	0.114	0.104	0.126
Catholic Europe	Llama-3.1-70B	2.145	2.000	2.289	0.462	0.449	0.477
Catholic Europe	Llama-3.1-8B	2.267	1.900	2.567	0.415	0.397	0.432
Catholic Europe	Qwen-2.5-72B	2.263	2.000	2.526	0.253	0.236	0.265
Catholic Europe	Qwen-2.5-7B	3.308	2.462	4.269	0.254	0.235	0.274
Confucian	Llama-3.1-70B	3.882	3.697	4.105	0.202	0.177	0.221
Confucian	Llama-3.1-8B	5.033	4.165	6.367	0.149	0.104	0.158
Confucian	Qwen-2.5-72B	4.408	4.211	4.658	0.154	0.130	0.172
Confucian	Qwen-2.5-7B	2.923	1.769	4.000	0.114	0.103	0.130
English-Speaking	Llama-3.1-70B	5.934	5.526	6.250	0.146	0.128	0.172
English-Speaking	Llama-3.1-8B	6.167	5.600	6.467	0.110	0.087	0.139
English-Speaking	Qwen-2.5-72B	6.632	6.329	6.855	0.088	0.071	0.108
English-Speaking	Qwen-2.5-7B	7.269	6.769	7.808	0.083	0.068	0.101
Latin America	Llama-3.1-70B	7.566	7.118	8.000	0.165	0.126	0.200
Latin America	Llama-3.1-8B	5.767	5.033	6.434	0.185	0.147	0.198
Latin America	Qwen-2.5-72B	5.763	5.013	6.513	0.146	0.087	0.200
Latin America	Qwen-2.5-7B	3.962	2.615	5.154	0.218	0.164	0.255
Orthodox Europe	Llama-3.1-70B	3.079	3.000	3.171	0.348	0.322	0.367
Orthodox Europe	Llama-3.1-8B	3.100	3.000	3.333	0.311	0.280	0.329
Orthodox Europe	Qwen-2.5-72B	3.053	2.934	3.171	0.205	0.184	0.220
Orthodox Europe	Qwen-2.5-7B	4.000	3.385	4.769	0.176	0.160	0.200
Protestant Europe	Llama-3.1-70B	1.092	1.000	1.184	0.477	0.463	0.488
Protestant Europe	Llama-3.1-8B	1.200	1.000	1.367	0.430	0.407	0.438
Protestant Europe	Qwen-2.5-72B	1.158	1.000	1.289	0.282	0.264	0.295
Protestant Europe	Qwen-2.5-7B	2.885	1.308	4.539	0.283	0.264	0.295
West & South Asia	Llama-3.1-70B	5.934	5.566	6.421	0.147	0.130	0.163
West & South Asia	Llama-3.1-8B	6.967	6.233	7.867	0.077	0.056	0.093
West & South Asia	Qwen-2.5-72B	7.500	7.211	7.921	0.055	0.036	0.067
West & South Asia	Qwen-2.5-7B	5.462	3.808	6.885	0.065	0.051	0.090

Table 5: Representational quality per cultural region and a model (along with the large models used in the analysis, we provide results for smaller counterparts). **Rank** is the within-model rank by mean per-city Spearman ρ_i averaged across layers, where 1 is the best, and 8 is the worst rank (Note, we drop bottom 5% layers by global RSA score to compute the mean rank). **Spearman** ρ is the mean per-city correlation between activation and geographic distance matrices (over the top 20% of layers by global RSA score). All 95% CIs are computed via two-level block-bootstrapped intervals.

Region	n_{paths}	PRECISION$_r$	RECALL$_r$	REC$_r^{\text{pres}}$	emit-conf.$_r$	nav-pass$_{45,r}$
`Qwen-2.5-72B`						
Protestant Europe	38	0.88	**1.00**	0.55	-0.39	0.74
Catholic Europe	63	0.95	0.95	0.79	-0.39	0.80
Orthodox Europe	52	**1.00**	0.92	0.62	-0.50	0.91
English-Speaking	9	**1.00**	0.67	0.22	**-0.35**	**1.00**
Confucian	28	**1.00**	**1.00**	0.39	-0.69	0.71
African-Islamic	**135**	0.95	0.94	**0.93**	-0.50	0.80
West & South Asia	46	0.80	**1.00**	0.65	-0.51	0.77
`Llama-3.1-8B`						
Protestant Europe	38	0.86	0.86	0.66	-0.57	0.72
Catholic Europe	63	0.94	0.89	0.76	-0.63	0.69
Orthodox Europe	52	0.92	0.92	0.73	-0.70	0.60
English-Speaking	9	**1.00**	0.67	0.22	**-0.43**	**1.00**
Confucian	28	0.75	**1.00**	0.43	-0.79	0.60
African-Islamic	**135**	0.90	0.89	**0.94**	-0.71	0.64
West & South Asia	46	0.71	**1.00**	0.74	-0.72	0.55
`Qwen-2.5-7B`						
Protestant Europe	38	**1.00**	0.86	0.50	-0.81	0.87
Catholic Europe	63	**1.00**	0.89	0.63	-0.73	0.92
Orthodox Europe	52	**1.00**	0.92	0.63	-0.80	0.88
English-Speaking	9	**1.00**	0.67	0.22	-0.82	**1.00**
Confucian	28	**1.00**	**1.00**	0.54	-0.93	0.61
African-Islamic	**135**	0.95	0.89	**0.90**	-0.72	0.82
West & South Asia	46	0.86	**1.00**	0.61	**-0.67**	0.79

Table 6: Per-region α=0 navigation summary for the non-main cells: `Qwen-2.5-72B` (large-model companion to the `Llama-3.1-70B` headline in Table 1), and the two small-model cross-scale / cross-family cells (`Llama-3.1-8B`, `Qwen-2.5-7B`). Column definitions and bolding rule are identical to Table 1.

α	dest	f-tgt	prob	hal.	valid
-2	0.04	0.000	0.0000	0.98	0.00
-1	0.89	0.014	0.0033	0.15	0.40
-0.5	0.98	0.046	0.0258	0.05	0.64
0	1.00	0.052	0.0398	0.02	0.77
$+0.5$	1.00	0.080	0.0617	0.01	0.63
$+1$	0.95	0.169	0.1364	0.05	0.24
$+2$	0.11	0.270	0.2057	0.69	0.00

Table 7: **Target probability mass rises monotonically with α while strict path-validity is non-monotone, defining a safe in-distribution steering window $\alpha \in [-0.5, +0.5]$ on Llama-3.1-8B L_{28}.** Intervention strength (α) sweep reference table (N=225). `dest`=fraction of trajectories reaching a valid destination; `f-tgt`=fraction of emissions in Protestant Europe; `prob`=mean per-step target probability mass; `hal.`=hallucination rate; `valid`=strict path-adjacency validity.

			Catholic ctrl		English ctrl	
Cell	α	Δ_{target}	cos	Δ_{spec}	cos	Δ_{spec}
Llama-3.1-8B L28	$+1$	$+0.097$ [$+0.079$, $+0.116$]	$+0.81$	$+0.094$ ✓	-0.15	$+0.134$ ✓
Qwen-2.5-7B L24	$+1$	$+0.198$ [$+0.173$, $+0.223$]	$+0.43$	$+0.212$ ✓	-0.01	$+0.186$ ✓

Table 8: **At 7–8B scale near-neighbor specificity holds against the Catholic-Europe control on both families, in contrast to the L70B L76 cell where it collapses; distance-dependent specificity emerges at scale rather than being intrinsic to the intervention.** Small-model supplement to Table 3: Δ_{target} and distance-dependent specificity at 7–8B (Llama-3.1-8B L28, Qwen-2.5-7B L24), all at α=$+1$. cos is computed inside each model's own activation space: `Llama`-family Protestant–Catholic cosine stays in the high-collinearity band ($+0.81$ at L8B L28, $+0.72$ at L70B L76); `Qwen`-family cosine is roughly half as large ($+0.43$ at L24, $+0.37$ at L76).

K Tables

Table 9: Cities grouped by cultural region and country.

Region	Country	Cities
African-Islamic	Algeria	Abou El Hassan, Adrar, Aflou, Ain El Hadjel, Algiers, Annaba, Aoulef, Arris, Azazga, Ain Bessem, Ain Defla, Ain El Bya, Ain Kercha, Ain Sefra, Ain Touta, Babor, Barika, Beni Saf, Berriane, Berrouaghia, Bir El Ater, Birine, Bordj Menaiel, Bou Saada, Bougaa, Boumahra Ahmed, Brezina, Bechar, Bejaia, Charef, Chellalat El Adhaouara, Cheria, Chorfa, Dar Chioukh, Didouche Mourad, Djamaa, El Abiodh Sidi Cheikh, El Achir, El Bayadh, El Eulma, El Hadjira, El Kala, El Meghaier, El Menia, El Milia, Feidh El Botma, Frenda, Guerara, Hammam Bou Hadjar, Hassi Bahbah, Hassi Messaoud, In Salah, Jijel, Khenchela, Ksar Chellala, Laghouat, Lardjem, M'Sila, Maghnia, Mascara, Mazouna, Mecheria, Mehdia Daira De Meghila, Melouza, Meskiana, Messaad, Metlili Chaamba, Mostaganem, Ouargla, Oued Fodda, Ouenza, Ouled Mimoun, Oum El Bouaghi, Reggane, Reguiba, Robbah, Rouached, Ras El Aioun, Ras El Oued, Saida, Sebdou, Sidi Bel Abbes, Sidi Khaled, Sidi Okba, Skikda, Souk Ahras, Tamanghasset, Taibet, Telerghma, Theniet El Had, Tiaret, Timimoun, Tindouf, Tipasa, Tolga, Zemoura, Zeribet El Oued, 'Ain Deheb, 'Ain El Melh, 'Ain El Turk
African-Islamic	Kenya	Awendo, Bondo, Bungoma, Busia, Chuka, El Wak, Eldama Ravine, Eldoret, Emali, Embu, Garba Tula, Garissa, Gilgil, Githunguri, Habaswein, Hola, Homa Bay, Isebania, Isiolo, Iten, Kabarnet, Kajiado, Kakamega, Kakuma, Kapenguria, Kapsabet, Karatina, Kericho, Kerugoya, Khwisero, Kiambu, Kihancha, Kikuyu, Kilifi, Kimilili, Kiminini, Kinoi, Kiserian, Kisii, Kisumu, Kitale, Kitengela, Kitui, Lafey, Lamu, Limuru, Lodwar, Lugulu, Maai Mahiu, Machakos, Mairo Inya, Makindu, Makueni Boma, Malaba, Malindi, Mandera, Maralal, Mariakani, Marsabit, Masalani, Maua, Mazeras, Mbale, Meru, Migori, Moi'S Bridge, Molo, Mombasa, Moyale, Mtwapa, Muhoroni, Mumias, Murang'A, Mwingi, Nairobi, Naivasha, Nakuru, Nanyuki, Narok, Njoro, Nyamira, Nyeri, Ogembo, Ol Kalou, Oyugis, Ramu, Rongai, Rongo, Ruiru, Siaya, Takaba, Taveta, Thika, Ukunda, Upper Gilgil, Voi, Wajir, Wanguru, Watamu, Webuye
African-Islamic	Tanzania	Arusha, Basotu, Biharamulo, Boma La Ngombe, Bukoba, Bunda, Chala, Chanika, Chato, Dodoma, Galappo, Hedaru, Ifakara, Igugunu, Igurusi, Ikungi, Ikwiriri, Iringa, Izazi, Kabanga, Kahama, Kasamwa, Kasulu, Kibaigwa, Kibakwe, Kibondo, Kidodi, Kilindoni, Kilosa, Kilwa Masoko, Kirando, Kiratu, Kondoa, Kyela, Laela, Lembeni, Lindi, Liwale, Lukuledi, Lushoto, Madizini, Mafinga, Mahanje, Makongolosi, Makuyuni, Malinyi, Manyoni, Mapinga, Maswa, Matiri, Mbulu, Mgandu, Mhango, Mkuranga, Mlangali, Morogoro, Mpanda, Mtwango, Mtwara, Mugumu, Muheza, Muriti, Mvomero, Mwandiga, Mwanhuzi, Mwanza, Nangomba, Nanyamba, Newala Kisimani, Ngerengere, Ngudu, Nguruka, Nshamba, Nsunga, Nungwi, Nyalikungu, Nzega, Ruangwa, Rujewa, Shelui, Shinyanga, Shirati, Sikonge, Sirari, Sokoni, Songea, Sumbawanga, Tarakea, Tingi, Tumbi, Tunduma, Tunduru, Urambo, Usevia, Ushirombo, Utengule, Uvinza, Uyovu, Wete, Zanzibar
Catholic Europe	Belgium	Aalter, Aarschot, Andenne, Arlon, Asse, Ath, Beveren, Bilzen, Binche, Blankenberge, Braine-L'Alleud, Braine-Le-Comte, Brussels, Chaudfontaine, Chatelet, Courcelles, Deinze, Denderleeuw, Dendermonde, Diest, Diksmuide, Eeklo, Ekeren, Eupen, Fleurus, Frameries, Geel, Gembloux, Genk, Geraardsbergen, Halle, Heist-Op-Den-Berg, Herentals, Herstal, Heusy, Heverlee, Hoogstraten, Huy, Ieper, Jambes, Kalmthout, Kasterlee, Knokke-Heist, Koersel, Koksijde, Lanaken, Lede, Lochristi, Lokeren, Lommel, Maaseik, Maasmechelen, Manage, Marche-En-Famenne, Mariakerke, Meise, Merelbeke, Mol, Mortsel, Mouscron, Neerpelt, Nijlen, Nivelles, Ostend, Oudenaarde, Peer, Poperinge, Puurs, Peruwelz, Riemst, Rixensart, Roeselare, Ronse, Rotselaar, Schilde, Seraing, Sint-Katelijne-Waver, Sint-Kruis, Sint-Niklaas, Sint-Truiden, Soumagne, Stekene, Tervuren, Tessenderlo, Tielt, Tienen, Tongeren, Torhout, Tournai, Turnhout, Vise, Walcourt, Waregem, Wervik, Westerlo, Zedelgem, Zemst, Zonhoven, Zottegem, Zwevegem
Catholic Europe	Czechia	Benesov, Beroun, Blansko, Bohnice, Brandys Nad Labem-Stara Boleslav, Brno, Bruntal, Bilina, Breclav, Cheb, Chomutov, Chrudim, Dvur Kralove Nad Labem, Decin, Frydek-Mistek, Havlickuv Brod, Havirov, Hlubocepy, Hodonin, Horni Mecholupy, Hradec Kralove, Hranice, Jablonec Nad Nisou, Jihlava, Jindrichuv Hradec, Jirkov, Jicin, Kadan, Karlovy Vary, Karvina, Kladno, Klatovy, Klasterec Nad Ohri, Kolin, Koprivnice, Kosire, Kralupy Nad Vltavou, Krnov, Kromeriz, Kutna Hora, Letnany, Liberec, Liben, Litomerice, Litvinov, Louny, Michle, Mlada Boleslav, Modrany, Most, Melnik, Neratovice, Novy Jicin, Nachod, Olomouc, Opava, Orlova, Ostrava, Ostrov, Otrokovice, Pardubice, Pelhrimov, Pilsen, Prague, Prostejov, Pisek, Prerov, Pribram, Rakovnik, Roznov Pod Radhostem, Slany, Sokolov, Stare Strasnice, Stary Bohumin, Strakonice, Svitavy, Teplice, Trutnov, Tabor, Trebic, Trinec, Uherske Hradiste, Uhersky Brod, Valasske Mezirici, Varnsdorf, Vsetin, Vyskov, Zlin, Znojmo, Usti Nad Labem, Usti Nad Orlici, Cerny Most, Ceska Lipa, Ceska Trebova, Ceske Budejovice, Cesky Tesin, Repy, Sumperk, Zatec, Zdar Nad Sazavou

Continued on next page

Table 9 – continued from previous page

Region	Country	Cities
Catholic Europe	Italy	Agropoli, Alba, Alghero, Aosta, Arezzo, Avezzano, Bagheria, Barcellona Pozzo Di Gotto, Bellaria-Igea Marina, Belluno, Benevento, Biella, Bressanone, Bronte, Campi Bisenzio, Canicatti, Canosa Di Puglia, Carbonia, Casale Monferrato, Cascina, Castrovillari, Catanzaro, Cecina, Chioggia, Chivasso, Civitavecchia, Cosenza, Crema, Crotone, Desenzano Del Garda, Enna, Fabriano, Faenza, Fano, Foggia, Gela, Genoa, Gragnano, Gravina In Puglia, Grosseto, Ischia, Isernia, Lancenigo-Villorba, Lecce, Lecco, Magenta, Massafra, Mazara Del Vallo, Mirandola, Mola Di Bari, Molfetta, Mondragone, Nuoro, Olbia, Oristano, Osimo, Pachino, Palagonia, Pallanza-Intra-Suna, Palmi, Perugia, Pescara, Pietrasanta, Piombino, Portogruaro, Potenza, Quartu Sant'Elena, Ragusa, Rieti, Rome, Rossano Stazione, Rovereto, Rovigo, Salsomaggiore Terme, Saluzzo, San Benedetto Del Tronto, San Giovanni Lupatoto, San Salvo, San Vito Dei Normanni, Sannicandro Garganico, Savona, Sciacca, Sestri Levante, Sezze, Siderno, Siena, Siracusa, Sondrio, Spoleto, Sulmona, Tortona, Trapani, Trieste, Udine, Ventimiglia, Viadana, Vibo Valentia, Vicenza, Vignola, Viterbo
Confucian	China	Aral, Babu, Bailin, Baiquan, Baise, Basuo, Beijing, Changli, Changsha, Chaohu, Chapchal, Chenghua, Chenzhou, Dalain Hob, Dalianwan, Dandong, Datong, Dongying, Dunhuang, Encheng, Erenhot, Erong, Fuyuan, Fuyun, Gar, Genhe, Guangyuan, Gujangbagh, Guozhen, Guyuan, Hami, Handan, Heihe, Hengshan, Hepingjie, Hongzhou, Huangpi, Huilong, Huizhou, Hunchun, Jinchang, Jinxiang, Jiuquan, Jiyuan, Kaitong, Kashgar, Korla, Laohekou, Lhoka, Lianjiang, Linghai, Linqiong, Lintan Chengguanzhen, Linyi, Liuzhi, Longquan, Luobei, Luohe, Lushui, Luliang, Mabai, Mangnai Zhen, Manzhouli, Mengcheng Chengguanzhen, Mengmao, Mujiayingzi, Nagqu, Nanchang, Nuofu, Nyingchi, Oymak, Pucheng, Qabqa, Qamdo, Qiemo, Qingdao, Qinnan, Rikaze, Salaqi, Sandaohezi, Sanming, Shangluo, Shangri-La, Shuanglong, Tacheng, Tieling, Turpan, Wenchang, Wuchang, Wuhai, Xichang, Xilinhot, Xilinji, Xiongjia, Yatou, Yongning, Yushu, Zhalantun, Zhicheng, Zhuji
Confucian	Japan	Aizu-Wakamatsu, Anan, Annaka, Asakuchi, Awaji, Chitose, Echizen, Fukuecho, Gero, Gion, Gose, Gotsucho, Hachinohe, Hagi, Hakodate, Hamanoichi, Hanawa, Hasaki, Hayashino, Hitoyoshi, Hotaka, Ikedacho, Iki, Imaichi, Ishigaki, Ishioka, Itogawa, Itoigawa, Itoman, Iyo, Izumo, Kamaishi, Kasugai, Kasumi, Katsuura, Kawaminami, Kochi, Komagane, Komatsu, Kunisaki, Kurihara, Kushimoto, Kushiro, Kutchan, Koshu, Mifune, Minamisatsuma, Mishohirajo, Miura, Miyakojima, Mombetsu, Moritake, Motomachi, Murakami, Muroran, Muroto-Misakicho, Mutsu, Nago, Nakashibetsu, Namie, Nayoro, Naze, Nemuro, Nichinan, Nishinoomote, Obihiro, Omaezaki, Owase, Rumoi, Sado, Shima, Shimo-Furano, Shimokizukuri, Shirone, Shizukuishi, Shizunai-Furukawacho, Shonai, Shobara, Susono, Takahagi, Takedamachi, Tanba, Tokyo, Toyohashi, Uonuma, Ushibukamachi, Wajima, Wakkanai, Watari, Yabase, Yachi, Yakumo, Yame, Yanai, Yatsuomachi-Higashikumisaka, Yurigaoka, Yurihonjo, Omihachiman, Omura, Owa
Confucian	South Korea	Andong, Ansan-Si, Anseong, Asan, Beolgyo, Boryeong, Boseong, Bucheon-Si, Busan, Buyeo, Changnyeong, Changsu, Changwon, Cheongju-Si, Cheongsong Gun, Chinch'On, Chinju, Chuncheon, Chungju, Daejeon, Damyang, Eisen, Gaigeturi, Ganghwa-Gun, Gapyeong, Gijang, Gimcheon, Gongju, Goseong, Gunsan, Gwangyang, Gyeongju, Gyeongsan-Si, Haenam, Heunghae, Hongch'On, Hongseong, Hwacheon, Hwado, Hwasun, Hwawon, Icheon-Si, Iksan, Imsil, Jangheung, Jangseong, Jeju City, Jenzan, Jeongeup, Jinan-Gun, Jumunjin, Kimhae, Kimje, Kinzan, Koch'Ang, Koesan, Kosong, Kunwi, Kurye, Kyosai, Miryang, Mokpo, Muan, Mungyeong, Munsan, Naju, Nangen, Neietsu, Okcheon, Osan, Sangju, Santyoku, Sejong, Seogwipo, Seongnam-Si, Seosan, Seoul, Sinan, Sokcho, Suncheon, Taesal-Li, Tangjin, T'Aebaek, Ulsan, Waegwan, Wanju, Wonju, Yanggu, Yangju, Yangp'Yong, Yangsan, Yecheon, Yeoju, Yeongam, Yeongdong, Yeonggwang, Yeongju, Yeosu, Yesan, Yonmu
English-Speaking	Australia	Adelaide Hills, Albany, Alexandra Hills, Alice Springs, Armidale, Australind, Bairnsdale, Ballarat, Banora Point, Bathurst, Bendigo, Bentleigh, Bracken Ridge, Broken Hill, Bundaberg, Burnie, Busselton, Byford, Caboolture, Cairns, Canberra, Carrum Downs, Castle Hill, Casula, Cessnock, Coffs Harbour, Coorparoo, Cronulla, Darwin, Deception Bay, Dee Why, Doreen, Dubbo, Ellenbrook, Engadine, Gawler, Geelong, Geraldton, Gladstone, Glenmore Park, Goulburn, Greenvale, Griffith, Hervey Bay, Hillside, Hobart, Labrador, Lara, Launceston, Lilydale, Logan City, Mackay, Maitland, Mandurah, Maroochydore, Maroubra, Maryborough, Mildura, Mitcham, Morley, Morphett Vale, Mount Druitt, Mount Gambier, Mount Isa, Mount Martha, Narre Warren, Newcastle, Orange, Pakenham, Palmerston, Paralowie, Port Macquarie, Preston, Redbank Plains, Robina, Rockhampton, Rockingham, Rowville, Scarborough, Shepparton, Strathfield, Sunbury, Sunshine West, Tamworth, Taree, The Gap, Toowoomba, Townsville, Traralgon, Tuggeranong Administrative District, Umina Beach, Upper Coomera, Wangaratta, Warragul, Warrnambool, Warwick, Willetton, Wodonga, Wollongong City Centre, Wyndham Vale

Continued on next page

Table 9 – continued from previous page

Region	Country	Cities
English-Speaking	United Kingdom	Aberdare, Aberdeen, Aberystwyth, Alloa, Arbroath, Ashford, Ashington, Aylesbury, Ayr, Banbury, Barrow In Furness, Barry, Bebington, Beverley, Bideford, Bishop Auckland, Bolton, Boston, Braunstone, Bridlington, Calne, Camberley, Carlisle, Carmarthen, Carterton, Chapeltown, Cheltenham, Chepstow, Christchurch, Coleraine, Deal, Derby, Derry, Droitwich, Dumfries, East Dereham, Eastbourne, Elgin, Felixstowe, Fylde, Gainsborough, Goole, Greenock, Grimsby, Harlow, Harpenden, Hayling Island, Hereford, Horsham, Huntingdon, Inverness, Isle Of Lewis, Keighley, Kempston, Kendal, Kettering, King'S Lynn, Kippax, Larne, Lisburn, Llandudno, London, Lowestoft, Macclesfield, March, Mendip, Musselburgh, Nantwich, Newark On Trent, Newbury, Newmarket, Newry, Omagh, Oswestry, Paignton, Penzance, Perth, Peterhead, Plymouth, Redcar, Ripon, Royston, Salisbury, Skegness, Southend-On-Sea, St Austell, Stamford, Sudbury, Sunderland, Sutton Coldfield, Swansea, Telford, Thetford, Tiverton, Tonbridge, Viewpark, Weymouth, Whitehaven, Winchester, Yeovil
English-Speaking	United States	Aberdeen, Altamont, Anchorage, Anthem, Atlantic City, Back Mountain, Bangor, Baton Rouge, Bellingham, Billings, Boise, Boone, Brawley, Brownsville, Brownwood, Bullhead City, Butte, Centralia, Chico, Converse, Cookeville, Corsicana, Corvallis, Del Rio, Dumas, East Moline, Elizabeth City, Elko, Enchanted Hills, Erlanger, Fairbanks, Fernley, Fort Leonard Wood, Garden City, Gardner, Gillette, Grand Forks, Grand Junction, Herriman, Hibbing, Hilo, Hilton Head Island, Hobbs, Idaho Falls, Juneau, Kalispell, Katy, Key West, Kirksville, Lake Worth Beach, Laramie, Laredo, Lincoln, Lompoc, Madera, Makakilo / Kapolei / Honokai Hale, Marquette, Marshfield, Mason City, Mattoon, Mckinleyville, Midway, Minot, New Britain, New Smyrna Beach, Newport Beach, Niagara Falls, Nogales, North Chicago, North Platte, Pacifica, Paducah, Pendleton, Plattsburgh, Pueblo, Ridgecrest, Ridgeland, Rock Springs, Romulus, Saint George, Sarasota, Searcy, Sioux Falls, Spokane, Sunland Park, Tahlequah, Talladega, Tallahassee, Texarkana, Traverse City, Tupelo, Warner Robins, Washington, Weirton Heights, Wenatchee, West Lynchburg, Wichita, Wichita Falls, Willmar, Wilmington
Latin America	Argentina	Arrecifes, Arroyito, Anatuya, Bahia Blanca, Buenos Aires, Caleta Olivia, Carmen De Patagones, Casilda, Catamarca, Catriel, Caucete, Chacabuco, Chajari, Chascomus, Chilecito, Colonia Wanda, Concepcion Del Uruguay, Concordia, Coronel Suarez, Cruz Del Eje, Curuzu Cuatia, Cutral-Co, Cordoba, Diamante, Dolores, El Bolson, El Soberbio, Esperanza, Esquel, Esquina, Famailla, Federal, Fontana, Formosa, General Jose De San Martin, General Pico, General Villegas, Gobernador Virasora, Goya, Gualeguay, Ituzaingo, Jardin America, Joaquin V. Gonzalez, Juan Jose Castelli, La Rioja, Laboulaye, Las Brenas, Lincoln, Mar Del Plata, Marcos Juarez, Mendoza, Mercedes, Morteros, Necochea, Olavarria, Palpala, Paso De Los Libres, Pirane, Plottier, Posadas, Profesor Salvador Mazza, Quitilipi, Rio Cuarto, Rio Gallegos, Rio Grande, Rio Tercero, Saladas, San Antonio Oeste, San Carlos De Bariloche, San Jorge, San Jose De Jachal, San Justo, San Luis, San Martin De Los Andes, San Nicolas De Los Arroyos, San Rafael, San Ramon De La Nueva Oran, San Vicente, Santa Elena, Santa Rosa, Santiago Del Estero, Tandil, Termas De Rio Hondo, Trelew, Tres Arroyos, Ushuaia, Veinticinco De Mayo, Venado Tuerto, Vera, Victoria, Villa Dolores, Villa Gesell, Villa Maria, Villa Mercedes, Villa Ocampo, Villa Regina, Villa Angela, Villaguay, Zapala, Zarate
Latin America	Mexico	Acaponeta, Acatlan De Osorio, Acambaro, Agua Dulce, Agua Prieta, Aguascalientes, Arriaga, Atoyac De Alvarez, Autlan De Navarro, Ayutla De Los Libres, Cabo San Lucas, Cacahoatan, Camargo, Cancun, Centro Familiar La Soledad, Champoton, Chetumal, Chiapa De Corzo, Chihuahua, Ciudad Constitucion, Ciudad De Armeria, Ciudad Del Carmen, Ciudad Fernandez, Ciudad Juarez, Ciudad Madero, Ciudad Mante, Ciudad Miguel Aleman, Ciudad Rio Bravo, Ciudad Sabinas Hidalgo, Ciudad Victoria, Comitan, Compostela, Costa Rica, Crucecita, El Salto, Empalme, Francisco I. Madero, Frontera, Guadalupe Victoria, Guamuchil, Hacienda Santa Fe, Heroica Caborca, Heroica Ciudad De Tlaxiaco, Huetamo De Nunez, Ixmiquilpan, Jose Cardel, La Orilla, La Paz, Libres, Linares, Loma Bonita, Los Mochis, Lazaro Cardenas, Madera, Manuel Ojinaga, Martinez De La Torre, Matehuala, Mazatlan, Mexico City, Miguel Aleman (La Doce), Motul, Navojoa, Nogales, Nueva Rosita, Nuevo Casas Grandes, Nuevo Laredo, Palenque, Parral, Parras De La Fuente, Petatlan, Peto, Piedras Negras, Pijijiapan, Pinotepa Nacional, Puerto Escondido, Puerto Penasco, Rodolfo Sanchez Taboada, Salinas De Hidalgo, Saltillo, San Fernando, San Jose Iturbide, Santa Maria Atzompa, Santiago Papasquiaro, Santo Domingo Tehuantepec, Soledad De Graciano Sanchez, Sombrerete, Tamazunchale, Teloloapan, Tijuana, Tizimin, Tlaltenango De Sanchez Roman, Tulancingo, Tulum, Uruapan, Valle De Bravo, Villahermosa, Victor Rosales, Yurecuaro, Zapotiltic, Alamo

Continued on next page

Table 9 – continued from previous page

Region	Country	Cities
Latin America	Peru	Abancay, Andahuaylas, Ayacucho, Ayaviri, Bagua Grande, Barranca, Barranco, Belen, Bellavista, Cajamarca, Callao, Camana, Catacaos, Chachapoyas, Chancay, Chaupimarca, Chepen, Chiclayo, Chilca, Chimbote, Chincha Alta, Chocope, Chongoyape, Chosica, Chulucanas, El Pedregal, Ferrenafe, Huacho, Huamachuco, Huancavelica, Huanta, Huaral, Huaraz, Huarmey, Huanuco, Ica, Ilave, Ilo, Jauja, Jaen, Juanjui, Juliaca, Junin, La Banda, La Breita, La Oroya, La Peca, La Rinconada, La Union, Lambayeque, Laredo, Lima, Mala, Minas De Marcona, Moche, Mollendo, Moquegua, Moyobamba, Nazca, Pacasmayo, Paijan, Paita, Paramonga, Picsi, Pimentel, Pisco, Piura, Puerto Maldonado, Puno, Patapo, Querecotillo, Rioja, San Fernando, San Vicente De Canete, Santa Ana, Santa Rosa, Santiago De Cao, Satipo, Sana, Sechura, Selva Alegre, Sicuani, Sullana, Tacna, Talara, Tambo Grande, Tambopata, Tarma, Tingo Maria, Tocache, Tumbes, Uchiza, Villa Poeta Jose Galvez Barrenechea, Viru, Wanchaq, Yanacancha, Yauri, Yunguyo, Yurimaguas, Zarumilla
Orthodox Europe	Greece	Acharnes, Agios Dimitrios, Agios Ioannis Rentis, Agrinio, Agia Paraskevi, Aigaleo, Alexandroupoli, Amaliada, Argyroupoli, Artemida, Aspropyrgos, Athens, Aigio, Chalkida, Chania, Chaidari, Chios, Cholargos, Corfu, Dhafni, Drama, Elefsina, Eleftherio - Kordelio, Elliniko, Florina, Galatsi, Giannitsa, Glyfada, Gerakas, Ilioupoli, Ioannina, Irakleio, Irakleion, Kaisariani, Kalamaria, Kalamata, Kamateron, Karditsa, Katerini, Kavala, Keratsini, Khalandrion, Kifisia, Kilkis, Komotini, Koropi, Korydallos, Kos, Kozani, Korinthos, Lamia, Livadeia, Larisa, Marousi, Melissia, Moskhaton, Mytilene, Megara, Naousa, Nea Erythraia, Nea Ionia, Nea Makri, Nea Smyrni, Oraiokastro, Orestiada, Palaio Faliro, Panorama, Peraia, Peristeri, Petroupolis, Piraeus, Polichni, Preveza, Ptolemaida, Pylaia, Patra, Perama, Pyrgos, Rethymno, Rodos, Salamina, Sparta, Serres, Thessaloniki, Thivai, Trikala, Tripoli, Volos, Voula, Vrilissia, Vari, Veroia, Xanthi, Agioi Anargyroi, Agios Athanasios, Alimos, Ano Liosia, Argos, Arta, Edessa
Orthodox Europe	Russia	Akhtubinsk, Alatyr', Aldan, Anadyr, Anna, Arzgir, Astrakhan, Aykhal, Baltiysk, Barabinsk, Belidzhi, Beloyarskiy, Bodaybo, Cherepanovo, Cherepovets, Chita, Chudovo, Dalnegorsk, Gornyak, Isil'Kul', Kandalaksha, Kartaly, Kinel'-Cherkassy, Kirov, Kizlyar, Kodinsk, Kolpashevo, Komsomolsk-On-Amur, Kondopoga, Korsakov, Kostomuksha, Kotel'Nikovo, Kotovo, Krasnokamensk, Krasnokamsk, Krasnovishersk, Kurgan, Kursk, Kyakhta, Kyzyl, Magadan, Manturovo, Mariinsk, Millerovo, Minusinsk, Mirny, Moscow, Mozhga, Nadym, Nar'Yan-Mar, Nikel, Nikolayevsk-On-Amure, Nizhnesortymskiy, Nizhneudinsk, Novoaleksandrovsk, Novozybkov, Orsk, Pechora, Pervouralsk, Raychikhinsk, Salekhard, Sasovo, Serdobsk, Severobaykal'Sk, Shimanovsk, Shuya, Sol'-Iletsk, Sovetskaya Gavan', Strezhevoy, Sukhinichi, Svetogorsk, Syktyvkar, Talnakh, Tara, Tarko-Sale, Tashtagol, Tobolsk, Tuapse, Turinsk, Tynda, Tyrnyauz, Ufa, Ukhta, Uray, Usol'Ye-Sibirskoye, Ust'-Kut, Velikiy Ustyug, Velikiye Luki, Vel'Sk, Vilyuchinsk, Vladivostok, Vyazemskiy, Vyshniy Volochek, Yagry, Yakutsk, Yarovoye, Yeniseysk, Yershov, Yeysk, Zelenogorsk
Orthodox Europe	Ukraine	Amvrosiivka, Antratsyt, Apostolove, Armyansk, Bakhmach, Balaklava, Bar, Berdyansk, Berestyn, Berezhany, Bila Tserkva, Bilopillia, Bohodukhiv, Bohuslav, Boryslav, Cherkasy, Chernihiv, Chernivtsi, Chervonohrad, Chervonopartyzansk, Chornomors'K, Chortkiv, Dniprorudne, Dolyna, Dolynska, Dzhankoi, Henichesk, Ivano-Frankivsk, Kamianets-Podilskyi, Kamyanske, Kerch, Kharkiv, Khmelnytskyi, Khmilnyk, Kiliya, Korabel, Kostopil, Kozyatyn, Kremenets, Kriukiv, Kupiansk, Kyiv, Lebedyn, Lozova, Luhansk, Lutsk, Mariupol, Melitopol, Mohyliv-Podilskyy, Mostyska, Mykolayiv, Myrhorod, Netishyn, Nizhyn, Nova Kakhovka, Novoukrayinka, Ovruch, Pavlohrad, Pereiaslav, Pervomaisk, Pervomaysk, Podilsk, Pokrovsk, Polohy, Polonne, Poltava, Pryluky, Pyatykhatky, Radomyshl, Rakhiv, Reni, Romny, Rozdilna, Sarny, Shostka, Shpola, Simferopol, Skadovsk, Sloviansk, Starobilsk, Sudak, Svatove, Tulchyn, Uman, Uzhhorod, Varash, Vilnyansk, Vinnytsya, Volnovakha, Volochysk, Volodymyr-Volynskyi, Vovchansk, Voznesensk, Vynnyky, Vynohradiv, Yasynuvata, Yevpatoriya, Zlatopil, Znamyanka, Zviahel
Protestant Europe	Germany	Aachen, Alsfeld, Ansbach, Bad Arolsen, Bad Neustadt An Der Saale, Bad Reichenhall, Bad Tolz, Bamberg, Berlin, Biberach An Der Riss, Bramsche, Brandenburg An Der Havel, Bruchkobel, Bruggen, Burg Bei Magdeburg, Burghausen, Boblingen, Cham, Cuxhaven, Dessau, Detmold, Dillingen An Der Donau, Dingolfing, Dinslaken, Eisenach, Eisenhuttenstadt, Elmschenhagen, Eltville, Emmerich, Erding, Erwitte, Flensburg, Forst, Friesoythe, Furstenwalde, Fussen, Gronau, Gorlitz, Gustrow, Heidelberg, Herborn, Hettstedt, Hilden, Idar-Oberstein, Ilmenau, Ingolstadt, Itzehoe, Karlstadt, Konstanz, Konz, Kronach, Konigslutter Am Elm, Kunzelsau, Lahr, Landsberg Am Lech, Langenhagen, Leuben, Luckenwalde, Marktredwitz, Meppen, Mittweida, Neumarkt In Der Oberpfalz, Neuruppin, Neustrelitz, Norden, Osterholz-Scharmbeck, Parchim, Passau, Petershagen, Pirmasens, Plauen, Rastatt, Ratekau, Regensburg, Ribnitz-Damgarten, Schluchtern, Schmallenberg, Schneverdingen, Schwarzenberg, Schwedt (Oder), Schwabisch Gmund, Seesen, Selm, Senftenberg, Sinzig, Sommerda, Torgau, Trossingen, Uelzen, Universitats- Und Hansestadt Greifswald, Volksdorf, Volklingen, Wangen, Weil Am Rhein, Wiehl, Wilhelmshaven, Wismar, Wittenberge, Witzenhausen, Zeitz

Continued on next page

Table 9 – continued from previous page

Region	Country	Cities
Protestant Europe	Sweden	Alingsas, Bjorlanda, Boden, Boo, Borlange, Boras, Bromma, Danderyd, Enkoping, Eskilstuna, Eslov, Falkenberg, Falkoping, Falun, Gamla Uppsala, Gustavsberg, Gavle, Halmstad, Haninge, Helsingborg, Huddinge, Hudiksvall, Huskvarna, Harnosand, Hasselby Villastad, Hassleholm, Jakobsberg, Jonkoping, Kalmar, Karlshamn, Karlskoga, Karlskrona, Karlstad, Katrineholm, Kinna, Kiruna, Kristianstad, Kristinehamn, Kumla, Kungsbacka, Kungalv, Koping, Landskrona, Lerum, Lidingo, Lidkoping, Linkoping, Ljungby, Ludvika, Lulea, Lund, Majorna, Malmo, Mariestad, Motala, Marsta, Molndal, Molnlycke, Nacka, Norrkoping, Norrtalje, Nykoping, Nassjo, Oskarshamn, Partille, Pitea, Rasunda, Sandviken, Skelleftea, Skovde, Sollentuna, Staffanstorp, Stockholm, Sundsvall, Sodertalje, Tensta, Trelleborg, Trollhattan, Tullinge, Tumba, Taby, Uddevalla, Umea, Upplands Vasby, Uppsala, Vallentuna, Varberg, Visby, Vanersborg, Varnamo, Vastervik, Vasteras, Vastra Frolunda, Vaxjo, Ystad, Angelholm, Akersberga, Orebro, Ornskoldsvik, Ostersund
Protestant Europe	The Netherlands	Aalten, Almere Stad, Amsterdam, Apeldoorn, Assen, Bergen, Bergen Op Zoom, Best, Borne, Borssele, Breda, Brummen, Cranendonck, Dalfsen, Delft, Delfzijl, Den Helder, Deventer, Doetinchem, Drachten, Dronten, Eersel, Eibergen, Elburg, Elst, Emmeloord, Emmen, Enkhuizen, Enschede, Epe, Geertruidenberg, Gennep, Goes, Goirle, Gorinchem, Hardenberg, Harlingen, Heerenveen, Helmond, Hilversum, Hoofddorp, Hoogeveen, Hoorn, Ijmuiden, Ijsselstein, Kampen, Kerkrade, Korrewegwijk, Leek, Leeuwarden, Lelystad, Leusden, Maarssen, Maassluis, Maastricht, Medemblik, Middelharnis, Mijdrecht, Noordwijk-Binnen, Oss, Oud-Beijerland, Papendrecht, Raalte, Rijssen, Roermond, Roosendaal, Rotterdam, Scheveningen, Sittard, Sneek, Stadskanaal, Staphorst, Steenwijk, Terneuzen, Tiel, Tubbergen, Urk, Veendam, Veghel, Velp, Venlo, Venray, Vlissingen, Volendam, Voorschoten, Vught, Waalwijk, Waddinxveen, Wageningen, Weesp, Wijchen, Wijk Bij Duurstede, Winschoten, Winterswijk, Woerden, Zaltbommel, Zandvoort, Zeewolde, Zeist, Zwolle
West & South Asia	Indonesia	Agats, Atambua, Balikpapan, Bangkalan, Banjar, Banjarmasin, Banyuwangi, Baturaja, Baubau, Bengkulu, Biak, Bima, Bintuni, Blangpidie, Bontang, Bungku, Cilacap, Daruba, Dobo, Dumai, Fakfak, Gorontalo, Indralaya, Indramayu, Jakarta, Jambi City, Jayapura, Kalabahi, Kalianget, Kawasi, Kedungwaru, Kendari, Klungkung, Kota Sambas, Labuan, Labuan Bajo, Lahewa, Lewoleba, Lhokseumawe, Lumajang, Luwuk, Majene, Malinau, Manggar, Manismata, Manokwari, Maros, Masohi, Maumere, Merauke, Muara Bungo, Muara Teweh, Mulia, Muntok, Nabire, Namlea, Nunukan, Padang, Padangsidempuan, Pagar Alam, Palangkaraya, Palu, Pangkalan Brandan, Pangkalpinang, Pekanbaru, Pontianak, Poso, Rantauprapat, Rantepao, Rembangan, Sabang, Sampit, Sanana, Sangkapura, Sengkang, Singkil, Sintang, Soe, Sorong, Sumbawa Besar, Sungai Penuh, Tahuna, Tambolaka, Tanjung, Tanjung Balai, Tanjung Pinang, Tanjung Selor, Tebingtinggi, Teluk Dalam, Teminabuan, Terbanggi Besar, Ternate, Timika, Tomohon, Tual, Waingapu, Waisai, Wamena, Weleri, Wonosari
West & South Asia	Israel	Acre, Afula, Arad, Ariel, Ashdod, Ashkelon, Bat Yam, Beersheba, Bet Shemesh, Bet She'An, Be'Er Ya'Aqov, Binyamina-Giv'At Ada, Bnei Brak, Baqa El Gharbiya, Daliyat Al Karmel, Dimona, Eilat, El'Ad, Er Reina, Et Taiyiba, Et Tira, Gan Yavne, Ganei Tikva, Gedera, Giv'At Shmuel, Hadera, Haifa, Herzliya, Hod Hasharon, HOlon, HUra, Jaffa, Jerusalem, Judeida Makr, Kadima Zoran, Kafr Kanna, Kafr Manda, Kafr Qari', Kafr Qasim, Karmi'El, Kfar Saba, Kfar Yona, Kiryat Gat, Kiryat Ono, Kuseifa, Lod, Maalot Tarshiha, Maghar, Majd El Kurum, Mevasseret Tsiyyon, Migdal Ha'Emeq, Modiin Ilit, Modi'In Makkabbim Re'Ut, Nahariyya, Nazareth, Nazerat 'Illit, Nesher, Ness Ziona, Netanya, Netivot, Ofaqim, Or Akiva, Or Yehuda, Pardes HAnna Karkur, Petah Tiqva, Qalansuwa, Qiryat Ata, Qiryat Bialik, Qiryat Mal'Akhi, Qiryat Shmona, Qiryat Tiv'On, Qiryat Yam, Ra'Anana, Rahat, Ramat Gan, Ramat Hasharon, Ramla, Rehovot, Rishon Letsiyyon, Rosh Ha'Ayin, Safed, Sakhnin, Sderot, Shefar'Am, Shoham, Tamra, Tel Aviv, Tel Sheva', Tiberias, Tirat Karmel, Umm El Fahm, West Jerusalem, Yavne, Yehud, Yirka, Yoqne'Am 'Illit, Yafa, Zikhron Ya'Aqov, 'Ar'Ara, 'Ar'Ara Banegev
West & South Asia	Philippines	Abuyog, Aglipay, Agoo, Al-Barka, Aparri, Asia, Atimonan, Aurora, Bangued, Bato, Bauang, Bayambang, Bignay Uno, Bislig, Bogo, Bongao, Borongan, Cabadbaran, Calabanga, Calapan, Calbayog City, Carigara, Catarman, Catbalogan, Cotabato, Culion, Dapitan, Davao, Digos, Dumaguete, Esperanza, Gingoog, Glan, Guihulngan, Iba, Ilagan, Iligan City, Iloilo, Inabanga, Ipil, Itogon, Jagna, Jolo, Jose Rizal, Kabacan, Kabankalan, Laoag, Laoang, Liloan, Lupon, Maasin, Magalang, Maganoy, Mahayag, Malungun, Mamburao, Manay, Manila, Mankayan, Manolo Fortich, Mansalay, Mapun, Mariveles, Masbate, Mercedes, Murcia, Narra, New Corella, Nia Valencia, Nunungan, Olongapo, Oroquieta, Paete, Pandan, Passi, Pinamalayan, Puerto Princesa, Recodo, Rizal, Romblon, Roxas City, Sablayan, Sagay, San Jose, Sanchez-Mira, Santa Catalina, Sitangkai, Solano, Sorsogon, Surallah, Surigao, Taal, Tabuk, Tagbilaran City, Talacogon, Tandag, Tuguegarao, Veruela, Virac, Wao