# The Anthropology of Food: How NLP can Help us Unravel the Food Cultures of the World

**Arij Riabi[1*], Sougata Saha[2], Monojit Choudhury[2]**
[1]INRIA Paris, [2]Mohamed bin Zayed University of Artificial Intelligence
[1]arij.riabi@devoteam.com, [2]{sougata.saha, monojit.choudhury}@mbzuai.ac.ae

## Abstract

Food carries cultural meaning beyond nutrition. It shapes identity, memory, and social norms, which makes it a central concern in anthropology. Given the diversity of food practices across cultures, analyzing them at scale while preserving their depth ("thick" descriptions) remains difficult for ethnographic methods, where Natural Language Processing (NLP) methods can help. Earlier NLP tools often captured only surface-level "thin" descriptions. Recent methods, especially Large Language Models (LLMs), create openings to recover cultural nuance. In this position paper, we outline research questions at the intersection of food anthropology and NLP, and discuss how LLMs can enable a scalable and culturally grounded anthropology of food. We present a case study examining what LLMs represent about global eating habits, which are often shaped by colonial histories and globalization. Our findings suggest that LLMs' internal representations recognize cultural clusters, such as shared food habits among formerly colonized regions, but fail to grasp the pragmatic and experiential aspects of food, like the worldwide spread of dishes like pizza or biryani. We conclude by highlighting some of the potential risks and gaps of using NLP for cultural analysis.

## 1 Introduction

> *"Dis-moi ce que tu manges, je te dirai ce que tu es"*– Brillat-Savarin (1842)

The famous aphorism by Jean Anthelme Brillat-Savarin links food and identity by suggesting that what we eat reflects who we are, our culture, values, and way of life. It frames eating not simply as a biological act but as an expression of the self. Anthropology has long embraced this view. Food is not merely sustenance but a symbolic system encoding identity, hierarchy, emotion, and collective memory (Fischler, 1988; Douglas, 1972; Sutton, 2010). It is studied both *as an artifact* - an object worthy of analysis in its own right - and *as a lens* to understand broader social, historical, and political processes (Mintz, 1985; Poulain, 2012). From the structure of meals that signal intimacy or exclusion (Douglas, 1972) to the ways sugar production reveals colonial power (Mintz, 1985), food functions as a cultural archive that is both deeply embodied and symbolically charged, including through the socially learned cultivation of tastes (Bourdieu, 1984). Such work exemplifies Geertz's notion of "thick description" (Geertz, 1973) of culture, which seeks to interpret meaning, context, and symbolism beyond observable behavior.

However, food, as an integral part of culture, shares many of culture's defining characteristics: it is ever-evolving, long-tailed (composed of diverse practices rather than a few dominant ones), and experiential (it is a human experience rather than mere factual knowledge) (Saha et al., 2025). Classically, anthropology relies on ethnography, a method grounded in deep, qualitative immersion within specific communities to uncover "thick" descriptions of meaning. While this approach provides rich, contextualized insights, it is inherently labor-intensive and difficult to scale or compare systematically across global populations. This makes it difficult for ethnographic approaches alone to capture the diversity and scale of food's cultural complexity, an area where computational methods, particularly NLP, can help. So far, most NLP studies of food have largely focused on recipes, sentiment, or trivia datasets (Hu et al., 2024; Palta and Rudinger, 2023; Zhou et al., 2025b), operating within what anthropology might call a "thin" descriptive paradigm. With the advent of LLMs, however, NLP systems can now engage with richer contextual and cultural dimensions of language, opening the possibility of scaling "thick" interpretation (Kommers et al., 2025). This shift presents a grow-

[*]Now at Devoteam.



*Proceedings of the First Workshop on Multilingual Multicultural Evaluation*, pages 76–98

ing opportunity for NLP to move beyond shallow text processing and toward enabling a deeper, large-scale understanding of how food encodes identity, power, and collective memory across societies.

In this paper, we argue for re-conceptualizing food in NLP along two complementary dimensions: (i) **Food as Artifact** - the subject of analysis, and (ii) **Food as Lens** - a proxy for studying identity, class, and sociohistorical change. We contrast how food has been studied in anthropology and NLP, and propose a set of research questions that emerging NLP methods, especially LLMs, can help explore more deeply. As an initial demonstration, we present a case study analyzing what LLMs implicitly represent about global daily eating habits and how these representations align with anthropological evidence on colonial histories and globalization. Overall, our contributions are:

- A systematic survey of recent NLP research on food, contrasted with anthropological accounts of food as symbolic and contextual.
- A two-dimensional conceptual framework to reorient NLP toward more culturally meaningful representations.
- A set of interdisciplinary research questions where NLP methods can help study societal concepts such as identity, behavior, status, and cultural change, using food as the lens.
- A case study demonstrating how LLMs can help frame novel anthropological hypotheses to study food and society.

## 2 A Primer on Food in Anthropology

> *To incorporate a food is, in both real and imaginary terms, to incorporate all or some of its properties: we become what we eat. Incorporation is a foundation for identity* - Fischler (1988)

Anthropology has long treated food as a central object of inquiry (as an artifact), not only for what it reveals about eating practices themselves but also for how it reflects broader cultural, social, and political processes (as a lens). In practice, most anthropological works encompass both perspectives - analyzing food as a tangible object of study and as a symbolic medium of meaning, as illustrated in Figure 1. Here, we highlight some existing works that capture these dual roles of food.

**Food, Identity, and Incorporation.** Food is viewed as a symbolic system key to sociality and identity (Mintz and Du Bois, 2002; Douglas, 1972), rather than mere sustenance. Fischler (1988)'s *principle of incorporation* frames eating as physical and symbolic assimilation: food internalizes cultural norms, affirms belonging, and marks boundaries of otherness (Grignon and Grignon, 1980).

**Food as Embodied and Everyday Practice.** Unlike music or dance, food is ubiquitous, multi-sensory, and both private and communal (Lupton, 1996; Fischler, 1988). Its taste, smell, and ritual regularity make it a vehicle for expressing identity, belonging, and exclusion (Caplan, 2013; Bourdieu, 1984). Food also encodes class, religion, gender, and power (Heldke, 1992), while carrying histories of migration, colonialism, and resistance (Pilcher, 2023; Wilk, 2006).

**Food as Social Code and Political Economy.** Douglas conceptualized meals as "codes" that signal social order through structure and participation (Douglas, 1972). Similarly, food classification systems reinforce group cohesion. At the political level, Mintz's history of sugar (Mintz, 1985) illustrates how commodities embody colonial power and economic dependency. Culinary categories (e.g., "cuisine du terroir") and gastronomic discourse also encode class, ideology, and resistance to homogenization (Nonini, 2013; Sutton, 2018).

**Food, Emotion, and Memory.** Food is deeply affective, eliciting disgust, nostalgia, comfort, and pleasure (Fischler, 1988; Lupton, 1996). Such emotions link eating to collective memory and personal identity. Proust's "Madeleine" (Proust, 1913) remains a canonical symbol of food's capacity to evoke involuntary memory. This affective power underscores food's role in cultural continuity across time and place (Holtzman, 2006).

**Beyond Symbolism: Sensory Anthropology and Embodiment.** Recent frameworks emphasize food as a lived, sensory practice rather than just symbolic code (Fox, 2003; Counihan, 2018). The TEP10 model (Sproesser et al., 2025) distinguishes what people eat (ingredients, origins) from how they eat (rituals, mealtime structures). Religious traditions, such as Catholic culinary discourse, frame pleasure and embodiment as integral to communal and spiritual life (Poulain, 2012). Sensory anthropology thus foregrounds food as a site where culture is enacted, felt, and moralized (Howes, 2010; Korsmeyer and Sutton, 2011).

**Food as a Cultural Lens for Computation.** Food, in Newmark's typology (Newmark, 2003), cuts across material culture, social organization, norms, and habits. It can be studied as a material object (tools, recipes) (Gremillion, 2011; Dalby, 2003; Alcock, 2005; Sutton, 2010; Staller, 2003) and as a lens onto ritual practice, hierarchy, and everyday behavior. Traditional anthropology offers thick, small-scale insights but struggles to compare cultures systematically. NLP, by contrast, can scale analyses of cookbooks, menus, and food discourse, revealing diachronic patterns, identity markers, and implicit norms. Embedding-based methods can uncover associations between food and identity, while LLMs offer a way to complement thick description with cross-cultural breadth (Kommers et al., 2025).

Collectively, such applications point to a growing space where NLP can enhance anthropology's capacity to analyze culture at both depth and scale.

# 3 Food for Cultural Analytics in NLP

To understand how food has been treated across NLP studies, we compiled a curated corpus by systematically gathering papers from multiple sources.

## 3.1 Survey Methodology

We focused on major NLP venues: ACL, EMNLP, EACL, and NAACL, conducting a systematic search across titles, abstracts, and full texts. We used a Boolean combination of the following keywords "cultural adaptation," or "culture" and "food," "cuisine," "recipe," "gastronomy," and "identity." We deliberately prioritize *CL venues to ensure methodological rigor, reproducibility, and alignment with the NLP community's standards for peer-reviewed research. To broaden our scope, we also searched Google Scholar[1] and Semantic Scholar[2] for interdisciplinary work that integrates perspectives from anthropology, food studies, and cultural sociology, using the keywords "food and identity," "food culture," and "food as social practice." We excluded works that focus only on nutrition or agriculture, resulting in 21 studies (Figure 1).

Although we did not restrict our search by publication date, we observed that most of the 21 publications were within the last three years (2023-2025), highlighting the recent emergence and growing momentum of interdisciplinary research at the intersection of NLP and culture. In a related survey on culture and NLP, Adilazuarda et al. (2024) analyzed 90 papers published between 2020 and 2024 and noted a clear shift in NLP research toward cultural awareness, particularly following the rise of LLMs. Liu et al. (2025) similarly highlighted this trend, showing an increasing emphasis on cultural adaptation and sensitivity across various NLP tasks. We observe this shift mirrored in the domain of food and NLP as well.

## 3.2 Survey Findings

**Food as Artifact, Not Lens.** Despite food's deep cultural resonance, NLP research has often reduced it to a static artifact. Many studies use food data to test factual recall, translation accuracy, or retrieval performance without engaging with its social meaning. For example, Hu et al. (2024) developed a cross-cultural recipe retrieval dataset, while Cao et al. (2024) focused on Chinese-English recipe translation. Both highlight technical challenges but remain centered on correctness, ingredient substitution, or retrieval accuracy. Similarly, multimodal work (Fu et al., 2020; Bagler and Singh, 2018) shows that while models recognize food semantics across text and images, they fail to account for cultural context. For instance, two dishes may look visually similar but carry entirely distinct ritual or regional significance.

**Food as Diagnostic Probe.** Several works use food as a diagnostic probe to reveal cultural bias or commonsense gaps in models. FORK (Palta and Rudinger, 2023) surfaces US-centric assumptions in commonsense QA through culinary questions. Zhou et al. (2025b) introduce MAPO, a multilingual dataset for ingredient prediction, which demonstrates how models stereotype or overgeneralize cultural food knowledge. While valuable, these approaches primarily use food as factual tests, rather than as a lens into lived cultural practice.

**Food in Cultural Evaluation Frameworks.** Several benchmarks incorporate food into large-scale cultural evaluation, but often remain at a surface level. Winata et al. (2025a) developed visual QA tasks focused on dish recognition and cuisine origin, while Pellegrini et al. (2021) highlights the subjective and context-dependent challenges of modeling culinary substitutions. Counterfactual recipe generation by Liu et al. (2022) advances this line by probing compositional generalization, showing how models struggle to simulate stylistic or culturally specific preparation methods. To-

[1] https://scholar.google.com/
[2] https://www.semanticscholar.org/

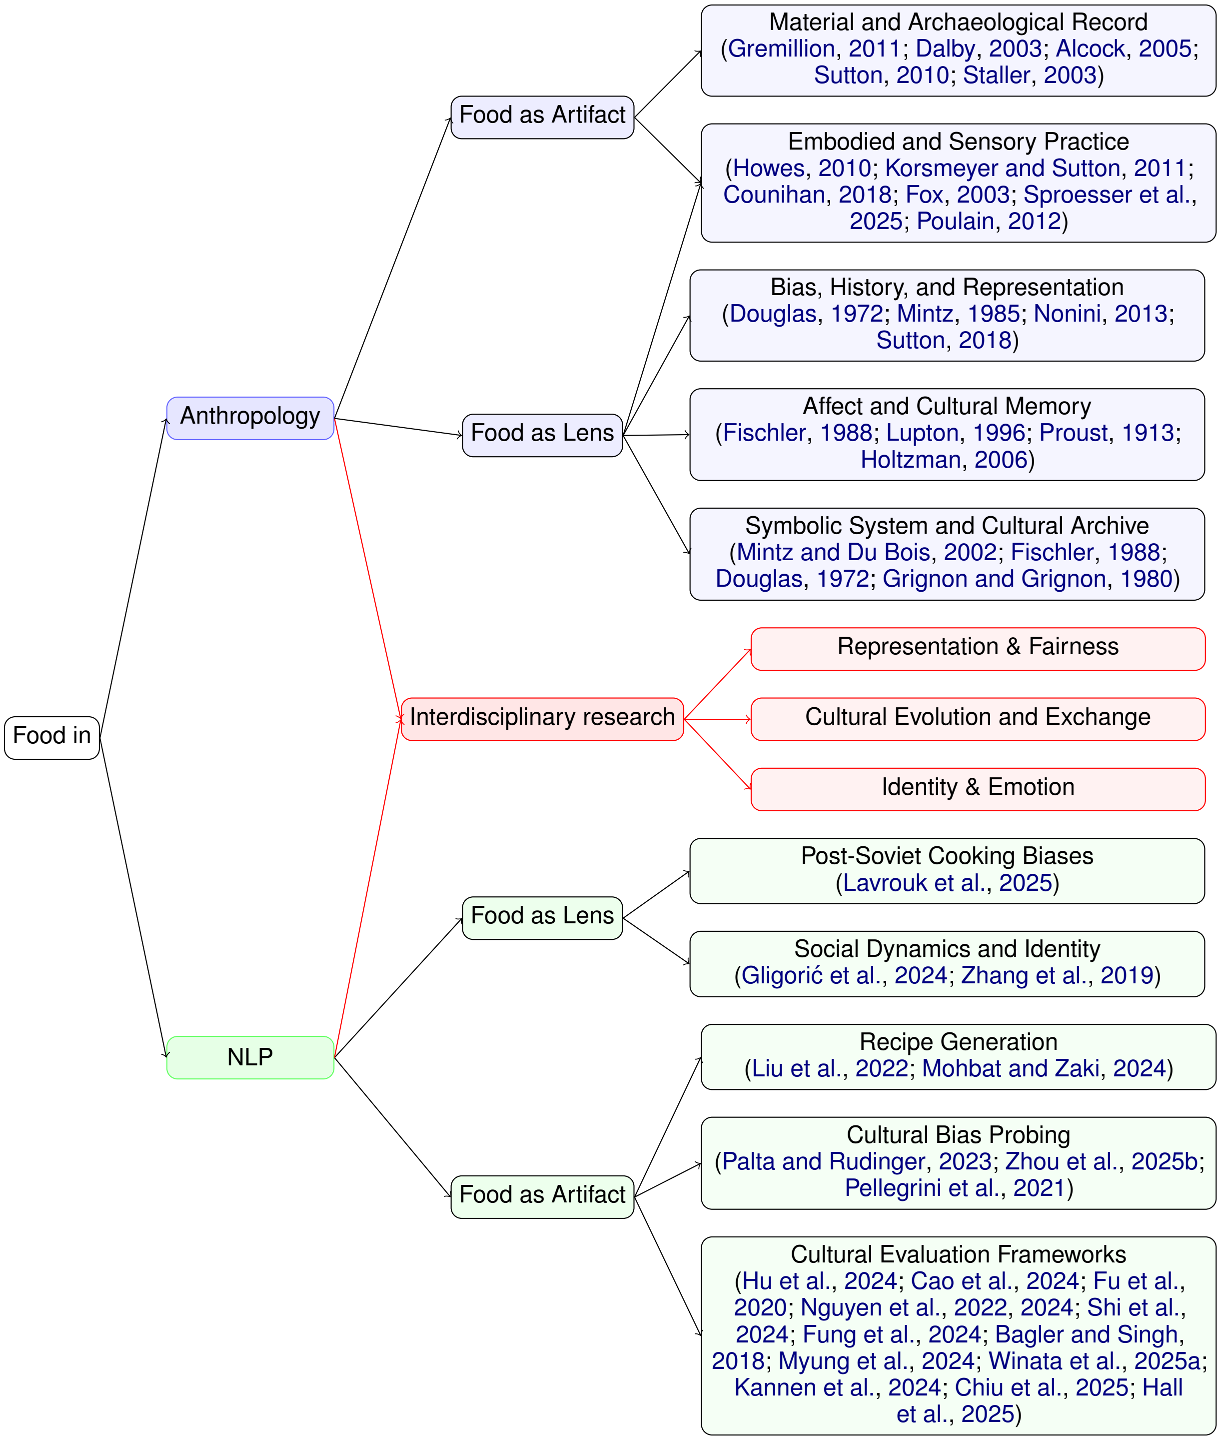


Figure 1: Taxonomy of the survey mapping research on food in NLP (highlighted in green) and anthropology (in violet), contrasting artifact-oriented with lens-based studies. The taxonomy illustrates our call for an interdisciplinary framework (in red) bridging computational and anthropological approaches.

gether, these studies highlight food's potential as a domain for evaluating procedural knowledge and cultural embodiment, yet they stop short of engaging deeper cultural theorization.

Food also consistently appears as an important category in cultural evaluation benchmarks (Nguyen et al., 2022, 2024; Shi et al., 2024; Fung et al., 2024; Myung et al., 2024; Chiu et al., 2025). Strikingly, Fung et al. (2024) shows that LLMs perform worse on cuisine-related reasoning than on domains such as education or holidays, since food often requires fine-grained, localized knowledge that is less documented, particularly for underrepresented communities. This illustrates the gap between models' generalized representations and the situated specificity of food practices.

**Food as Social and Behavioral Practice.** Few works treat food as a social practice. Gligorić et al. (2024) studies food purchase mimicry on a university campus, showing how choices reflect social alignment and symbolic inclusion. Their findings resonate with Fischler (1988)'s claim that eating is not just ingestion but assimilation of cultural values. Such perspectives position food as a *behavioral proxy* for social belonging, suggesting opportunities for NLP to model culture as practice, performance, and interaction.

**Bias, History, and Representation.** Recent work also highlights how LLMs reproduce cultural biases in food knowledge. Lavrouk et al. (2025) demonstrate that models misattribute dish origins, particularly in multilingual contexts. For instance, Russian prompts lead to systematic over-association with Russian cuisine, reflecting how linguistic dominance shapes cultural representation in LLMs. Zhang et al. (2019), focusing on food images, caution against assuming national datasets as mono-cultural. Their study links visual aesthetics to user engagement (likes, bookmarks, ratings) and calls for interdisciplinary collaboration with anthropologists to interpret cultural meaning. Together, these works underscore the socio-historical and representational dimensions of food in NLP.

**Parallel Streams in Nutrition and Health.** We also acknowledge parallel NLP research that studies food's nutrition and health aspects. Such studies use NLP to improve dietary outcomes and environmental sustainability (McCarthy, 2025), classify foods, recipes and assess diets (Zhou et al., 2025a; Rezayi et al., 2025), and extract dietary patterns (Choi et al., 2022). While these works are valuable for health and industry, since they treat food as a technical object and not as a cultural or social practice, they fall outside the scope of this paper.

### 3.3 Research on Culture and LLMs

Adilazuarda et al. (2024) survey how "culture" is represented and measured in LLMs, showing that most studies rely on abstract or *demographic proxies* and primarily test for factual correctness - the "thin" description of culture - leaving its experiential and contextual dimensions underexplored. We extend this discussion by focusing on food, an integral and embodied aspect of culture. Many existing approaches, often using prompting techniques such as *socio-demographic prompting* (Li et al., 2024; AlKhamissi et al., 2024; Wan et al., 2023), instruct LLMs to assume diverse cultural identities and exhibit behaviors representative of an average person from the assumed culture. A model is then considered culturally aware if its generated behavior aligns with expected responses (Kharchenko et al., 2024; LI et al., 2024; Dawson et al., 2024). However, several works emphasize that culture should not be reduced to static facts or trivia but understood as a dynamic, situated, and context-dependent practice, which LLMs can help understand at scale (Zhou et al., 2025c; Kommers et al., 2025). Similarly, Saha et al. (2025) advocates for developing *meta-cultural competent* models that adaptively engage with unfamiliar cultural contexts instead of relying solely on static data.

Overall, NLP research has mostly examined food as an *artifact*, a *diagnostic probe*, or a *benchmark category*. Building on insights from anthropology and the social sciences, we propose that NLP can also treat food as a *lens*, to study broader cultural patterns, meanings, and social relations.

## 4 Towards Cross-Disciplinary Research

NLP methods, particularly LLMs, enable cultural analysis at a scale that traditional anthropological and social science approaches cannot easily achieve (Kommers et al., 2025). Bridging these fields, we outline ten research questions along three broad dimensions (illustrated in Figure 1) through which LLMs can uncover global food patterns, empirically revisit theories in structural anthropology (Strauss, 1974; Hénaff, 1998) and computational social science (Lazer et al., 2009; Goel and Bagler, 2022; Ahnert, 2013; Edelmann et al., 2020).

### 4.1 Identity and Emotion

**How does food-related discourse reflect emerging identities & groups?** Food practices signal inclusion or exclusion, shaping groups and collective identities (e.g., veganism, nationalism). NLP can track how communities construct belonging through hashtags, reviews, or recipe-sharing forums, revealing negotiated identity markers.

**How do food preferences reflect lifestyle, value systems, or behavioral patterns?** Shifts from traditional to modern food practices mirror broader societal change (Sproesser et al., 2022). NLP can analyze online posts and blogs to capture how food-related language reflects urbanization, global influences, individual lifestyle choices (e.g., clean eating, comfort food), and evolving cultural identity.

**How are emotions and affect encoded in food-related language across cultures?** Food evokes strong emotions such as nostalgia, pride, disgust, etc, which shape identity and memory (Lupton, 1996; Holtzman, 2006). Sentiment analysis, emotion tagging, and metaphor mining can uncover cross-cultural variation and support more culturally aware emotion detection systems.

### 4.2 Cultural Evolution and Exchange

**How do food narratives reflect postcolonial influences & cultural exchange?** Food discourse often carries traces of colonial histories through ingredients, preparation methods, and naming conventions. Analyzing large-scale recipe corpora or food blogs can uncover how colonial legacies persist or evolve, revealing implicit power asymmetries and cultural adaptations. For example, words like "tacos," "kimchi," or "couscous" may appear in different cultural embeddings, traditional or modern eating, depending on regional context and temporal framing. This variation can be captured through dynamic embedding models or temporal word sense induction, offering insights into how culinary terms shift meaning and connotation over time and across communities (Sproesser et al., 2025). Our case study further illustrates how constructing country-country similarity networks from LLM probabilities enables systematic testing of anthropological hypotheses, such as the role of colonial histories or globalization in shaping food practices, on a scale far beyond traditional ethnographies.

**How have recipes evolved over time under cultural, economic, and environmental pressures?** Recipes are living cultural artifacts, continuously adapted across time and space. Diachronic NLP analysis of historical cookbooks and recipe websites can detect ingredient substitutions, preparation shifts, and linguistic trends that reflect socio-economic and ecological change.

**How are food rituals and communal practices reflected in text, and what can they reveal about social cohesion?** Food rituals reinforce collective norms. NLP can analyze ceremonial menus, festival descriptions, or religious texts to uncover recurring structures of ritualistic eating (e.g., Passover, Eid, Onam feasts, etc), which reflect temporality and shared group values.

**How does food-related discourse reflect emerging cultural shifts around health, sustainability, and environmental consciousness?** Language around plant-based diets, organic food, zero-waste cooking, etc, reflects evolving cultural values tied to global health and environmental movements. NLP methods can analyze trends in social media, blogs, or menus to trace how these discourses spread, who participates in them, and how they intersect with local traditions or resistance. This offers insight into how cultural norms around food are negotiated and transformed in response to global challenges, which can in turn inform policy or behavioral interventions.

**How can LLMs simulate social dynamics in food discourse (e.g., norm shifts, dietary trends)?** Recent LLM advancements have demonstrated their potential as *world models* (Shavit et al., 2023; Acharya et al., 2025), indicating the possibility of simulating human behavior at scale. Studies have shown that LLMs can approximate social behavior in multi-agent environments and generative simulations (Park et al., 2023; Zhou et al., 2023; Wang et al., 2024; Zhou et al., 2024). Although these models are limited in scope and lack deep generalization (Saha and Choudhury, 2025), they point towards a promising direction: the potential to simulate cultural dynamics, including food-related behaviors. LLMs can be fine-tuned or prompted to model hypothetical cultural changes (e.g., how food taboos emerge/fade). Although these simulations cannot replace anthropological expertise, they can support hypothesis testing by letting researchers study sociolinguistic drift, normative change, and symbolic meaning at scale.

### 4.3 Linguistic Variation and Representation

Food names reflect hierarchies and regional identities. For example, the multiple ways of expressing "breakfast" in Malay, Mandarin, and Tamil, such as "morning rice" or "layering the belly", illustrate how meals are culturally framed rather than just names, highlighting challenges of translation ambiguity in NLP (Poulain et al., 2023).

**How do dialects and local food terms affect model understanding and cultural representation?** Ingredient names and local terminology often encode social or regional identity as culture plays a role in shaping food semantics (Mazzuca and Majid, 2023). Studying how LLMs interpret or generate these can reveal representational biases and inform improvements in cultural adaptation. For example, ingredient names ("bhindi" vs "okra", "aubergine" vs "eggplant", etc) often vary by region and carry cultural significance. NLP can surface biases, mismatches, or translation failures in handling this diversity. This could also enable studying the dimensions along which cultures vary: aboutness, common ground, and linguistic form and style (Hershcovich et al., 2022). For example, analyzing Google reviews across different price ranges and cuisine categories - from fine dining to fast food - could help uncover variations in language registers, dialect usage, and cultural framing.

**How can NLP mitigate representational disparities for underrepresented cuisines and communities?** Food practices often reflect regional and cultural identities that do not neatly align with nation-states. For example, while countries like China, India, Mexico, and Brazil show strong regional–national alignment (Sproesser et al., 2022), the U.S. exhibits distinct ethnic culinary clusters (e.g., African American, Latin American) that transcend a single national food identity and may even align with transnational groups. These patterns highlight the need for NLP models to capture intra-national diversity rather than reducing multiethnic cuisines to monolithic national categories.

## 5 Case Study

To illustrate how LLMs can enrich our understanding of food as a cultural lens, we present a case study examining how textual representations of food reflect patterns of identity, belonging, and social change. Anthropological work shows that both colonization and geographic proximity shape eating habits (Douglas, 1972; Mintz, 1985; Fischler, 1988). Yet in today's globalized world, dishes often transcend their place of origin. For instance, in 2024, DoorDash, a widely used food-delivery app in the USA and Canada, reported Indian dishes such as naan, butter chicken, and biryani among the most ordered items in the region (sources in Table 1). Food ordering platforms such as Talabat and Zomato reported pizza, an Italian dish, as a staple in Oman, the UAE, and India. We ask two questions inspired by anthropology: (Q1) **What country-level clusters of food practices do LLMs produce, and do these clusters reflect colonial histories?** (Q2) How connected are these clusters in light of contemporary globalization - **do LLMs mirror lived cultural practices, or do they reproduce more static, fact-based representations?**

### 5.1 Setup

Several studies show that online food-ordering behavior reflects a combination of familiarity and convenience, both of which are closely tied to cultural exposure and dietary norms (Rozin, 2005; Warde, 2016; Zhang et al., 2025; Keeble et al., 2020). While convenience influences ordering frequency, familiarity largely determines what individuals feel comfortable consuming or associating with their identity. Thus, what people choose to order online serves as a meaningful behavioral proxy for cultural familiarity, particularly at the dish level, where preferences reflect social and cultural embeddedness rather than mere accessibility.

To answer the two questions, we adapted the Food Choice Questionnaire (FCQ) (Steptoe et al., 1995), focusing on two dimensions: **convenience** (foods easily available nearby) and **familiarity** (familiar foods). We prompted two instruction-tuned LLMs - Llama-3.1-8B-Instruct (Dubey et al., 2024) and Gemma-2-9B-it (Team et al., 2024) - across 170 countries, and observed their log probabilities for tokens corresponding to a curated list of approximately 3,500 food items. For each model and dimension, we constructed a country$\times$food matrix $M$ by collating the token log-probabilities (converted to probabilities using softmax), and computed a country-country similarity matrix $S_{cc} = MM^{\top}$. We detected communities by constructing a weighted mutual $k$-NN graph on $S_{cc}$ with $k = 10$, and applied the Leiden community detection algorithm (Traag et al., 2018), which optimizes modularity $Q$ (Newman and Girvan, 2004), a measure that compares intra-cluster connectivity with a random baseline. Further details in Section A.1.

### 5.2 Results and Discussion

To answer Q1, we compared the clusters against a curated mapping between countries and their historical colonizers (details in Section A.1) and used Fisher's exact test (Fisher, 1922) to evaluate whether LLM-derived food-based communities aligned with colonial histories within continents. Across both models and questions, the country-country graphs have high modularity ($Q \in [0.79, 0.83]$) with 10-11 communities and far more intra than inter-community edges (detailed statistics in Table 2). **The clusters significantly align**

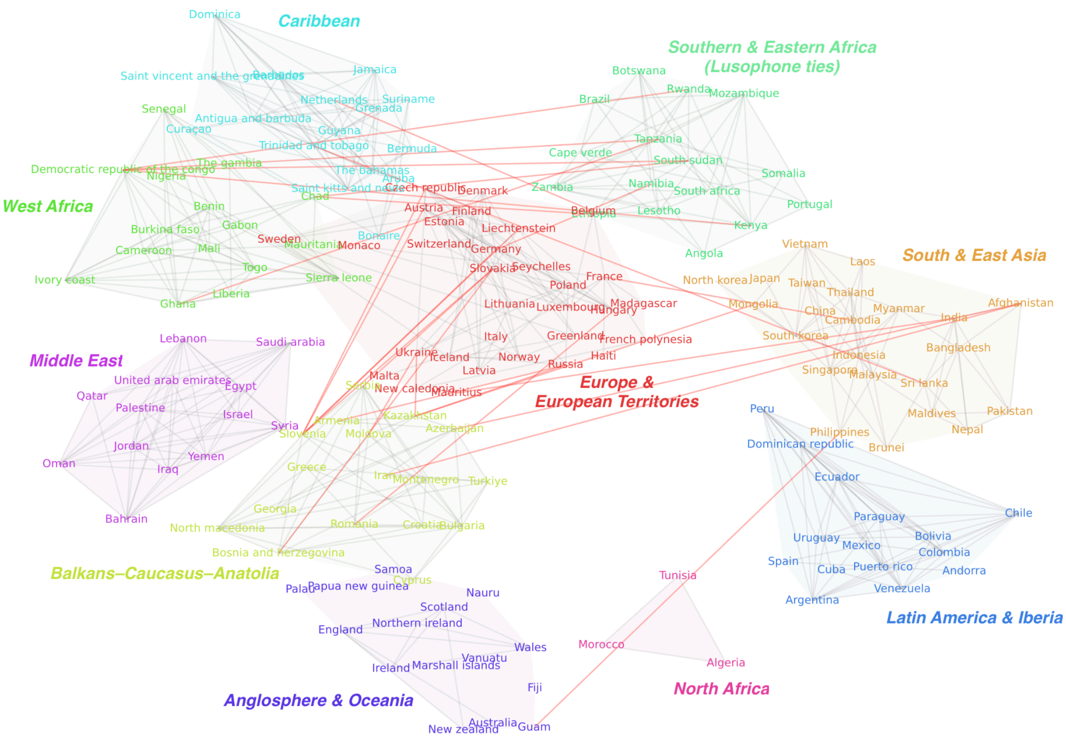


Figure 2: **Country communities from Llama (convenience).** Weighted mutual-$k$NN graph ($k = 10$) on $S_{cc} = MM^{\top}$; Leiden partition on modularity. High modularity ($Q = 0.829$) with coherent regional clusters. Intra-community edges shown in gray, inter-community edges in red. Community details in Table 3. **Labels of the communities are heuristic; they reflect the majority of countries in each cluster.** Node colors are used solely for visual separation of clusters and do not carry additional semantic meaning.

**with colonial histories**: Fisher tests show higher odds ratios (OR) that countries sharing a colonizer fall in the same community (Llama: OR = 3.4 for convenience, 2.6 for familiarity; Gemma: 3.7 and 2.7). Map-based visualisations confirm tight regional blocks and few cross-block ties (Fig 6 - 9).

Figure 2 illustrates Llama's convenience-based clusters, where countries were grouped into 10 communities. Interestingly, **we observe some countries occupying *brokerage positions*, acting as inter-community bridges**. From a structural perspective, such bridging roles are crucial for the diffusion of ideas, practices, and goods across cultural boundaries, as noted in classical network studies of innovation and knowledge transfer (Gould, 1993; Burt, 2004). While this remains a hypothesis that requires validation with real-world food exchange and migration patterns, our results suggest that LLMs capture relational structures that resonate with known mechanisms of cultural flow.

Beyond colonial clustering, Figure 2 also reveals **striking disconnections and over-connections**. For example, Morocco, Tunisia, and Algeria, commonly grouped as the Maghreb, appear fragmented, while Arab countries are split into disconnected sub-clusters. **Europe shows dense interconnections**, plausibly reflecting its overrepresentation in textual corpora, while **Latin America forms a cluster with few external ties**, despite strong historical and migratory links to Europe. These anomalies suggest hypotheses: that LLM-derived clusters overemphasize colonial and linguistic segmentation, underrepresent globalization and migration, and reproduce biases in training data.

Some observed patterns may reflect artifacts of the training corpora rather than real-world dynamics. For instance, brokerage roles could shift if we examined an Arabic LLM trained on more region-specific data. This underscores both the promise and the limits of LLMs: while they encode patterns of cultural knowledge, these patterns may be shaped as much by training data distributions as by real-world dynamics. Nonetheless, analyzing such structures can generate testable hypotheses for validation against migration records, trade flows, or ethnographic accounts, thereby opening pathways for computational methods to complement anthropological research. We observe similar insights for both dimensions and models (Figures 4, 3, 5).

To address Q2, we examined food ordering data (2021-2024) from DoorDash, Talabat, and Zomato for Australia, Canada, India, New Zealand, Oman, the UAE, and the USA. Table 1 shows that **pizza is a global staple** (e.g., Australia, India, UAE), while **Indian dishes such as naan, butter chicken, and biryani are frequently ordered** in Canada, New Zealand, India, and the USA, alongside local favorites like Wagyu beef dishes and California rolls. These trends highlight how globalization reshapes lived food practices beyond colonial and historical divides. The persistence of colonial clustering in LLM-derived communities suggests that models draw on historical categories embedded in training data as the basis for cultural organization. However, **sparse inter-cluster connectivity shows that models underrepresent transnational food flows** (migration-driven exchanges, globalized staples, etc) despite their prominence in the real world, reflecting the difference between knowledge in training data and the fluid practices of everyday eating.

Visualizing these clusters geographically reveals that countries and their food communities appear largely disconnected, highlighting the need for anthropologically grounded LLM evaluations that capture culture as lived and dynamic rather than merely factual artifacts (see Figure 6). Additional visualizations across models and dimensions further confirm these patterns of regional cohesion and sparse global connectivity (Figures 8,7,9).

## 6 Discussion and Conclusion

Although our proposed re-framing opens interdisciplinary possibilities, it also raises practical and conceptual challenges, a few of which are below:

**1. Do NLP systems need to engage with symbolic and experiential aspects of food?** One might argue that tasks aligned with factual correctness, such as recognizing dishes, generating recipes, or identifying cuisine origins, are sufficient for models. However, if NLP aspires to develop models that truly "understand" culture, it must capture the symbolic and experiential dimensions of food. Current benchmarks that treat cultural knowledge as static trivia fail to evaluate these deeper capabilities; moving from food as artifact to lens is thus a necessary step for robust cultural evaluation of LLMs. Computational approaches simulating societies (Epstein, 2012; Shavit et al., 2023) further emphasize the need to model the interconnectedness of food and society, beyond mere factual recall.

**2. Are computational methods a replacement for cultural expertise?** Culturally adept NLP models are not substitutes for anthropologists, sociologists, or cultural historians. Instead, they serve as complementary tools, enabling empirical, scalable analysis (Newell and Simon, 2007; Polak, 2016), uncovering implicit cultural patterns, and testing hypotheses across large corpora (Kozlowski et al., 2019; Arseniev-Koehler and Foster, 2022), as further demonstrated by our case study. Nonetheless, careful reflection on methodological and epistemological limits is essential to avoid over-claiming what models can "know" about food and culture.

**3. Safety, stereotyping, and cultural bias.** Treating food as a cultural lens introduces ethical concerns. NLP models can perpetuate stereotypes, essentialize cuisines, or reinforce colonial hierarchies, reflecting biases in training data (Sheng et al., 2019; Gehman et al., 2020; Nadeem et al., 2021; Bender et al., 2021). In applications like virtual assistants, recipe recommenders, or educational tools, these biases can shape user perceptions, for example, by associating certain cuisines with poverty, hygiene, or spiciness. Western-centric framing in prior studies also marginalizes non-Western culinary traditions (Palta and Rudinger, 2023; Cao et al., 2024; Zhou et al., 2025b). Addressing these issues requires culturally diverse datasets, transparent documentation, and participatory design.

**4. Integrating cultural reflexivity and relativism.** Anthropology stresses reflecting on one's own cultural assumptions ("cultural reflexivity") and understanding practices in their own contexts ("cultural relativism") (Davies, 2012; Bourdieu, 2004; Deer, 2014; Boas, 1941). For example, assumptions linking political hierarchy to elaborate cuisines are not universal, as shown by Gellner's ethnography in Africa (Gellner, 1985). Similarly, NLP can benefit from reflexive and relativist principles, promoting models that account for meaning as situated and context-dependent, leading to more culturally competent and inclusive systems.

In conclusion, food is a rich cultural artifact that encodes values, identities, and social practices, offering NLP a unique opportunity to contribute to a deeper, anthropological understanding of culture through advanced computational methods such as LLMs. To advance this vision, we outline research directions at the intersection of food studies and NLP and illustrate their potential through an exploratory case study, and discuss their challenges.

## Limitations

While this paper highlights new interdisciplinary directions, several limitations must be acknowledged. **First**, our survey of NLP and mainly anthropology literature is necessarily selective and focuses primarily on English-language sources and major conferences. There is extensive anthropological work on food practices in non-Western contexts and diverse languages that we do not fully capture here. **Second**, many of the NLP studies we reviewed rely on digital corpora (e.g., online reviews, social media, cookbooks), which may over-represent urban, affluent, or globally connected communities while under-representing rural or marginalized voices. Food and its encompassing rituals and social dynamics are deeper, which digital sources might fail to capture. **Third**, our focus on symbolic and social dimensions of food intentionally excludes nutritional, agricultural, and product-level applications, which are important but outside our scope. Future work might consider how these perspectives intersect, particularly when addressing health disparities, fitness, and food security. **Fourth**, our case study, limited to two models and a small set of questions, is not intended as an exhaustive account of food-related patterns. Rather, it illustrates the feasibility of using LLMs for knowledge and pattern discovery through a network analysis approach, while acknowledging that other methods could equally be applied. Moreover, our analysis primarily used English food item names; incorporating local terms may reveal different cultural patterns, an important direction for future validation. **Finally**, the three dimensions proposed in Section 4 - "identity and emotion", "cultural evolution and exchange", and "linguistic variation and representation" - broadly capture most of the pertinent research directions at the intersection of food anthropology and NLP. However, the ten research questions themselves are not exhaustive. More questions are possible that belong to the three dimensions.

## Acknowledgements

This research was supported by Microsoft Accelerate Foundation Models Research (AFMR) Grant. Arij Riabi was partly funded by the BPI "Scribe" project.

## References

Deepak Bhaskar Acharya, Karthigeyan Kuppan, and B Divya. 2025. Agentic ai: Autonomous intelligence for complex goals–a comprehensive survey. *IEEE Access*.

Muhammad Adilazuarda, Sagnik Mukherjee, Pradhyumna Lavania, Siddhant Singh, Alham Aji, Jacki O'Neill, Ashutosh Modi, and Monojit Choudhury. 2024. Towards measuring and modeling "culture" in llms: A survey. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 15763–15784.

Sebastian E Ahnert. 2013. Network analysis and data mining in food science: the emergence of computational gastronomy. *Flavour*, 2:1–3.

Joan P Alcock. 2005. *Food in the ancient world*. Bloomsbury Publishing USA.

Badr AlKhamissi, Muhammad ElNokrashy, Mai Alkhamissi, and Mona Diab. 2024. Investigating cultural alignment of large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12404–12422, Bangkok, Thailand. Association for Computational Linguistics.

Alina Arseniev-Koehler and Jacob G Foster. 2022. Machine learning as a model for cultural learning: Teaching an algorithm what it means to be fat. *Sociological Methods & Research*, 51(4):1484–1539.

Ganesh Bagler and Navjot Singh. 2018. Data-driven investigations of culinary patterns in traditional recipes across the world. In *2018 IEEE 34th International Conference on Data Engineering Workshops (ICDEW)*, pages 157–162.

Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.

Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008.

Franz Boas. 1941. Race, language and culture. *The Journal of Nervous and Mental Disease*, 94(4):513–514.

Pierre Bourdieu. 1984. *Distinction: A Social Critique of the Judgement of Taste*. Harvard University Press.

Pierre Bourdieu. 2004. *Science of science and reflexivity*. Polity.

Jean Anthelme Brillat-Savarin. 1842. *Physiologie du Gout, ou, Méditations de gastronomie transcendante...* Charpentier.

Ronald S Burt. 2004. Structural holes and good ideas. *American journal of sociology*, 110(2):349–399.

Yong Cao, Yova Kementchedjhieva, Ruixiang Cui, Antonia Karamolegkou, Li Zhou, Megan Dare, Lucia Donatelli, and Daniel Hershcovich. 2024. Cultural adaptation of recipes. *Transactions of the Association for Computational Linguistics*, 12:80–99.

Pat Caplan. 2013. *Food, health and identity*. Routledge.

Yu Ying Chiu, Liwei Jiang, Bill Yuchen Lin, Chan Young Park, Shuyue Stella Li, Sahithya Ravi, Mehar Bhatia, Maria Antoniak, Yulia Tsvetkov, Vered Shwartz, and Yejin Choi. 2025. CulturalBench: A robust, diverse and challenging benchmark for measuring LMs' cultural knowledge through human-AI red-teaming. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 25663–25701, Vienna, Austria. Association for Computational Linguistics.

Insu Choi, Jihye Kim, and Woo Chang Kim. 2022. Dietary pattern extraction using natural language processing techniques. *Frontiers in Nutrition*, Volume 9 - 2022.

Carole M Counihan. 2018. *The anthropology of food and body: Gender, meaning and power*. Routledge.

Andrew Dalby. 2003. *Food in the Ancient World from A to Z*. Psychology Press.

Charlotte Aull Davies. 2012. *Reflexive ethnography: A guide to researching selves and others*. Routledge.

Fiifi Dawson, Zainab Mosunmola, Sahil Pocker, Raj Abhijit Dandekar, Rajat Dandekar, and Sreedath Panat. 2024. Evaluating cultural awareness of llms for yoruba, malayalam, and english. *Preprint*, arXiv:2410.01811.

Cécile Deer. 2014. Reflexivity. In *Pierre Bourdieu*, pages 195–208. Routledge.

Mary Douglas. 1972. Deciphering a meal. In *Implicit meanings: Essays in anthropology*, pages 249–275. Routledge.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Achim Edelmann, Tom Wolff, Danielle Montagne, and Christopher A Bail. 2020. Computational social science and sociology. *Annual review of sociology*, 46(1):61–81.

Joshua M Epstein. 2012. *Generative social science: Studies in agent-based computational modeling*. Princeton University Press.

Claude Fischler. 1988. Food, self and identity. *Social Science Information*, 27(2):275–292.

Ronald A Fisher. 1922. On the interpretation of $\chi$ 2 from contingency tables, and the calculation of p. *Journal of the royal statistical society*, 85(1):87–94.

Robin Fox. 2003. Food and eating: an anthropological perspective. *Social Issues Research Centre*, 2003:1–21.

Han Fu, Rui jin Wu, Chenghao Liu, and Jianling Sun. 2020. Mcen: Bridging cross-modal gap between cooking recipes and dish images with latent variable model. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14558–14568.

Yi Fung, Ruining Zhao, Jae Doo, Chenkai Sun, and Heng Ji. 2024. Massively multi-cultural knowledge acquisition & lm benchmarking. *Preprint*, arXiv:2402.09369.

Clifford Geertz. 1973. *The Interpretation of Cultures*. Basic Books, New York.

Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. 2020. RealToxicityPrompts: Evaluating neural toxic degeneration in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3356–3369, Online. Association for Computational Linguistics.

Ernest Gellner. 1985. No haute cuisine in africa. *Relativism and the Social Sciences*.

Kristina Gligorić, Arnaud Chiolero, Emre Kıcıman, Ryen W White, Eric Horvitz, and Robert West. 2024. Food choice mimicry on a large university campus. *PNAS Nexus*, 3(12):pgae517.

Mansi Goel and Ganesh Bagler. 2022. Computational gastronomy: A data science approach to food. *Journal of Biosciences*, 47(1):12.

Roger V Gould. 1993. Collective action and network structure. *American sociological review*, pages 182–196.

Kristen J Gremillion. 2011. *Ancestral appetites: food in prehistory*. Cambridge University Press.

Claude Grignon and Christiane Grignon. 1980. Styles d'alimentation et goûts populaires. *Revue française de sociologie*, pages 531–569.

Siobhan Mackenzie Hall, Samantha Dalal, Raesetje Sefala, Foutse Yuehgoh, Aisha Alaagib, Imane Hamzaoui, Shu Ishida, Jabez Magomere, Lauren Crais, Aya Salama, and Tejumade Afonja. 2025. The human labour of data work: Capturing cultural diversity through world wide dishes. *Proc. ACM Hum.-Comput. Interact.*, 9(7).

Lisa M Heldke. 1992. Food politics, political food. *Cooking, eating, thinking: transformative philosophies of food*, pages 301–327.

Marcel Hénaff. 1998. *Claude Levi-Strauss and the making of structural anthropology*. U of Minnesota Press.

Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza Fierro, Katerina Margatina, Phillip Rust, and Anders Søgaard. 2022. Challenges and strategies in cross-cultural NLP. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6997–7013, Dublin, Ireland. Association for Computational Linguistics.

Jon D. Holtzman. 2006. Food and memory. *Annual Review of Anthropology*, 35(Volume 35, 2006):361–378.

David Howes. 2010. *Sensual relations: Engaging the senses in culture and social theory*. University of Michigan Press.

Tianyi Hu, Maria Maistro, and Daniel Hershcovich. 2024. Bridging cultures in the kitchen: A framework and benchmark for cross-cultural recipe retrieval. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1068–1080, Miami, Florida, USA. Association for Computational Linguistics.

Nithish Kannen, Arif Ahmad, Marco Andreetto, Vinodkumar Prabhakaran, Utsav Prabhu, Adji Bousso Dieng, Pushpak Bhattacharyya, and Shachi Dave. 2024. Beyond aesthetics: Cultural competence in text-to-image models. In *Advances in Neural Information Processing Systems*, volume 37, pages 13716–13747. Curran Associates, Inc.

Matthew Keeble, Jean Adams, Gary Sacks, Lana Vanderlee, Christine M White, David Hammond, and Thomas Burgoine. 2020. Use of online food delivery services to order food prepared away-from-home and associated sociodemographic characteristics: a cross-sectional, multi-country analysis. *International journal of environmental research and public health*, 17(14):5190.

Julia Kharchenko, Tanya Roosta, Aman Chadha, and Chirag Shah. 2024. How well do llms represent values across cultures? empirical analysis of llm responses based on hofstede cultural dimensions. *Preprint*, arXiv:2406.14805.

Cody Kommers, Drew Hemment, Maria Antoniak, Joel Z. Leibo, Hoyt Long, Emily Robinson, and Adam Sobey. 2025. Meaning is not a metric: Using llms to make cultural context legible at scale. *Preprint*, arXiv:2505.23785.

Carolyn Korsmeyer and David Sutton. 2011. The sensory experience of food. *Food, Culture & Society*, 14(4):461–475.

Austin C Kozlowski, Matt Taddy, and James A Evans. 2019. The geometry of culture: Analyzing the meanings of class through word embeddings. *American Sociological Review*, 84(5):905–949.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, pages 611–626.

Anton Lavrouk, Tarek Naous, Alan Ritter, and Wei Xu. 2025. What are foundation models cooking in the post-soviet world? In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 20698–20720, Suzhou, China. Association for Computational Linguistics.

David Lazer, Alex Pentland, Lada Adamic, Sinan Aral, Albert-László Barabási, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, and 1 others. 2009. Computational social science. *Science*, 323(5915):721–723.

CHENG LI, Mengzhuo Chen, Jindong Wang, Sunayana Sitaram, and Xing Xie. 2024. CultureLLM: Incorporating cultural differences into large language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

Huihan Li, Liwei Jiang, Jena D Hwang, Hyunwoo Kim, Sebastin Santy, Taylor Sorensen, Bill Yuchen Lin, Nouha Dziri, Xiang Ren, and Yejin Choi. 2024. Culture-gen: Revealing global cultural perception in language models through natural language prompting. *arXiv preprint arXiv:2404.10199*.

Chen Cecilia Liu, Iryna Gurevych, and Anna Korhonen. 2025. Culturally aware and adapted NLP: A taxonomy and a survey of the state of the art. *Transactions of the Association for Computational Linguistics*, 13:652–689.

Xiao Liu, Yansong Feng, Jizhi Tang, Chengang Hu, and Dongyan Zhao. 2022. Counterfactual recipe generation: Exploring compositional generalization in a realistic scenario. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7354–7370, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Deborah Lupton. 1996. Food, the body and the self.

Claudia Mazzuca and Asifa Majid. 2023. The semantic representation of food is shaped by cultural experience. *Language and Cognition*, 15(4):651–669.

Danielle I. McCarthy. 2025. Nutritional intelligence in the food system: Combining food, health, data and ai expertise. *Nutrition Bulletin*, 50(1):142–150.

Sidney W Mintz. 1985. *Sweetness and power: The place of sugar in modern history*. Penguin.

Sidney W. Mintz and Christine M. Du Bois. 2002. The anthropology of food and eating. *Annual Review of Anthropology*, 31(Volume 31, 2002):99–119.

Fnu Mohbat and Mohammed J. Zaki. 2024. Llava-chef: A multi-modal generative model for food recipes. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, CIKM '24, page 1711–1721, New York, NY, USA. Association for Computing Machinery.

Junho Myung, Nayeon Lee, Yi Zhou, Jiho Jin, Rifki Afina Putri, Dimosthenis Antypas, Hsuvas Borkakoty, Eunsu Kim, Carla Perez-Almendros, Abinew Ali Ayele, Víctor Gutiérrez-Basulto, Yazmín Ibáñez García, Hwaran Lee, Shamsuddeen Hassan Muhammad, Kiwoong Park, Anar Sabuhi Rzayev, Nina White, Seid Muhie Yimam, Mohammad Taher Pilehvar, and 3 others. 2024. Blend: A benchmark for llms on everyday knowledge in diverse cultures and languages. In *Advances in Neural Information Processing Systems*, volume 37, pages 78104–78146. Curran Associates, Inc.

Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. StereoSet: Measuring stereotypical bias in pretrained language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5356–5371, Online. Association for Computational Linguistics.

Allen Newell and Herbert A Simon. 2007. Computer science as empirical inquiry: Symbols and search. In *ACM Turing award lectures*, page 1975.

M. E. J. Newman and M. Girvan. 2004. Finding and evaluating community structure in networks. *Phys. Rev. E*, 69:026113.

Peter Newmark. 2003. A textbook of translation.

Tuan-Phong Nguyen, Simon Razniewski, Aparna S. Varde, and Gerhard Weikum. 2022. Extracting cultural commonsense knowledge at scale. *Proceedings of the ACM Web Conference 2023*.

Tuan-Phong Nguyen, Simon Razniewski, and Gerhard Weikum. 2024. Cultural commonsense knowledge for intercultural dialogues. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, CIKM '24, page 1774–1784, New York, NY, USA. Association for Computing Machinery.

Donald M Nonini. 2013. The local-food movement and the anthropology of global systems. *American Ethnologist*, 40(2):267–275.

Shramay Palta and Rachel Rudinger. 2023. FORK: A bite-sized test set for probing culinary cultural biases in commonsense reasoning models. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9952–9962, Toronto, Canada. Association for Computational Linguistics.

Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22.

Chantal Pellegrini, Ege Özsoy, Monika Wintergerst, and Georg Groh. 2021. Exploiting food embeddings for ingredient substitution. In *International Conference on Health Informatics*.

Jeffrey M Pilcher. 2023. *Food in world history*. Routledge.

Paweł Polak. 2016. Computing as empirical science-evolution as a concept. *Studies in Logic, Grammar and Rhetoric*, 48(1).

Jean-Pierre Poulain. 2012. The sociology of gastronomic decolonization. *The gaze of the west and framings of the east. Frontiers of Globalization Series. Palgrave Macmillan, London*, pages 218–232.

Jean-Pierre Poulain, Elise Mognard, Jacqui Kong, Jan Li Yuen, Laurence Tibère, Cyrille Laporte, Fong-Ming Yang, Anindita Dasgupta, Pradeep Kumar Nair, Neethiahnanthan Ari Ragavan, and Ismail Mohd Noor. 2023. Much more than food: The malaysian breakfast, a socio-cultural perspective. *Sustainability*, 15(3).

Marcel Proust. 1913. *À la recherche du temps perdu*. Grasset.

Saed Rezayi, Zhengliang Liu, Zihao Wu, Chandra Dhakal, Bao Ge, Haixing Dai, Gengchen Mai, Ninghao Liu, Chen Zhen, Tianming Liu, and Sheng Li. 2025. Exploring new frontiers in agricultural nlp: Investigating the potential of large language models for food applications. *IEEE Transactions on Big Data*, 11(3):1235–1246.

Paul Rozin. 2005. The meaning of food in our lives: a cross-cultural perspective on eating and well-being. *Journal of nutrition education and behavior*, 37:S107–S112.

Sougata Saha and Monojit Choudhury. 2025. User behavior prediction as a generic, robust, scalable, and low-cost evaluation strategy for estimating generalization in LLMs. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 11047–11065, Vienna, Austria. Association for Computational Linguistics.

Sougata Saha, Saurabh Kumar Pandey, and Monojit Choudhury. 2025. Meta-cultural competence: Climbing the right hill of cultural awareness. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8025–8042, Albuquerque, New Mexico. Association for Computational Linguistics.

Yonadav Shavit, Sandhini Agarwal, Miles Brundage, Steven Adler, Cullen O'Keefe, Rosie Campbell, Teddy Lee, Pamela Mishkin, Tyna Eloundou, Alan Hickey, and 1 others. 2023. Practices for governing agentic ai systems. *Research Paper, OpenAI*.

Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2019. The woman worked as a babysitter: On biases in language generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3407–3412, Hong Kong, China. Association for Computational Linguistics.

Weiyan Shi, Ryan Li, Yutong Zhang, Caleb Ziems, Sunny Yu, Raya Horesh, Rogério Abreu De Paula, and Diyi Yang. 2024. CultureBank: An online community-driven knowledge base towards culturally aware language technologies. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4996–5025, Miami, Florida, USA. Association for Computational Linguistics.

G. Sproesser, M. B. Ruby, C. S. Akotia, M. dos Santos Alvarenga, R. Bhangaonkar, I. Furumitsu, S. Imada, G. Kaptan, M. Kaufer-Horwitz, C. Fischler, P. Rozin, H. T. Schupp, and B. Renner. 2025. Food cultures: Understanding traditional and modern eating. In H. Meiselman, editor, *Handbook of Eating and Drinking: Interdisciplinary Perspectives*. Springer. In Press.

Gudrun Sproesser, Matthew B. Ruby, Naomi Arbit, Charity S. Akotia, Marle dos Santos Alvarenga, Rachana Bhangaokar, Isato Furumitsu, Xiaomeng Hu, Sumio Imada, Gülbanu Kaptan, Martha Kaufer-Horwitz, Usha Menon, Claude Fischler, Paul Rozin, Harald T. Schupp, and Britta Renner. 2022. Similar or different? comparing food cultures with regard to traditional and modern eating across ten countries. *Food Research International*, 157:111106.

John E Staller. 2003. An examination of the palaeobotanical and chronological evidence for an early introduction of maize (zea mays l.) into south america: a response to pearsall. *Journal of Archaeological Science*, 30(3):373–380.

Andrew Steptoe, Tessa M Pollard, and Jane Wardle. 1995. Development of a measure of the motives underlying the selection of food: the food choice questionnaire. *Appetite*, 25(3):267–284.

C Levi Strauss. 1974. Structural anthropology. *Persona & Derecho*, 1:571.

David Sutton. 2018. Cooking in theory: Risky events in the structure of the conjuncture. *Anthropological Theory*, 18(1):81–105.

David E. Sutton. 2010. Food and the senses. *Annual Review of Anthropology*, 39(Volume 39, 2010):209–223.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, and 1 others. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.

Vincent Antonio Traag, Ludo Waltman, and Nees Jan van Eck. 2018. From louvain to leiden: guaranteeing well-connected communities. *Scientific Reports*, 9.

Yixin Wan, Jieyu Zhao, Aman Chadha, Nanyun Peng, and Kai-Wei Chang. 2023. Are personalized stochastic parrots more dangerous? evaluating persona biases in dialogue systems. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9677–9705, Singapore. Association for Computational Linguistics.

Ruiyi Wang, Haofei Yu, Wenxin Zhang, Zhengyang Qi, Maarten Sap, Yonatan Bisk, Graham Neubig, and Hao Zhu. 2024. SOTOPIA-$\pi$: Interactive learning of socially intelligent language agents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12912–12940, Bangkok, Thailand. Association for Computational Linguistics.

Alan Warde. 2016. *The practice of eating*. John Wiley & Sons.

Richard Wilk. 2006. *Home cooking in the global village: Caribbean food from buccaneers to ecotourists*. Berg.

Genta Indra Winata, Frederikus Hudi, Patrick Amadeus Irawan, David Anugraha, Rifki Afina Putri, Wang Yutong, Adam Nohejl, Ubaidillah Ariq Prathama, Nedjma Ousidhoum, Afifa Amriani, Anar Rzayev, Anirban Das, Ashmari Pramodya, Aulia Adila, Bryan Wilie, Candy Olivia Mawalim, Cheng Ching Lam, Daud Abolade, Emmanuele Chersoni, and 32 others. 2025a. WorldCuisines: A massive-scale benchmark for multilingual and multicultural visual question answering on global cuisines. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3242–3264, Albuquerque, New Mexico. Association for Computational Linguistics.

Genta Indra Winata, Frederikus Hudi, Patrick Amadeus Irawan, David Anugraha, Rifki Afina Putri, Wang Yutong, Adam Nohejl, Ubaidillah Ariq Prathama, Nedjma Ousidhoum, Afifa Amriani, Anar Rzayev, Anirban Das, Ashmari Pramodya, Aulia Adila, Bryan Wilie, Candy Olivia Mawalim, Cheng Ching Lam, Daud Abolade, Emmanuele Chersoni, and 32 others. 2025b. WorldCuisines: A massive-scale benchmark for multilingual and multicultural visual question answering on global cuisines. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3242–3264, Albuquerque, New Mexico. Association for Computational Linguistics.

Qing Zhang, Christoph Trattner, Bernd Ludwig, and David Elsweiler. 2019. Understanding cross-cultural visual food tastes with online recipe platforms. *Proceedings of the International AAAI Conference on Web and Social Media*, 13(01):671–674.

Yunke Zhang, Yiran Fan, Peijie Liu, Fengli Xu, and Yong Li. 2025. Cyber food swamps: Investigating the impacts of online-to-offline food delivery platforms on healthy food choices. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 19, pages 2260–2272.

Huixue Zhou, Lisa S. Chow, Lisa Harnack, Satchidananda Panda, Emily N.C. Manoogian, Minchen Li, Yongkang Xiao, and Rui Zhang. 2025a. Nutrirag: Unleashing the power of large language models for food identification and classification through retrieval methods. *medRxiv*.

Li Zhou, Taelin Karidi, Wanlong Liu, Nicolas Garneau, Yong Cao, Wenyu Chen, Haizhou Li, and Daniel Hershcovich. 2025b. Does mapo tofu contain coffee? probing LLMs for food-related cultural knowledge. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 9840–9867, Albuquerque, New Mexico. Association for Computational Linguistics.

Naitian Zhou, David Bamman, and Isaac L. Bleaman. 2025c. Culture is not trivia: Sociocultural theory for cultural NLP. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 25869–25886, Vienna, Austria. Association for Computational Linguistics.

Xuhui Zhou, Zhe Su, Tiwalayo Eisape, Hyunwoo Kim, and Maarten Sap. 2024. Is this the real life? is this just fantasy? the misleading success of simulating social interactions with LLMs. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 21692–21714, Miami, Florida, USA. Association for Computational Linguistics.

Xuhui Zhou, Hao Zhu, Leena Mathur, Ruohong Zhang, Haofei Yu, Zhengyang Qi, Louis-Philippe Morency, Yonatan Bisk, Daniel Fried, Graham Neubig, and 1 others. 2023. Sotopia: Interactive evaluation for social intelligence in language agents. *arXiv preprint arXiv:2310.11667*.

# A Appendix

## A.1 Case Study Details

The Food Choice Questionnaire (FCQ) (Steptoe et al., 1995) captures food-related behaviors across nine dimensions: familiarity, convenience, health, mood, sensory appeal, natural content, price, weight control, and ethical concern. For our study, we selected two dimensions: (i) **Convenience:** Represented by the question *"The food in country that can be bought in shops close to where someone lives or works is?"* (ii) **Familiarity:** represented by the question *"The food in country that people usually eat is?"*

We prompted two LLMs - Llama-3.1-8B-Instruct (Dubey et al., 2024) and Gemma-2-9B-it (Team et al., 2024) - for 170 countries ($C$), using a curated list of approximately 3,500 ($N$) food items collected from resources such as TasteAtlas[3] and cross-country cuisine knowledge bases (Winata et al., 2025b). For each prompt, we recorded the log probabilities of tokens corresponding to the $N$ food items. If a food item spanned multiple tokens, we averaged across them to obtain its final probability. This produced, for each model, an $C \times N$ matrix ($M$), where each row represents the probability distribution of $N$ food items for a given country. We derived a country-country weighted adjacency matrix ($S_{cc}$) by computing $MM^\top$, which encodes the relatedness of countries based on food probabilities. To retain meaningful structure, we build a *mutual* $k$-NN graph with $k = 10$, and applied a community detection algorithm to identify clusters of countries.

Finally, we curated a mapping between countries and their historical colonizers and used Fisher's exact test to evaluate whether the LLM-derived food-based communities significantly aligned with colonial histories within continents.

All experiments were run using vLLM (Kwon et al., 2023) with a batch size of 8 and temperature set at 1 on a 48 GB NVIDIA RTX 6000 Ada Generation GPU with FP16 quantization, and took approximately 12 hours to execute. Below is a sample prompt for Llama.

**Prompt for logprobs prediction for a Question**

<|begin_of_text|><|start_header_id|>system<|end_header_id|>

Cutting Knowledge Date: December 2023
Today Date: 26 Jul 2024

<|eot_id|><|start_header_id|>user<|end_header_id|>

Answer the given question related to food items in a country. Generate only one food item. Do not output any extra text. [Question] <|eot_id|><|start_header_id|>assistant|<|end_header_id|>

[3] https://www.tasteatlas.com/

| Country | Year | Source | Top Ordered Dishes | URL |
|---|---|---|---|---|
| Australia | 2023 | Doordash | Chicken dishes, Pizza, Burgers, Wagyu beef dishes, Chips | https://downloads.ctfassets.net/trvmqu12jq2l/5fTVhyjeP5pNS29PrXNaA5/5c859c8d0567564f7e881692afd0a22b/2024-Restaurant-Alcohol-Online-Ordering-Trends-Report-AU-NZ.pdf |
| New Zealand | 2023 | Doordash | Chicken dishes, Chips, Burgers, Naan, Noodle dishes | https://downloads.ctfassets.net/trvmqu12jq2l/5fTVhyjeP5pNS29PrXNaA5/5c859c8d0567564f7e881692afd0a22b/2024-Restaurant-Alcohol-Online-Ordering-Trends-Report-AU-NZ.pdf |
| USA | 2024 | Doordash | French fries, Garlic naan, Tacos, Pad thai, Miso soup, California roll, Chicken tikka masala, Edamame, Chips, Burrito | https://assets.ctfassets.net/trvmqu12jq2l/5fTVhyjeP5pNS29PrXNaA5/2c94dfa2b45b5ba60cf844d4dd598617/2025-DoorDash-Delivery-Trends-Report.pdf |
| Canada | 2024 | Doordash | Naan, Butter chicken, Miso soup, California roll, Tandoori roti, Burrito, Dynamite roll, Poutine, Chicken biryani, Steamed rice | https://downloads.ctfassets.net/trvmqu12jq2l/5fTVhyjeP5pNS29PrXNaA5/9117580f0d1b1ef9e49aa0a19ce4eec6/2025-DoorDash-Canada-Delivery-Trends-Report.pdf |
| UAE | 2021 | Talabat | French fries, chicken sandwich, Burgers, potato wedges, Pizza | https://hotelandcatering.com/foodservice/talabat-unveils-uaes-most-ordered-food-and-grocery-items-in-2021-and-youll-be-surprised |
| Oman | 2023 | Talabat | Fried Chicken, Burgers, sandwiches | https://www.zawya.com/en/press-release/companies-news/talabat-unveils-customer-ordering-trends-across-food-groceries-and-donations-throughout-2023-lhwjfcrb |
| India | 2024 | Zomato | Biryani, Pizza | https://world.hey.com/bhari/india-food-trends-2024-from-zomato-34ea104e |

Table 1: Country-wise top-ordered dishes between 2021 and 2024. DoorDash is a U.S.-based food delivery platform operating primarily in North America, Australia, and New Zealand, which publishes annual regional trends reports. Talabat is one of the largest food delivery services in the Middle East and North Africa (UAE, Kuwait, Saudi Arabia, Oman, etc.). Zomato, headquartered in India, is a major platform across South Asia offering food delivery and restaurant discovery. For Talabat and Zomato, statistics were drawn from trusted media reports (e.g., NDTV, Gulf News, Zawya) that cite the platforms' annual "year-in-review" data, as these platforms do not publish standalone public reports.

| Model | Question | Nodes | Edges | # Communities | $Q$ | Drawn intra / inter$^{\dagger}$ |
|---|---|---|---|---|---|---|
| Llama | Convenience | 165 | 567 | 10 | 0.829 | 482 / 28 |
| | Familiarity | 169 | 568 | 11 | 0.831 | 488 / 23 |
| Gemma | Convenience | 166 | 554 | 11 | 0.821 | 460 / 38 |
| | Familiarity | 168 | 523 | 10 | 0.796 | 426 / 44 |

Table 2: **Country–country networks.** Mutual $k$-NN ($k$=10) on $S_{cc}=MM^{\top}$; communities of size $< 3$ removed; $Q$ from Leiden. $^{\dagger}$Edges drawn after the 90% weight quantile filter used *only* for visualization.

**Leiden algorithm** We used the Leiden algorithm (Traag et al., 2018) for community detection. Leiden is an iterative graph partitioning method that optimizes a quality function such as modularity. It improves upon the Louvain algorithm (Blondel et al., 2008) by ensuring that all detected communities are well connected and by modifying the communities to avoid fragmentation.

The algorithm proceeds in three main phases:

1. **Local moving of nodes:** Nodes are moved to neighboring communities if this improves the quality function (e.g., modularity).

2. **Refinement:** Each community is further refined into subcommunities to ensure internal connectivity and prevent the formation of disconnected clusters.

3. **Aggregation:** A new aggregate network is created where each refined community becomes a node. The process is then repeated on this coarser network.

These steps are iterated until no further improvements are possible. Unlike Louvain, Leiden ensures that all resulting communities are internally connected, and when run iteratively, it converges to partitions that are locally optimal and free of disconnected components.

**Country-colonizer mapping.** We curate colonizer labels from CEPII's Geo dataset[4] (`geo_cepii.xls`) and map coordinates from Google's canonical countries file[5]. For each country, we compute a *colonizer set* as the union

[4] https://www.cepii.fr/CEPII/en/bdd_modele/bdd_modele_item.asp?id=6

[5] https://developers.google.com/public-data/docs/canonical/countries_csv

of all nonempty entries in `colonizer1..4` and `short_colonizer1..3`. We restrict Fisher tests to within-continent pairs, testing the association between *same colonizer* (nonempty intersection of colonizer sets) and *same community*. Restricting to the same continent controls for geography. Countries on one continent often share ingredients, trade links, and migration routes, so we compare like with like. This checks whether a shared colonizer still predicts being in the same community after accounting for regional similarity. It prevents inflated odds from obvious continent clustering (for example, Europe with Europe).

| Community label | Countries | Pairs sharing colonizers count | Colonizers |
|---|---|---|---|
| Europe & European Territories | Austria, Belgium, Czech republic, Denmark, Estonia, Finland, France, French polynesia, Germany, Greenland, Haiti, Hungary, Iceland, Italy, Latvia, Liechtenstein, Lithuania, Luxembourg, Madagascar, Malta, Mauritius, Monaco, New caledonia, Norway, Poland, Russia, Seychelles, Slovakia, Sweden, Switzerland, Ukraine | 40 | Austria, Denmark, France, Germany, Hungary, Netherlands, Russian federation, Spain, Sweden, United kingdom |
| South & East Asia | Afghanistan, Bangladesh, Brunei, Cambodia, China, India, Indonesia, Japan, Laos, Malaysia, Maldives, Mongolia, Myanmar, Nepal, North korea, Pakistan, Philippines, Singapore, South korea, Sri lanka, Taiwan, Thailand, Vietnam | 53 | China, France, India, Japan, Netherlands, Spain, United kingdom, United states of america |
| Balkans–Caucasus–Anatolia | Armenia, Azerbaijan, Bosnia and herzegovina, Bulgaria, Croatia, Cyprus, Georgia, Greece, Iran, Kazakhstan, Moldova, Montenegro, North macedonia, Romania, Serbia, Slovenia, Turkiye | 69 | Austria, Greece, Hungary, Russian federation, Serbia and montenegro, Türkiye, United kingdom |
| West Africa | Benin, Burkina faso, Cameroon, Chad, Democratic republic of the congo, Gabon, Ghana, Ivory coast, Liberia, Mali, Mauritania, Nigeria, Senegal, Sierra leone, The gambia, Togo | 61 | Belgium and luxembourg, France, United kingdom |
| Southern & Eastern Africa (Lusophone ties) | Angola, Botswana, Brazil, Cape verde, Ethiopia, Kenya, Lesotho, Mozambique, Namibia, Portugal, Rwanda, Somalia, South africa, South sudan, Tanzania, Zambia | 36 | Belgium and luxembourg, Egypt, Germany, Italy, Netherlands, Portugal, South africa, United kingdom |
| Caribbean | Antigua and barbuda, Aruba, Barbados, Bermuda, Bonaire, Curaçao, Dominica, Grenada, Guyana, Jamaica, Netherlands, Saint kitts and nevis, Saint vincent and the grenadines, Suriname, The bahamas, Trinidad and tobago | 65 | France, Netherlands, Spain, United kingdom |
| Latin America & Iberia | Andorra, Argentina, Bolivia, Chile, Colombia, Cuba, Dominican republic, Ecuador, Mexico, Paraguay, Peru, Puerto rico, Spain, Uruguay, Venezuela | 78 | Brazil, France, Haiti, Spain, United states of america |
| Anglosphere & Oceania | Australia, England, Fiji, Guam, Ireland, Marshall islands, Nauru, New zealand, Northern ireland, Palau, Papua new guinea, Samoa, Scotland, Vanuatu, Wales | 31 | Australia, France, Germany, Japan, New zealand, United kingdom, United states of america |
| Middle East | Bahrain, Egypt, Iraq, Israel, Jordan, Lebanon, Oman, Palestine, Qatar, Saudi arabia, Syria, United arab emirates, Yemen | 42 | France, Poland, Spain, Türkiye, United kingdom |
| North Africa | Algeria, Morocco, Tunisia | 3 | France, Türkiye |

Table 3: Llama's convenience communities: countries, shared-colonizer pairs, and colonizers. **Labels are heuristic; they reflect the majority of countries in each cluster**.

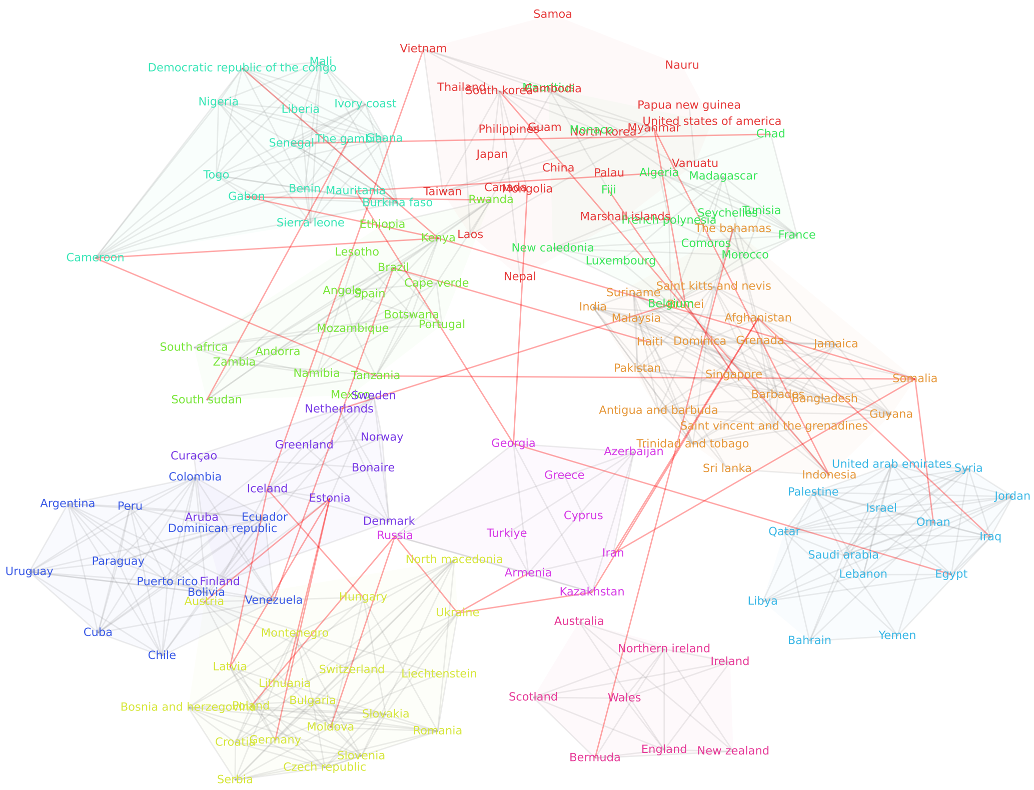


Figure 3: **Country communities from Gemma (convenience).** Weighted mutual-$k$NN graph ($k = 10$) on $S_{cc} = MM^{\top}$; Leiden partition on modularity. High modularity ($Q = 0.821$) with coherent regional clusters. Intra-community edges shown in gray, inter-community edges in red. Node colors are used solely for visual separation of clusters and do not carry additional semantic meaning.

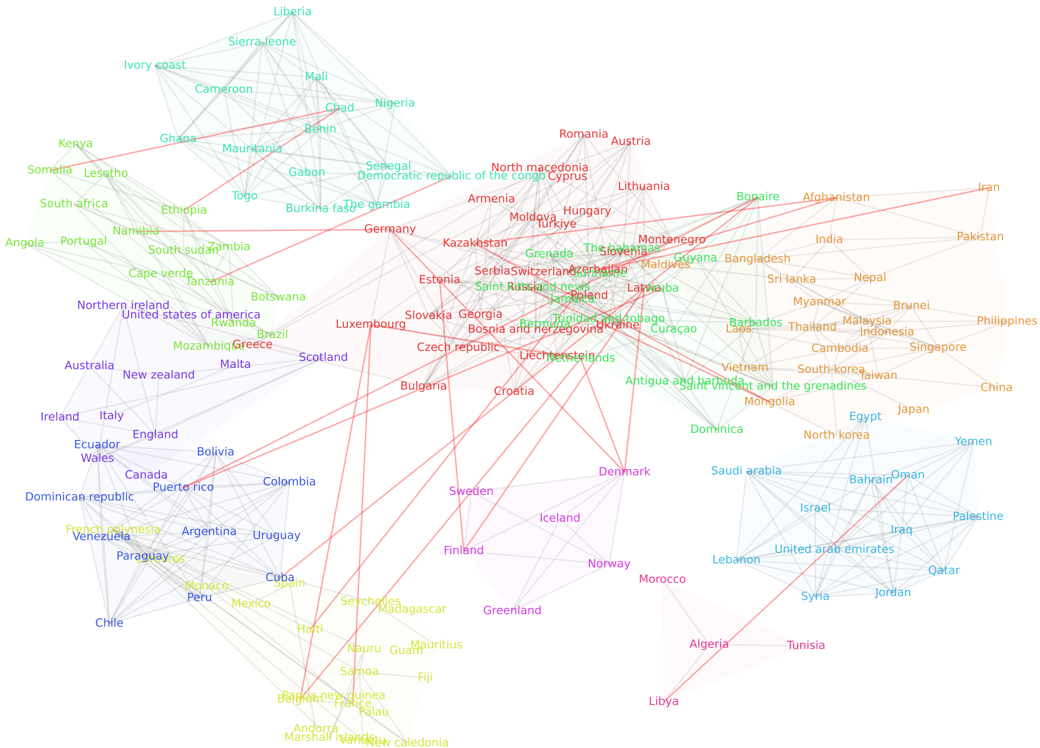


Figure 4: **Country communities from Llama (familiarity).** Weighted mutual-$k$NN graph ($k = 10$) on $S_{cc} = MM^{\top}$; Leiden partition on modularity. High modularity ($Q = 0.831$) with coherent regional clusters. Intra-community edges shown in gray, inter-community edges in red; light hulls mark communities. Node colors are used solely for visual separation of clusters and do not carry additional semantic meaning.

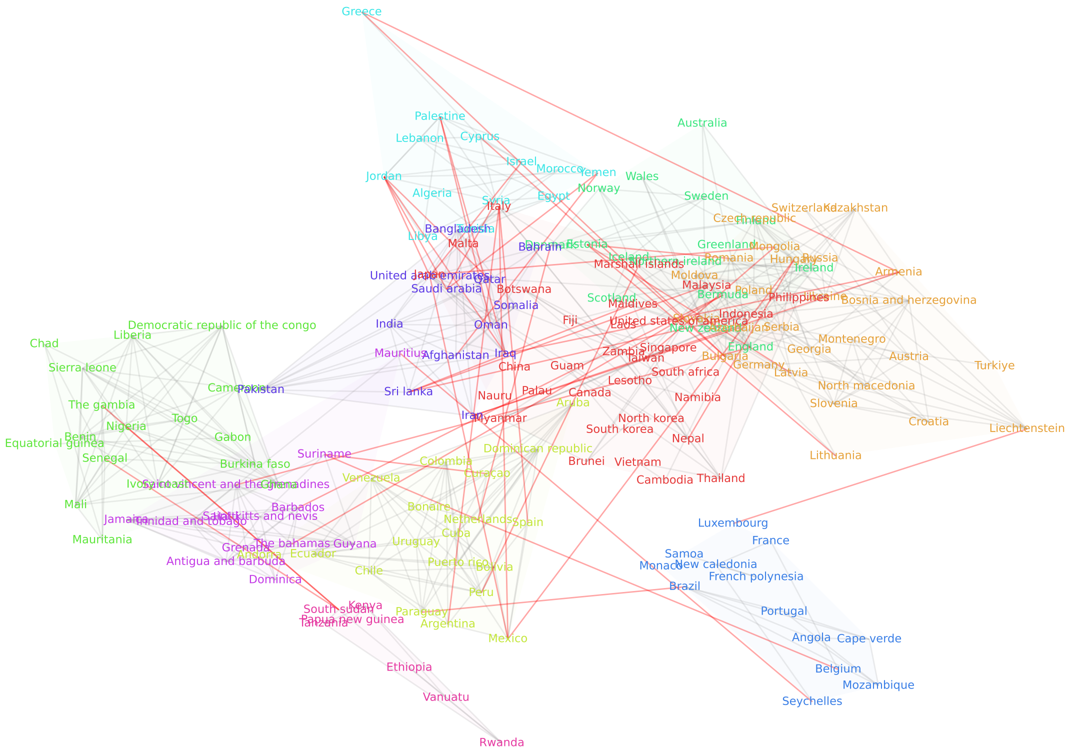


Figure 5: **Country communities from Gemma (familiarity).** Weighted mutual-$k$NN graph ($k = 10$) on $S_{cc} = MM^\top$; Leiden partition on modularity. High modularity ($Q = 0.796$) with coherent regional clusters. Intra-community edges shown in gray, inter-community edges in red. Node colors are used solely for visual separation of clusters and do not carry additional semantic meaning.

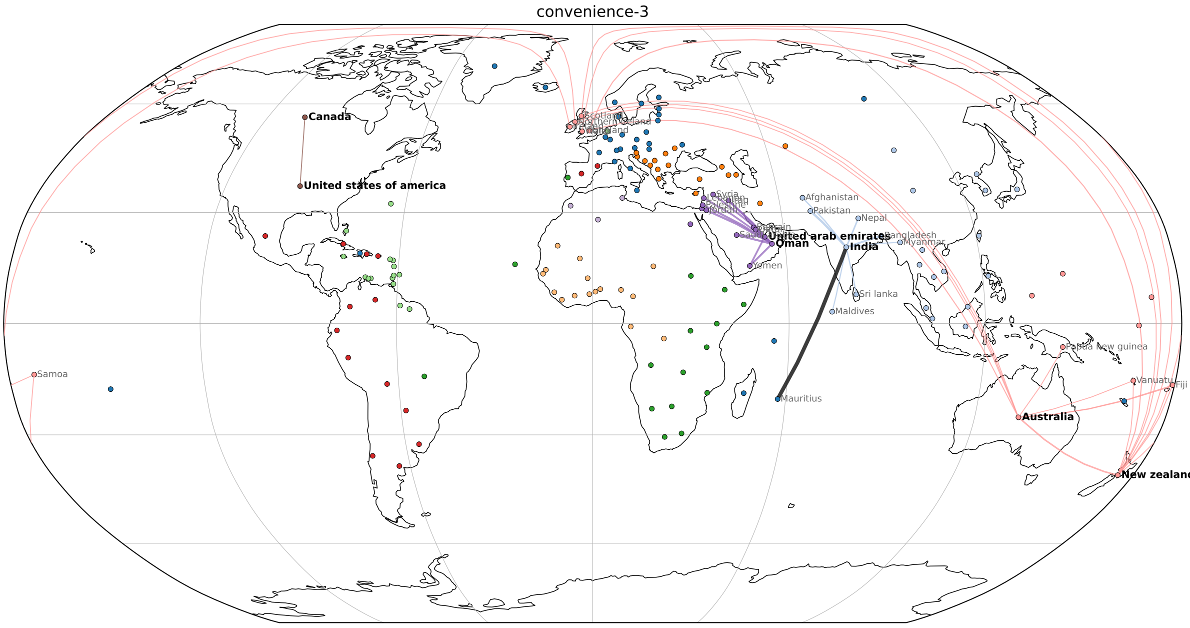


Figure 6: **Country food-similarity graph on world map (Llama, convenience).** Intra-community edges use the community color with weight-scaled thickness; inter-community edges are dark gray. Labels are shown only for focal countries and their neighbors. This map projection highlights how model-inferred communities align with geographic and colonial patterns, while inter-community ties remain sparse. Node colors are used solely for visual separation of clusters and do not carry additional semantic meaning.

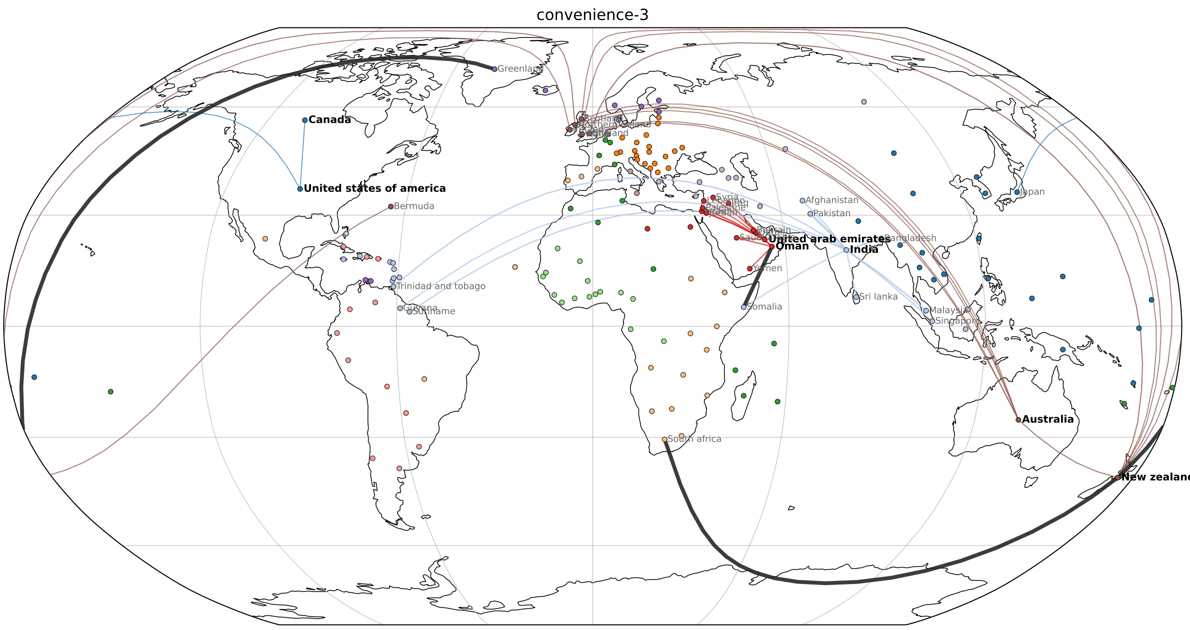


Figure 7: **Country food-similarity graph on world map (Gemma, convenience).** Node colors are used solely for visual separation of clusters and do not carry additional semantic meaning.

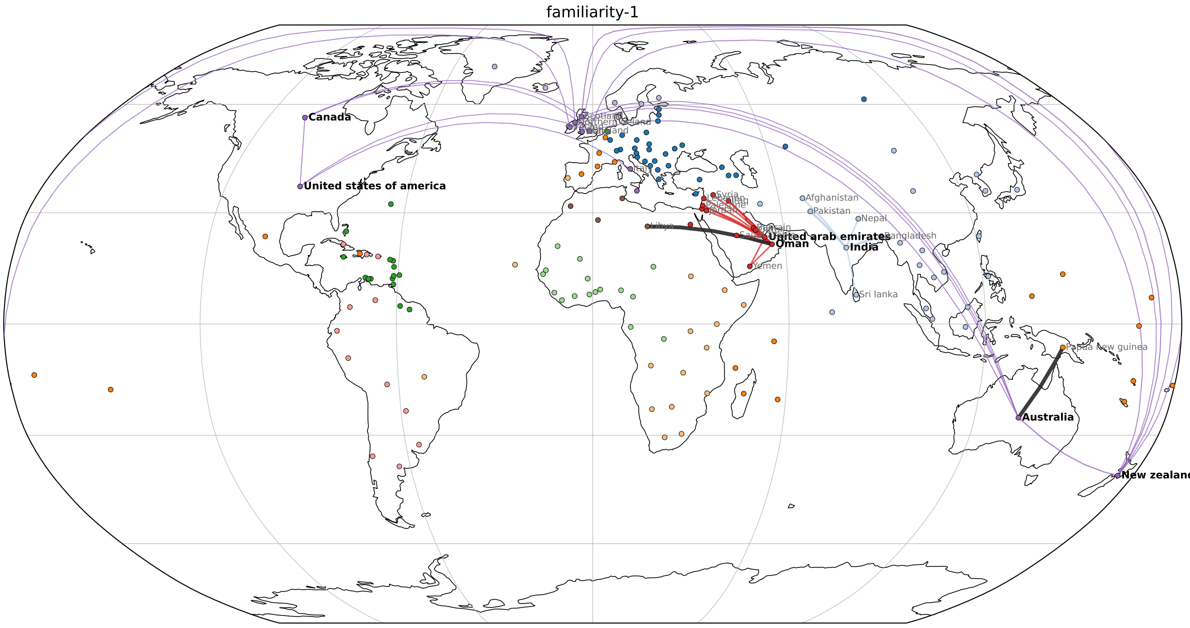


Figure 8: **Country food-similarity graph on world map (Llama, familiarity).** Node colors are used solely for visual separation of clusters and do not carry additional semantic meaning.

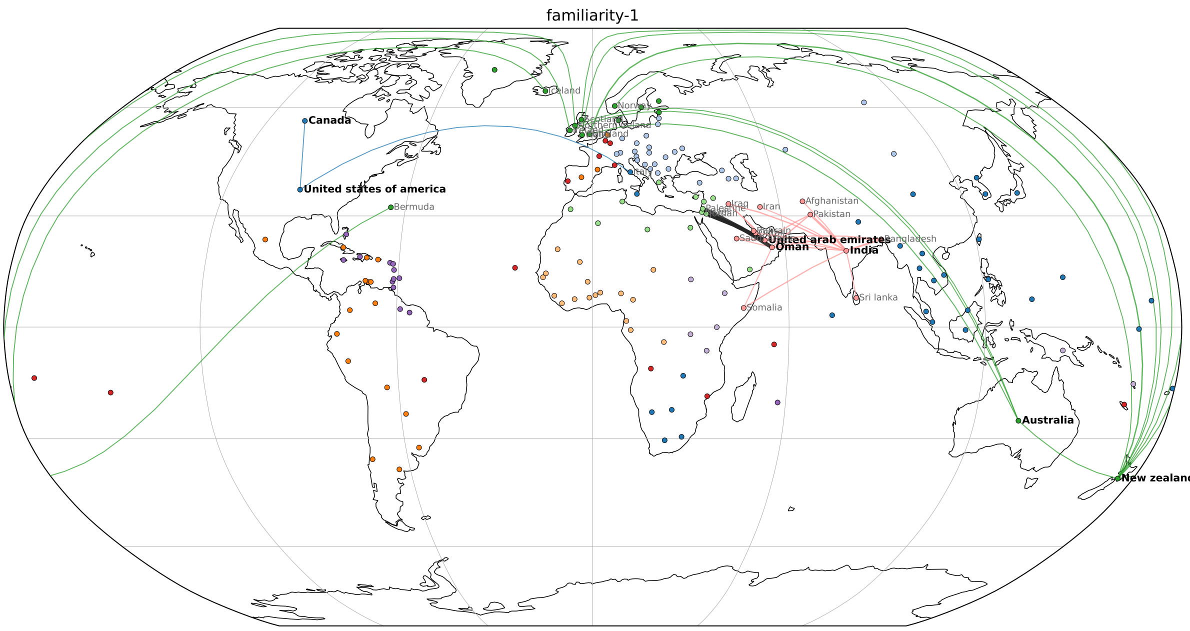


Figure 9: **Country food-similarity graph on world map (Gemma, familiarity).**Node colors are used solely for visual separation of clusters and do not carry additional semantic meaning.