

1 WikiGap: Promoting Epistemic Equity by Surfacing Knowledge Gaps Between 2 English Wikipedia and other Language Editions 3

4 ZINING WANG, University of British Columbia, Canada and Vector Institute for AI, Canada
5

6 YUXUAN ZHANG, University of British Columbia, Canada
7

8 DONGWOOK YOON, University of British Columbia, Canada
9

10 NICHOLAS VINCENT, Simon Fraser University, Canada
11

12 FARHAN SAMIR, University of British Columbia, Canada and University of Toronto, Canada
13

14 VERED SHWARTZ, University of British Columbia, Canada and Vector Institute for AI, Canada
15

The figure shows a screenshot of the WikiGap interface. At the top, there is a navigation bar with 'Article' and 'Talk' buttons. Below the navigation bar, the title 'Peking duck' is displayed. A 'Source Language Filter' button is highlighted in yellow and labeled 'D4: Source Language Filter →'. To the right of the filter, there is a 'Logs' button and a '30 Facts' button. Below the filter, there is a search bar with a magnifying glass icon. A dropdown menu shows language options: 'All' (selected), 'Chinese', 'French', and 'Russian'. The main content area displays a paragraph about Peking duck, with several words underlined in blue, indicating they are cross-lingual gaps. A 'D5: Cross-lingual Fact Search →' button is positioned next to the paragraph. The paragraph text discusses the history and preparation of Peking duck. Below the paragraph, there is a 'History [edit]' section. A 'D1: In-text Gap Markers' button is highlighted in yellow and positioned next to the history section. The text in this section also contains underlined words. To the right of the history section, there is a 'D3: Fact Provenance Card' button. Further down, there are sections for 'French' and 'D2: Cross-Lingual Fact Panel'. Each section has a 'View on Chinese Wikipedia' button.

42 Fig. 1. The WikiGap interface embeds cross-lingual facts into English Wikipedia via five key design elements (D1–D5), supporting
43 in-place access, traceability, and multilingual engagement.
44

45 Authors' Contact Information: Zining Wang, University of British Columbia, Canada and Vector Institute for AI, Canada, zining.wang@ubc.ca; Yuxuan
46 Zhang, University of British Columbia, Canada; Dongwook Yoon, University of British Columbia, Canada; Nicholas Vincent, Simon Fraser University,
47 Canada; Farhan Samir, University of British Columbia, Canada and University of Toronto, Canada; Vered Shwartz, University of British Columbia, Canada
48 and Vector Institute for AI, Canada.
49

50 Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not
51 made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components
52 Manuscript submitted to ACM

With more than 11 times as many pageviews as the next largest edition, English Wikipedia dominates global knowledge access relative to other language editions. Readers are prone to assuming English Wikipedia as a superset of all language editions, leading many to prefer it even when their primary language is not English. Other language editions, however, comprise complementary facts rooted in their respective cultures and media environments, which are marginalized in English Wikipedia. While Wikipedia's user interface enables switching between language editions through its Interlanguage Link (ILL) system, it does not reveal to readers that other language editions contain valuable, complementary information. We present WIKIGAP, a system that surfaces complementary facts sourced from other Wikipedias within the English Wikipedia interface. Specifically, by combining a recent multilingual information-gap discovery method [52] with a user-centered design, WIKIGAP enables access to complementary information from French, Russian, and Chinese Wikipedia. In a mixed-methods study (n=21), WIKIGAP significantly improved fact-finding accuracy, reduced task time, and received a 32-point higher usability score relative to Wikipedia's current ILL-based navigation system. Participants reported increased awareness of the availability of complementary information in non-English editions and reconsidered the completeness of English Wikipedia. WIKIGAP thus paves the way for promoting epistemic equity across language editions. Taken together, these findings show how interface design can redistribute reader attention and engagement across Wikipedia's multilingual knowledge production communities, supporting epistemic equity.

ACM Reference Format:

Zining Wang, Yuxuan Zhang, Dongwook Yoon, Nicholas Vincent, Farhan Samir, and Vered Shwartz. 2026. WikiGap: Promoting Epistemic Equity by Surfacing Knowledge Gaps Between English Wikipedia and other Language Editions. 1, 1 (January 2026), 29 pages. <https://doi.org/10.1145/nnnnnnnnnnnnnn>

1 Introduction

Wikipedia hosts more than 300 language editions, inviting contributions from many distinct linguistic and geographic regions. The investment in creating multiple language editions reflects Wikipedia's principles of democratizing both access to and contribution of information. These language editions are sociogeographically situated; it is generally well understood in Wikipedia scholarship that they are not mere translations of one another [26]. Instead, the language editions tend towards reflecting the knowledge that is contextually relevant, often based on cultural experience and localized media diets [4, 34].

However, this heterogeneity is not well understood by the broader public that engage with Wikipedia, whether readers or contributors [43]. Rather, the popular understanding of multilingual Wikipedia is that there is a global consensus on what information is reliable and noteworthy [26]. Moreover, English Wikipedia is by and large perceived as the most comprehensive and complete account of this consensus, containing as much as or more information than any other, making other languages the minority. This *English-as-superset* assumption is widespread, even among people who speak multiple languages proficiently; in a 2011 report, 93% of all editors surveyed attested to reading English Wikipedia, even though it was a primary language for only 52% of them [34, 60]. In the same vein, 20% of English Wikipedia's visitors were based in India, even more than readers from the United Kingdom. English Wikipedia is so widely understood to be an epistemic authority that is commonly treated as a source of neutral and complete information in and of itself by editors contributing to minority language editions [28].

This epistemic assumption is consequential not only for readers' understanding of Wikipedia but also for the organization of its peer production communities. Wikipedia is sustained by interconnected yet unevenly resourced

of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

Manuscript submitted to ACM

105 volunteer communities, each centered on different language editions [54]. Prior research has shown that volunteer
106 engagement in Wikipedia depends on motivational factors such as enjoyment, opportunities for learning, and perceived
107 recognition within a contributor community, beyond shared ideological commitments to free and open knowledge
108 [3, 46]. When English Wikipedia is widely perceived as the most complete and authoritative source, reader attention
109 and contributor effort become disproportionately concentrated in the English-language community. This concentration
110 structurally amplifies the conditions that prior work has shown to support sustained participation, while making those
111 conditions less accessible in minority language editions. One byproduct of this imbalance is a reduction in incentives
112 for volunteer contributions to those communities [37], undermining Wikipedia’s ethos as a globally and cooperatively
113 produced multilingual knowledge infrastructure. In this work, we ask how interface design can redistribute attention
114 and engagement across Wikipedia’s multilingual peer production communities, promoting epistemic equity.¹
115

116 Prior works have sought to incorporate contributions in minority language editions by consolidating information
117 from all editions into a single unified knowledge base [e.g., 1, 14]. However, this conceptual formulation fails to *center*
118 minority language editions and the unique volunteer contributions they contain. It continues to perpetuate the ideology
119 of there being a global consensus on noteworthy and reliable information, a universal knowledge base. As such, it does
120 little to challenge the status quo of English Wikipedia as an objective and complete epistemic authority.
121

122 In our work we take a social constructionist view of knowledge [22] in approaching the problem by improving
123 engagement with minority language editions. That is, we center necessarily incomplete and sociogeographically
124 positioned curation of knowledge. To highlight the positionality (rather than objectivity) of English Wikipedia, we aim
125 to highlight its content gaps relative to other language editions. At the same time, we balance this intervention against
126 being minimally disruptive to the current Wikipedia reading design [20]. We present WIKI GAP, an augmentation of
127 the Wikipedia reading experience that surfaces sentence-level factual differences from English to 3 other language
128 editions of Wikipedia (French, Russian, and Chinese). The system combines a computational backend (INFOGAP) that
129 identifies multilingual content gaps [52] with a browser-based interface that presents these differences through in-page
130 highlights and a sidebar display (Figure 1).
131

132 We evaluated WIKI GAP through a mixed-methods user study using a fact-finding task focused on culturally rich
133 food articles – a domain where cross-lingual disparities are particularly salient [35]. Our results show that participants
134 strongly preferred WIKI GAP over Wikipedia’s default interlanguage links system (ILLS), with System Usability Scale
135 (SUS) scores 32 points higher on average. In terms of performance, participants were significantly more accurate and
136 faster at answering factual questions when using WIKI GAP, suggesting that surfacing multilingual information directly
137 within the English interface not only improves user satisfaction but also enhances learning efficiency.
138

139 Through a Post-Study interview, we find that participants, prior to using WIKI GAP, were largely unaware of content
140 gaps in English Wikipedia, demonstrating that the problematic *English-as-superset* epistemic assumption persists more
141 than a decade after it was identified in seminal studies by Hecht [26]. Participants were overwhelmingly appreciative
142 of WIKI GAP’s capacity to surface these substantive information gaps during the reading of English Wikipedia pages,
143 and expressed openness and enthusiasm about interacting with multilingual content in English Wikipedia. Moreover,
144 a few participants with prior editing experience described how the system could support their editorial workflows
145 by enabling source verification and facilitating the integration of multilingual content. These findings underscore
146 WIKI GAP’s broader value as a bridge between communities of editors and readers across languages.
147

153
154 ¹Throughout this paper, we use epistemic equity as a design lens, focusing specifically on how interface design can redistribute attention and engagement
155 across Wikipedia’s multilingual peer production communities, rather than addressing broader institutional or political dimensions of epistemic equity.
156

157 In Section 6, we reflect on how WIKIGAP challenges the English-as-superset assumption by redistributing reader
 158 attention and engagement across Wikipedia’s multilingual peer production communities. We further interpret the
 159 system as a boundary object [56, 57] that supports articulation work between readers and editors across language
 160 editions. We contribute a validated system design for surfacing fact-level cross-lingual information gaps in Wikipedia.
 161 The system combines a computational pipeline that detects and integrates missing facts from non-English editions
 162 with an unobtrusive interface for English readers. A mixed-methods evaluation (n=21) shows that this design improves
 163 fact-finding accuracy and efficiency, outperforms Wikipedia’s interlanguage-link system in usability, and increases
 164 engagement with multilingual knowledge—challenging assumptions about English Wikipedia’s epistemic completeness.²
 165

166 2 Background and Related Work

167 Wikipedia’s multilingual architecture aspires to democratize knowledge access across languages, yet its design and
 168 community practices systematically privilege English content. In §2.1, we examine how this “English-as-superset”
 169 assumption [26] persists despite empirical evidence to the contrary, and how current Wikipedia interface reinforces
 170 epistemic hierarchies that undermine the platform’s multilingual vision by concentrating attention and participation
 171 within a single language edition. In §2.2, we explain that cross-linguistic variation should be surfaced rather than
 172 obscured by drawing literature from prior CSCW research on asset-based design [23, 33, 63] and feminist epistemology
 173 [7, 22]. Then, we review prior technical interventions designed to surface cross-linguistic knowledge differences,
 174 evaluating them along two critical dimensions: the granularity of differences they expose and their integration with
 175 readers’ existing browsing practices to support engagement across language communities. This analysis reveals
 176 persistent limitations that motivate our WIKIGAP design.

177 2.1 English-as-Superset and Epistemic Inequity in Multilingual Wikipedia

178 Wikipedia embodies a central tension in collaborative knowledge systems: although it is architecturally multilingual
 179 and aspires to provide knowledge to “every single person on the planet in their own language” [11], its design and
 180 community practices systematically privilege English content [16, 51]. This imbalance shapes how knowledge is
 181 organized and encountered on the platform, influencing both which perspectives become visible and how reader
 182 attention is distributed across language-specific peer production communities. At the interface level, this dynamic
 183 obscures cross-linguistic differences and reinforces what Hecht [26] term the “English-as-superset” assumption—the
 184 belief that English Wikipedia exhaustively represents the knowledge available across other language editions.

185 CSCW scholarship on epistemic equity [2, 8] helps explain why this assumption matters: collaborative systems
 186 should support the visibility and engagement of diverse epistemic standpoints rather than subsuming them under a
 187 dominant perspective. In Wikipedia’s multilingual context, this means recognizing that language communities produce
 188 knowledge shaped by distinct cultural, geographic, and historical conditions that are not uniformly visible across
 189 language editions

190 **2.1.1 English Wikipedia Is Treated as the Default, Yet Knowledge Gaps Persist.** Given English’s status as a global lingua
 191 franca, the English edition of Wikipedia has far outpaced other language editions in size and coverage. As a result, it is
 192 often implicitly treated as a superset of their knowledge. Hecht [26] identified this “English-as-superset” viewpoint,
 193 which, though rarely stated explicitly, remains prevalent—even among researchers who rely on multilingual Wikipedia
 194 to train text-generation and text-embedding models. However, substantial knowledge gaps persist between English
 195 and other language editions.

196 ²Code and data will be made publicly available upon publication.

209 and other language editions. Hecht and Gergle [25] showed that English Wikipedia lacks coverage on many topics
210 documented in other languages, while more recent work demonstrates that even when topics overlap, English articles
211 often omit facts present elsewhere [52].
212

213 *2.1.2 From Knowledge Gaps to Epistemic Invisibility.* From an epistemic equity perspective, such variation is not an error
214 or deficiency, but an expected outcome of situated knowledge production. One interpretation of Wikipedia’s multilingual
215 vision assumes a language-agnostic set of facts that can be uniformly translated across languages—a view that treats
216 language as a neutral vessel for meaning, abstracted from socio-geographic context [18, Chapter 4]. Yet much of
217 Wikipedia’s content concerns culturally and geographically situated histories, practices, and identities [61]. Accordingly,
218 language editions are shaped by the communities that produce them [30] and exhibit a well-documented “self-focus
219 bias” toward regionally relevant topics [9, 24, 52]. This emphasis on positionality echoes feminist epistemological
220 arguments that knowledge is always produced from particular social, cultural, and linguistic standpoints [22]. When
221 such epistemic differences are rendered invisible, they not only obscure situated knowledge but also limit opportunities
222 for readers and editors to encounter and engage with the communities that produce it.
223

224 The “English-as-superset” assumption can thus contribute to cycles of epistemic invisibility [2], in which visibility
225 and legitimacy are recursively reinforced within collaborative systems. On Wikipedia, these dynamics manifest in both
226 editorial and readership practices. Bilingual editors often consult English Wikipedia to determine what content belongs
227 in other language editions, treating it as a reference point [28]. Wikipedia’s emphasis on verifiability further privileges
228 English-language sources, creating structural barriers for knowledge rooted in other linguistic or cultural contexts [54].
229 Although multilingual editors can bridge across editions, they constitute only a small fraction of Wikipedia’s editor
230 base [21].
231

232 Readers similarly assume that different language editions contain equivalent information (Section 5). As engagement
233 with smaller editions declines, editors in those communities receive less visibility and recognition, weakening key
234 motivations for participation [3, 46]. This feedback loop reduces incentives for contribution and undermines Wikipedia’s
235 vision as a collaboratively produced multilingual knowledge infrastructure [37].
236

237 *2.1.3 Interface Design and the Obscuring of Epistemic Difference.* These epistemic hierarchies are not only reflected in
238 community norms and editorial practices, but are also enacted through interface design. Wikipedia primarily supports
239 cross-lingual navigation through Interlanguage Links (ILLs), presented as a drop-down menu on each article. While ILLs
240 enable access to other language editions, the interface provides no cues indicating that articles may differ substantially
241 in content. By framing language switching as a matter of translation rather than epistemic difference, the design
242 implicitly suggests a shared underlying knowledge base, limiting opportunities for readers to engage with the work
243 of other language communities. Yet prior work shows that many unique facts remain distributed across non-local
244 language editions, even for regionally associated topics [52]. As a result, Wikipedia’s current interface renders epistemic
245 differences structurally invisible to readers, making it difficult to engage with other language editions.
246

247 Together, this body of work motivates a focus on design interventions that challenge the “English-as-superset”
248 assumption by making cross-linguistic differences visible and supporting engagement across Wikipedia’s multilingual
249 peer production communities.
250

251 **2.2 Prior Work Targeting Cross-Lingual Knowledge Dissemination in Wikipedia**

252 *2.2.1 Surfacing vs. Consolidating Cross-Linguistic Differences.* Prior work on cross-linguistic variation in Wikipedia has
253 generally followed two contrasting orientations. One line of work treats differences across language editions primarily
254

261 as obstacles to information access, motivating efforts to consolidate multilingual knowledge into a single, unified
 262 representation [1, 14]. While such approaches aim to improve coverage and navigability, they implicitly assume that
 263 cross-linguistic variation should be minimized, and risk obscuring the distinct perspectives embedded in distributed
 264 knowledge production by enforcing uniform representations [53].
 265

266 In contrast, a second orientation emphasizes surfacing cross-linguistic differences as meaningful expressions of
 267 positionality [22]. From this perspective, variation across language editions is not a shortcoming to be resolved, but
 268 an expected and valuable outcome of situated knowledge production. This framing aligns closely with our earlier
 269 discussion of epistemic equity, which calls for collaborative systems to recognize and sustain diverse ways of knowing
 270 rather than subsume them under a dominant perspective.
 271

272 When translated into design terms, this orientation corresponds to an asset-based approach [23, 33, 63]. Rather
 273 than treating heterogeneity as a problem to be corrected, asset-based design views existing differences as resources
 274 for sense-making and learning. Applied to multilingual Wikipedia, this suggests that interfaces should help readers
 275 perceive and engage with cross-linguistic differences, instead of rendering them invisible through consolidation or
 276 assumed equivalence. Building on this perspective, our work extends prior efforts by exploring how interface design
 277 can surface positionality within everyday Wikipedia browsing practices.
 278

279
 280 2.2.2 *Prior Multilingual Interface Designs on Surfacing Cross-lingual Knowledge Differences.* Several systems have
 281 attempted to move beyond Wikipedia’s default ILLs by providing interfaces that surface content from multiple language
 282 editions. Manypedia [39] displays two language editions side-by-side using machine translation, but it does not
 283 explicitly highlight factual inconsistencies or information gaps between them. WikiCompare [50] takes a step further
 284 by identifying topic-level differences. When a user reads an article in one language (e.g., Hindi), the interface highlights
 285 subtopics that are discussed in another language edition (e.g., English) but missing in the current version; clicking a
 286 highlight redirects the user to the corresponding English section. While informative, WikiCompare still operates at
 287 the topic level, rather than identifying specific factual discrepancies across languages. Moreover, neither Manypedia
 288 nor WikiCompare formally evaluates how such cross-lingual surfacing affects readers’ comprehension, behavior, or
 289 information-gathering efficacy.
 290

291 The most ambitious multilingual interface to date is Omnipedia [6], which visualizes how different language commu-
 292 nities conceptualize a topic by aggregating anchor-link structures across editions. This design allows simultaneous
 293 comparison of multiple languages but comes at the cost of requiring users to adopt an entirely new interaction paradigm.
 294 Users must access Omnipedia through its dedicated search interface, which contrasts with how most people encounter
 295 Wikipedia—via search engines and in-page browsing [59]. As a result, Omnipedia’s workflow is difficult to integrate
 296 unobtrusively into readers’ existing habits on Wikipedia.
 297

298 Complementing such reader-centered systems, the WIKIPEDIA DIVERSITY OBSERVATORY (WDO) [44] offers editor-
 299 facing dashboards that quantify cultural and geographic content gaps across language editions and suggest missing but
 300 culturally relevant articles. However, WDO also operates outside Wikipedia’s native reading and editing interfaces and
 301 focuses only on article-level differences—an even coarser granularity than tools such as WikiCompare.
 302

303 In summary, prior systems demonstrate the value of surfacing cross-lingual heterogeneity, yet they share several
 304 infrastructural limitations that restrict their practical impact: (1) they primarily expose topic-level rather than sentence-
 305 level differences; (2) they rely on workflows external to the natural reading experience, limiting their ability to integrate
 306 multilingual knowledge unobtrusively. As Grudin’s classic analysis of groupware failures highlights [20], systems that
 307 impose new routines or additional work misaligned with existing practices frequently struggle to achieve sustained
 308

adoption; and (3) with the exception of Omnipedia, they lack empirical evaluation of how such designs actually affect readers.

2.2.3 *Advances in multilingual factual knowledge retrieval.* One of the challenges faced by prior systems that work at the article or topic-level was lack of access to computational tools that can recognize finer-grained differences between articles. Recent advances in multilingual NLP enable automatic detection of sentence-level factual gaps, which we leverage for developing our user-centric WIKIGAP design. Specifically, we rely on INFOGAP [52, Sec. 3.2]. The INFOGAP paper corroborated through a large-scale analysis that English Wikipedia lacks facts from other language editions, including facts that may have a broader appeal. For example, Apple CEO Tim Cook’s association with the Russia-Ukraine war is only mentioned on Russian Wikipedia, but is conceivably noteworthy to the wider array of readers that engage with English Wikipedia. Building on INFOGAP, our WIKIGAP system embeds those unique facts within Wikipedia’s native interface, overcoming limitations from the aforementioned interfaces.

3 WikiGap System

The goal of WIKIGAP is to challenge the English-as-superset assumption by making complementary knowledge from other language editions visible to readers during everyday Wikipedia use. Rather than treating multilingual differences as gaps to be corrected, WIKIGAP is designed to surface them in context, allowing readers to encounter cross-lingual knowledge while browsing English Wikipedia pages. We combine a computational pipeline that identifies sentence-level knowledge gaps across languages with a reader-facing interface that presents these differences through in-situ annotations and structured fact cards.

We begin with outlining the design requirements informed by user needs and foundational theories that shaped the WikiGap interface (§3.1). We then step behind the scenes to introduce INFOGAP [52], which we use to detect language-exclusive facts lacking in English Wikipedia (§3.2). Furthermore, we describe how we extended and adapted INFOGAP to detect and deliver language-exclusive facts, including how we filter, translate, and bind them to UI components at runtime (§3.3). We conclude this section by outlining the system implementation details (§3.4).

3.1 Designing WIKIGAP

We describe the exploratory interview we conducted to understand users' needs, from which we developed the design requirements (§3.1.1). Then, we describe the core UI elements that we derived from the requirements (§3.1.2).

3.1.1 Design Requirements Informed by User Needs and Theory. Guided by an asset-based perspective that shifts the focus from correcting gaps to making differences visible and interpretable, we combined findings from our exploratory interviews with established theories of reading and information seeking to define the design requirements for WIKIGAP. These requirements share an overarching goal: enabling readers to access cross-lingual knowledge while preserving the familiar experience of browsing English Wikipedia. The findings from the exploratory interview, requirements, and design elements are summarized in Figure 2 and detailed below.

Exploratory Interviews on Multilingual Wikipedia Use. We conducted exploratory interviews with four Computer Science graduate students (2 women, 2 men), all regular Wikipedia users; three spoke an additional language (Korean, Russian, or Hindi). Each 30-minute session followed an unstructured, conversational format guided by open-ended questions about their multilingual Wikipedia usage, perceptions of the Wikipedia ILL system, and ideas for improving cross-lingual content presentation. Participants reviewed an English Wikipedia article on *mooncake*—a topic unfamiliar

365 to them—and reflected on how multilingual facts could be surfaced in this context. They proposed interface features to
 366 better support multilingual reading habits. We took notes during the sessions and synthesized them using thematic
 367 coding to derive the following design requirements for WIKIGAP.
 368

369

370

371

372

373

374

375

376

377

378

379

380

381

382

383

384

385

386

387

388

389

390

391

392

393

394

395

396

397

398

399

400

401

402

403

404

405

406

407

408

409

410

411

412

413

414

415

416

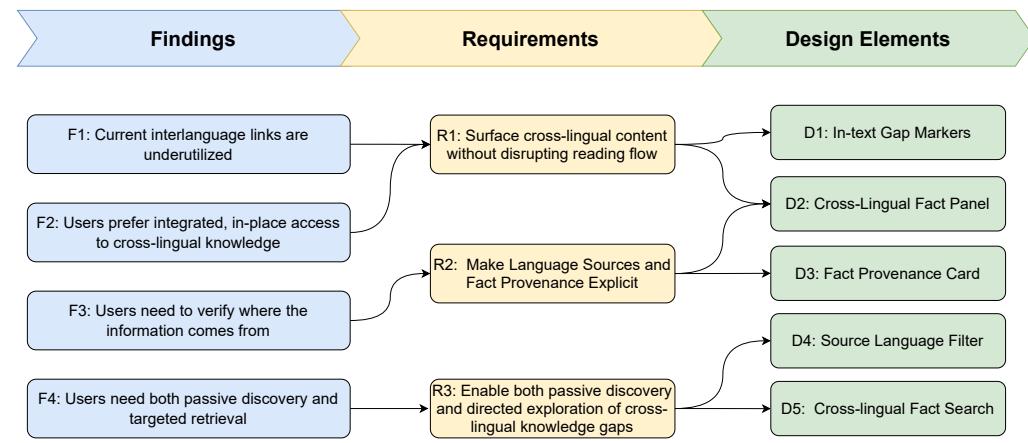


Fig. 2. How the findings from our preliminary interviews (F1-F4) informed WIKIGAP’s design requirements (R1-R3) and core design elements (D1-D5).

R1: Surface cross-lingual content without disrupting reading flow. This requirement directly addresses participants’ tendency to overlook multilingual content due to the high cost of switching language editions. Participants expressed a strong desire to access multilingual information without breaking their reading experience. They reported rarely using Wikipedia’s ILLs (F1), noting the disorienting nature of full-page switches: “*Even if I click the other language, I don’t know where to look for the thing I care about. It’s like starting over.*” Such page jumps impose high navigation costs [49], requiring users to reorient to a new layout and language, which discourages exploration.

Instead, users preferred lightweight, non-intrusive cues – similar to annotations in Grammarly or Google Docs – that subtly signal the presence of additional information (F2). This aligns with Anchored Annotation Theory [38], which finds that tying annotations to specific text locations improves comprehension and recall. Our design anchors cross-lingual facts directly to the English article using click-triggered underlines and a collapsible panel, allowing users to access additional facts only when interested. This approach follows the Spatial Contiguity Principle [29], which emphasizes placing related information close together. Collectively, these design choices preserve reading flow while enabling low-friction, in-place access to multilingual content.

R2: Ensure clear language attribution and traceability of multilingual facts to support trust and credibility. This requirement reflects readers’ need to assess the provenance of multilingual facts, especially when such information is absent from the English article. In our interviews, a recurring concern among participants was the trustworthiness of information surfaced from other language editions—particularly when such facts were not present in the English version. Participants emphasized that they wanted to know “*where the information comes from*,” including which language edition it originated from and the ability to read it in its original context (F3), because it is crucial for accessing information’s

417 credibility. This requirement highlights the importance of traceability and attribution in multilingual information access.
 418 Readers must be able to understand both the linguistic source and the original context of each surfaced fact. This
 419 entails more than just showing a translation, but rather also requires visual and interactive mechanisms that reinforce
 420 provenance and credibility.
 421

422 To meet this need, we present each fact with an explicit, color-coded language label and allow readers to trace each
 423 fact back to its source by opening the corresponding language edition and viewing the sentence in context. These
 424 strategies support transparency and accountability, giving readers confidence in the surfaced content and encouraging
 425 deeper multilingual exploration.
 426

427 **R3: Support both passive discovery and active exploration of cross-lingual knowledge gaps (customizable
 428 exploration).** This requirement reflects the need to support different levels of engagement, recognizing that readers
 429 vary in how much cross-lingual information they want to encounter. User feedback revealed a range of interaction
 430 preferences: users want flexible access modes—both passive discovery and targeted retrieval (F4). Different participants
 431 had different expectations for how cross-lingual facts should be surfaced. Some favored passive discovery, where the
 432 system automatically highlights interesting differences without requiring user input. Others preferred more active
 433 control, such as being able to search for specific content or filter information by language or topic. They discussed
 434 concern over potential information overload if too many facts were presented at once.
 435

436 This requirement reflects that readers should be able to modulate how much cross-lingual content they see and
 437 which languages they want to include in their exploration. From a cognitive standpoint, this helps reduce information
 438 overload and supports a more goal-directed behavior. It necessitates features like search, filtering, and toggling between
 439 languages – functions that empower users without requiring them to dig through entire articles in other languages.
 440

441 **3.1.2 Core Design Elements.** Based on our design process and the identified requirements (R1-R3), we developed five
 442 core design elements for WIKI GAP, as illustrated in Figure 1.
 443

444 **D1: In-text Gap Markers.** This design addresses R1 (non-disruptive presentation) and R2 (traceability) to make
 445 multilingual differences visible at the moment of reading without interrupting readers' flow, we implemented a subtle
 446 underlining system that signals the presence of additional information from other language editions. Each marker is
 447 country-color coded and appears as a patterned underline within the English article. When clicked, the side panel
 448 retrieves the fact relevant to that sentence, where the extension identifies complementary or missing information
 449 from another language edition. The underline then becomes more prominent, providing interaction feedback and
 450 inviting further exploration. Each language is associated with a distinct color – red for Chinese, blue for French, and
 451 green for Russian. This color coding is applied across both in-text highlights and corresponding elements in the side
 452 panel, reinforcing source attribution as emphasized in R2. As we elaborate in Sec. 3.3.2, only significant knowledge
 453 gaps identified by INFOGAP are highlighted, rather than minor variations or differences in wording. This selective
 454 highlighting prevents information overload and ensures that users' attention is drawn to substantive cross-lingual
 455 differences.
 456

457 **D2: Cross-Lingual Fact Panel.** To fulfill R1 and R2, we also developed a margin-anchored sidebar that displays the
 458 facts, translated to English and organized by their source language. When a user interacts with an in-text highlight,
 459 the corresponding fact card is revealed in the sidebar, which appears on the right side of the article. The sidebar is
 460 collapsible and can be pinned for extended browsing or hidden entirely to maintain a minimalist reading view. Each
 461 language section presents its facts in dedicated groups. Each fact is shown in English, accompanied by attribution
 462

469 information, including the source language tag and link to the original article. The close spatial relationship between
 470 in-text cues and sidebar content supports R1’s emphasis on minimizing reading disruption.
 471

472 **D3: Fact Provenance Card.** At the heart of WIKIGAP’s cross-lingual surfacing system is the fact card interface, which
 473 presents individual facts from other language editions in a compact, standardized format. Each fact card includes: (1) the
 474 target fact, translated into English, (2) a color-coded badge indicating the source language, and (3) a hyperlink labeled
 475 “View on [Language] Wikipedia” that opens the source article in a new tab with the original sentence highlighted. This
 476 design primarily addresses R2 (transparent attribution) by balancing at-a-glance comprehension with direct access to
 477 original context.
 478

479 **D4: Source Language Filter.** To support R3 (customizable exploration), we introduced a language filter function
 480 located at the top of the sidebar. These filters allow users to specify which language(s) they want to view content from,
 481 enabling personalization based on their linguistic background and interests. The filters use both language names and
 482 colors to maximize clarity. This functionality directly supports both passive and active interaction modes (R3).
 483

484 **D5: Cross-lingual Fact Search.** To further support R3, specifically the directed exploration function, we added a
 485 search box feature that allows users to perform keyword-based queries across multilingual facts. When a user enters a
 486 search term, relevant fact cards from the selected languages are retrieved and displayed in the sidebar. This feature
 487 caters to users with goal-oriented information-seeking tasks, such as verifying a specific claim or comparing facts across
 488 language editions. While D1 and D2 emphasize passive and contextual discovery, the search box fulfills a complementary
 489 role by enabling intentional, focused retrieval. Together, these elements provide a spectrum of access modes that reflect
 490 users’ varying preferences for how and when to engage with multilingual content.
 491

492 3.2 The INFOGAP Pipeline

493 While the interface makes multilingual differences visible to readers, INFOGAP provides the backend capability that
 494 enables such differences to be identified at sentence-level granularity. It was a previously developed computational
 495 pipeline introduced by Samir et al. [52], which detects factual misalignments between Wikipedia articles across
 496 languages. This subsection introduces the core components of the original INFOGAP system, which we use to identify
 497 sentence-level knowledge gaps between English and other language editions. In the following section, we will describe
 498 how we extended this pipeline to support additional languages and adapted its outputs for integration into the WIKIGAP
 499 interface.
 500

501 INFOGAP is a state-of-the-art LLM-based pipeline that takes an article L_s in the *source language* and the respective
 502 article L_t in the *target language* and returns the sets of common facts and facts that are exclusive to one article. As
 503 illustrated in Figure 3, given a Wikipedia article in the source language (L_s , e.g., English) and the respective article in
 504 another language (L_t , e.g., French), INFOGAP identifies three categories of facts: those shared across editions, those unique
 505 to the source, and those unique to the target (i.e., language-exclusive facts). For example, “*Oolong is a semi-oxidized*
 506 *Chinese tea*” is shared across languages, while “*It is served in US restaurants*” may appear only in English.
 507

508 The pipeline leverages a multilingual LLM and sentence embeddings to perform cross-lingual comparison in three
 509 stages, as illustrated in Figure 4:
 510

511 (1) **Fact decomposition.** Each paragraph in L_s and L_t is decomposed into atomic factual statements using prompts
 512 issued to a multilingual LLM. The output includes both the fact and its paragraph index for downstream
 513 alignment.
 514



542
 543
 544

- (2) **Multilingual alignment.** Each fact is encoded using a multilingual sentence encoder. For every fact in L_t , the top three nearest neighbors in L_s are retrieved based on cosine similarity. This reduces noise and narrows the search space when looking for equivalent facts.
- (3) **Alignment verification.** For each fact in L_t , the LLM is prompted to determine whether it is inferable from any of its nearest neighbors in L_s . If a matching fact is found in L_s , these facts are considered aligned, otherwise, the fact is labeled as a *knowledge gap* in L_s .

554 3.3 Extending and Integrating INFOGAP for WIKIGAP

555 In this section, we outline the conceptual modifications and integration steps that enable INFOGAP to serve as the
 556 backend for WIKIGAP. We adapted and extended the existing INFOGAP pipeline [52] in two key ways: (1) by enabling
 557 new language (Chinese) support beyond the original study, and (2) by building an integration layer that transforms
 558 INFOGAP’s output into interactive interface components usable within our system.

559
 560
 561
 562 3.3.1 *Rationale for Language Selection.* In this study, we intentionally focused on English and three additional language
 563 editions, namely Russian, French, and Chinese. We chose the first two because the original INFOGAP pipeline was
 564 developed for detecting cross-lingual gaps in biographical articles between English and either Russian or French. We
 565 included them because they have been validated in the research to ensure accuracy. To generalize the method to a new
 566 domain and an additional language, we incorporated support for Chinese into the fact decomposition and verification
 567 modules, as some co-authors were proficient in Chinese. This involved custom preprocessing for tokenization, crafting
 568 prompts, and conducting human evaluation on Chinese texts to confirm that INFOGAP performs comparably to its
 569 performance on the original languages studied by Samir et al. [52].

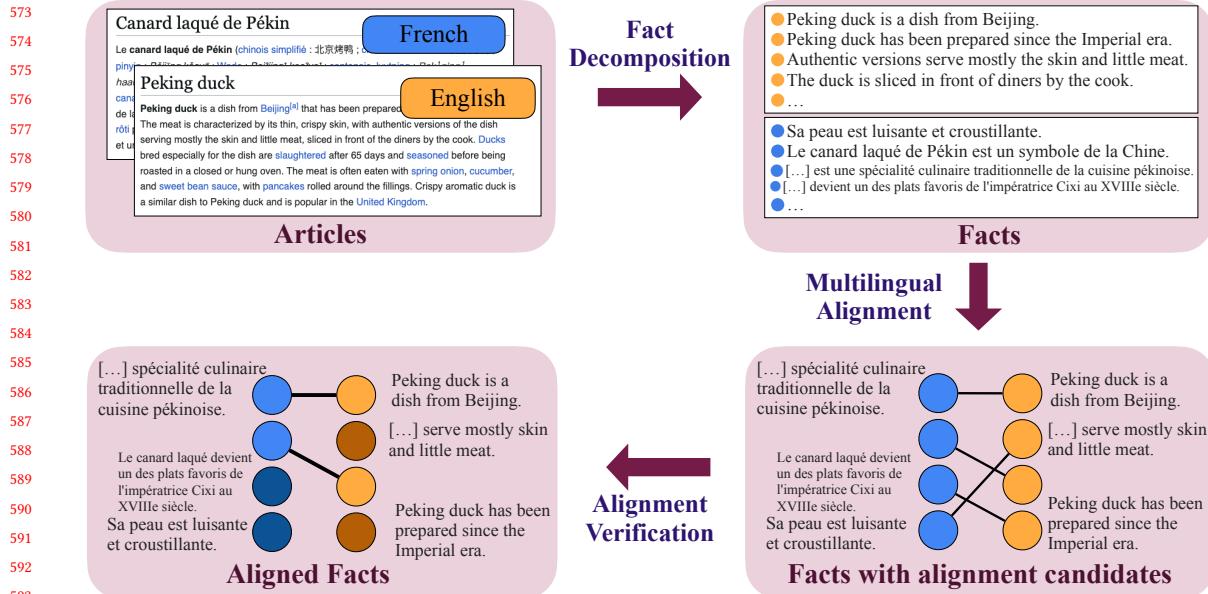


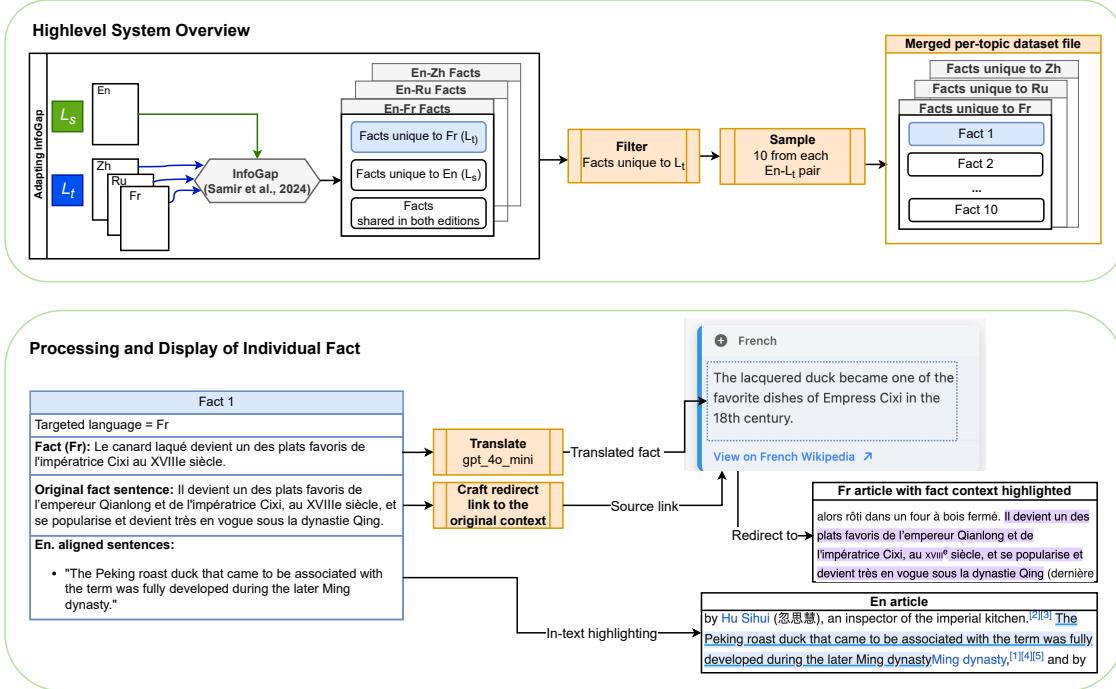
Fig. 4. Overview of the INFOGAP backend pipeline for cross-lingual fact alignment, reproduced from Samir et al. [52] and adapted. For additional technical details, see the original paper.

3.3.2 *Language attribution and fact selection.* Our focus is on *knowledge gaps* – facts that appear in the L_t version (e.g., French, Russian, or Chinese) of an article but *not* in L_s (English). We therefore filter the INFOGAP output to retain only gap facts. To ensure interface responsiveness, we use a precomputed output from INFOGAP for each article topic. For every topic, we prepare a unified dataset that includes exclusive facts from three target languages (French, Russian, and Chinese). These facts are grouped by language code (fr, zh, ru) to support language-specific interaction features such as language filtering (D4). To reduce cognitive overload, we limit the number of displayed facts to **10** per language, proportionally sampling based on section-level gap counts. If fewer than 10 gaps exist for a language, we display all available facts.

3.3.3 *Translations and traceability.* Because WIKIGAP displays cross-lingual content within the English Wikipedia page to readers who we do not assume read other languages, we translate the gap facts from their original language (L_t) into English for presentation in the fact cards. To support source traceability and deeper exploration, we create a direct *link-to-highlight* for each fact. An encoded version of the original sentence in L_t is appended to the target article's URL, sending readers to the exact sentence in context when they click on the "View on [L_t] Wikipedia" in the card (D3).

3.3.4 *In-text highlights.* For each gap fact, INFOGAP outputs the most semantically related English sentence from the source article, even when the fact itself is absent in English. We use these sentences as anchors for subtle in-text highlights (D1), which allow readers to preview cross-lingual differences inline without disrupting their reading flow. Figure 5 illustrates how such outputs map onto WIKIGAP's interface components. For example, the French Wikipedia article for *Peking Duck* mentions that the dish became a favorite of Empress Cixi in the 18th century – a detail missing from the English article. INFOGAP aligns this fact to the closest English sentence "The Peking roast duck that came to be

625 associated with the term was fully developed during the later Ming dynasty", which provides a semantically-relevant
 626 anchor for surfacing the French-exclusive fact in context.
 627



654 **Fig. 5. Overview of system implementation and data processing pipeline.** *Top:* A high-level overview of the data stream in the
 655 WIKIGAP system. We adapted the INFOGAP pipeline to support Chinese-language input alongside existing language pairs, followed
 656 by post-processing steps to standardize and merge datasets by topic. Orange process blocks indicate components we developed to
 657 enable proper integration and display of multilingual facts in the UI. *Bottom:* The data structure and rendering flow for an individual
 658 fact. This illustrates how each multilingual fact is transformed – through translation, alignment, tagging, and contextual linking –
 659 into an interactive component in the WIKIGAP interface.

662 3.4 System Implementation

664 We implemented WIKIGAP as a browser-based system that overlays multilingual content directly onto English Wikipedia
 665 articles. The system interface was developed as a Chrome extension using standard web technologies – HTML, CSS, and
 666 vanilla JavaScript – while the underlying data pipeline generates and serves precomputed JSON files derived from the
 667 INFOGAP framework. When running INFOGAP, we used GPT-4o for fact decomposition and fact alignment verification,
 668 and computed sentence embeddings using the multilingual LaBSE model [15, 48]. For each English Wikipedia topic,
 669 INFOGAP produces three separate datasets, each containing factual gaps between English and one of three target
 670 languages (fr, zh, ru). We then merge these files by topic into a unified, standardized JSON file named after the English
 671 article. This consolidation supports runtime efficiency: when a user visits an English topic page, the extension loads
 672 the corresponding file based on the topic title. Each file contains multilingual facts translated into English which are
 673 generated using GPT-4o-mini [47]. The overall system workflow is illustrated in Figure 5.

677 Table 1. Description of study metrics and corresponding data collection and analysis. All metrics were collected at the individual
 678 participant level and compared across WIKIGAP and control conditions.

680 Metrics	681 Data Collected	682 Data Analysis
683 Performance		
684 Quiz Accuracy	685 Percentage of correct answers in each condition	686 Paired t-test comparing WIKIGAP vs. control conditions
Completion Time	687 Time (in minutes) to complete each quiz session	688
689 Preference		
690 Usability	691 System Usability Scale (SUS) score in each condi- 692 tion	693 Paired t-test comparing SUS scores between conditions
694 User Feedback	695 Open-ended responses and interview transcripts	696 Thematic coding to extract perceived usability, utility, and 697 challenges
698		
699		

700 The extension interacts with the live English Wikipedia page, dynamically injecting highlights and sidebar content
 701 based on the current topic the user is browsing (the corresponding dataset is selected for rendering the data on the UI
 702 based on that topic). For each fact card, we generate external links that redirect readers to the original sentence in the
 703 source-language Wikipedia article. This is done by appending an encoded version of the original sentence to the article's
 704 base URL, allowing users to jump directly to the fact in context. Additionally, the search function is implemented by
 705 scanning both the fact body and language label for matches with the user's input and dynamically hiding non-matching
 706 entries.

707 **4 Evaluation Methods**

708 To examine whether surfacing cross-lingual knowledge gaps within the English Wikipedia interface can meaningfully
 709 support multilingual engagement, we conducted a mixed-methods user study combining quantitative performance
 710 measures with qualitative feedback. The study assesses usability, knowledge acquisition, and users' perceptions of the
 711 system and the cross-lingual content it surfaces. Table 1 summarizes the study metrics and analysis.³

712 **4.1 Participants**

713 Twenty-one people participated in the study. Of these, eight were student volunteers from the HCI course where this
 714 project originated, and the remaining thirteen participants were recruited through Upwork, and were compensated
 715 at a rate of 20 USD for the user study. Participants ranged in age from 18 to 44 and represented diverse cultural and
 716 ethnic backgrounds (see Table 2). All participants reported that they typically use the English version of Wikipedia.
 717 Twelve had previously contributed to Wikipedia to varying degrees, with two considering themselves frequent editors.
 718 Participants were required to meet the following inclusion criteria: (i) regular use of English Wikipedia, (ii) comfort
 719 with using a Chrome browser, and (iii) willingness to contribute user experience feedback.

720 The study followed a within-subject design, in which each participant was assigned two different topics out of the
 721 total five topics (Sec. 4.3). Each participant read one article using the WIKIGAP extension and another article using
 722 the default Wikipedia interface (with ILLs) as a control. We selected interlanguage links (ILLs) as the control to reflect
 723 the default, in-context Wikipedia experience and to simulate an extension built on top of existing knowledge-seeking
 724 workflows; moreover, prior multilingual interface designs are not publicly available for deployment. The assignment of
 725 treatment and control conditions was randomized for each participant.

726 ³The study protocol received approval from our institutional REB.

729
730
731 Table 2. Participant self-reported cultural background.
732
733
734
735
736
737
738
739

Cultural / Ethnic background	% of participants
East Asian	24%
White or European	19%
South Asian	14%
Hispanic or Latin American	14%
Black or African	10%
Middle Eastern or North Southeast Asian	10%
African	5%
Other	5%

740
741
742
743
744 4.2 Tasks
745

746 Participants completed an open-book reading quiz consisting of 10 multiple-choice questions in each condition. For each
 747 quiz, participants were provided with four Wikipedia articles on the same topic—specifically, the English article and its
 748 corresponding versions in French, Russian, and Chinese. They were instructed to rely solely on these four articles to
 749 answer the quiz questions and were not allowed to use external resources such as Google Search. Each participant
 750 completed the quiz twice: once using the standard Wikipedia interface (control condition) and once using the WIKIGAP
 751 extension (treatment condition). For every question, participants were also asked to indicate the language edition in
 752 which they found the answer.

753
 754 *Quiz Construction.* We constructed one quiz per article topic using facts extracted by INFOGAP and surfaced through
 755 the WIKIGAP interface. From the 30 multilingual facts identified for each article, we manually selected a subset of 10,
 756 ensuring a roughly equal number of facts from each of the three language editions—French (fr), Russian (ru), and Chinese
 757 (zh)—to maintain balanced representation. For a fair comparison, we ensured that all quiz questions are answerable
 758 from InfoGap facts and by extension from the respective articles in French, Russian, and Chinese.

759
 760 We then prompted a large language model GPT-4o [48] to generate a multiple-choice question with four answer
 761 choices from each fact. The questions were manually reviewed and edited to ensure clarity, factual correctness, and
 762 alignment with the presented content. In addition to selecting the correct answer, participants were asked to indicate
 763 the language edition in which they found the supporting fact. Although all quiz facts were sourced from non-English
 764 Wikipedia editions, we note that some participants may have inferred the correct answer from contextual clues in the
 765 English article. Nonetheless, the quizzes were designed to evaluate whether participants engaged with multilingual
 766 content. If participants had guessed entirely at random, the expected accuracy would be 25%, given the four-choice
 767 format. Example questions are included in the Appendix A.

768
 769 *Performance Metrics.* We recorded each participant’s task completion time and quiz accuracy score under both
 770 conditions. Quiz accuracy captured how much multilingual knowledge participants were able to obtain, while completion
 771 time reflected how quickly they could locate and interpret the relevant information. Taken together, these two measures
 772 provide an indication of learning efficiency and allowed us to quantitatively compare the effectiveness of WIKIGAP
 773 with the standard Wikipedia interface.

781 Table 3. Number of knowledge gap facts discovered per food topic and language by INFOGAP, and the total number of facts shown
 782 in the WikiGap interface. For each topic, up to 10 facts per language were selected to ensure balanced representation and reduce
 783 cognitive overload. *Injera* had only 8 gaps in Chinese, resulting in 28 total facts shown (*).

Food Topic	Russian (ru)	French (fr)	Chinese (zh)	Facts Shown in WikiGap
Wiener schnitzel	65	20	62	30
Peking duck	13	23	69	30
Paella	28	62	52	30
Philippine adobo	21	15	84	30
Injera	10	14	8	28*

4.3 Materials and Topic Assignment

794 *Selected topics.* We focus on food articles because food serves as a culturally rich lens through which to explore
 795 knowledge gaps across language editions in Wikipedia. Culinary practices are deeply embedded in regional customs,
 796 social histories, and national identities, making food a meaningful proxy for cultural knowledge [36, 62]. Analyzing
 797 food allows researchers to capture both shared and divergent cultural representations, such as differences in dish
 798 names, preparation styles, and regional associations [62]. This cultural complexity makes food an ideal domain for
 799 surfacing factual asymmetries. We selected five culturally specific dishes: *Injera* (Ethiopia), *Paella* (Spain), *Philippine*
 800 *Adobo* (Philippines), *Peking Duck* (China), and *Wiener Schnitzel* (Austria). Each of these dishes has a dedicated Wikipedia
 801 article and distinct cultural heritage. Table 3 summarizes the number of language-exclusive facts detected per topic and
 802 the total number shown in the WIKIGAP interface.

803 All five dishes satisfy the following practical constraints: (i) they originate outside the English-speaking world,
 804 providing a comparable degree of cultural distance from the English source language, (ii) their English Wikipedia article
 805 is moderately sized (1,000–2,000 words), and (iii) the WIKIGAP extension can surface ~30 multilingual facts for each.⁴

806 *Topic assignment.* Each participant is assigned two distinct topics, one for each condition (control and WIKIGAP). Topic
 807 assignment was designed to minimize cultural confounds and familiarity-driven performance differences. Specifically,
 808 the food topic chosen for each participant could not be strongly linked to the participant’s self-reported cultural or
 809 ethnic background. To reduce individual-performance bias, each topic was assigned in every condition to **exactly four**
 810 different participants.

4.4 Procedure

811 The lead investigator conducted the study remotely over a 1–1.5 hour recorded Zoom call. At the start, participants
 812 received a zip file containing the WIKIGAP extension and were instructed to install it on their Google Chrome browser.
 813 A short tutorial was provided by the researcher for both the treatment and control conditions. During the treatment
 814 condition, we began the round by introducing the WIKIGAP’s functionality and allowed participants to explore the
 815 extension on their own until they felt ready to proceed. In the control condition, we demonstrated how to access
 816 language versions on Wikipedia using ILLs, and how to use Chrome’s Google Translate function to read the content
 817 in English. At the beginning of each round, one of the five topics was assigned to the participant according to the
 818 topic assignment rules. After completing each round, the participant completed a System Usability Scale (SUS) after
 819 finishing the quiz. We used the SUS to measure ease-of-use and users’ preference on with or without the WIKIGAP

820 ⁴*Oolong tea* was included in the pilot with the first two participants but was dropped because its cultural context overlapped substantially with *Peking*
 821 *Duck*. It was therefore replaced with *Philippine Adobo* to preserve cultural diversity in topic representation.

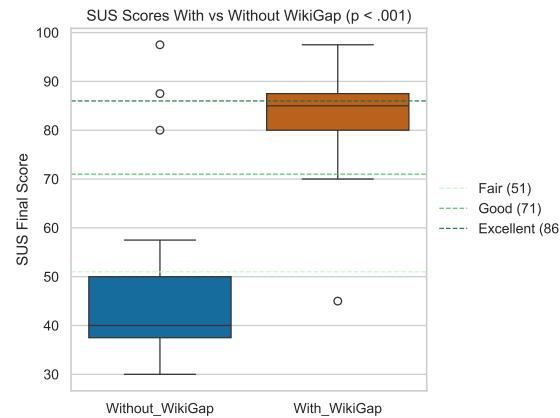


Fig. 6. Box plot showing the System Usability Scale (SUS) scores for each condition. Blue represents the control condition (no WIKIGAP), and orange represents the treatment condition (with WIKIGAP). Horizontal dashed lines represent standard usability benchmarks in varying shades of green: light green for *Fair* usability (SUS > 51), medium green for *Good* usability (SUS > 71), and dark green for *Excellent* usability (SUS > 86).

extension. To conclude, the first twelve participants completed an open-ended questionnaire, while the remaining nine took part in a semi-structured interview. Both instruments covered the same six areas: overall impressions of the system, presentation of multilingual information, perceived cultural impact, trust in the augmented facts, perceptions of Wikipedia's reliability after using WIKIGAP, and future inclination to explore multilingual content (see Appendix B for the full list of questions).

5 Findings

We report findings from our mixed-method evaluation of WIKIGAP. Together, these findings show how making multilingual knowledge differences visible within the English Wikipedia interface affects users' performance, preferences, and engagement with non-English content. Across quantitative measures and qualitative feedback, participants consistently described shifts in how they accessed information, perceived Wikipedia's completeness, and engaged with multilingual knowledge during routine fact-finding.

1. Making Multilingual Differences Visible Improves User Preference and Reading Experience (D1, D2, D3).

We evaluated user preference using the System Usability Scale (SUS). The average SUS scores for each condition are presented in Figure 6, along with general SUS cutoff scores for fair, good, and excellent usability [5]. Overall, WIKIGAP achieved a substantially higher usability score compared to the control. WIKIGAP falls within the "excellent" usability category, while the default Wikipedia page with ILL scored below the "fair" threshold. A paired-samples *t*-test revealed a significant difference in scores between the WIKIGAP condition ($M = 82.5$, $SD = 10.9$) and the control condition ($M = 48.2$, $SD = 18.6$), $t(20) = 6.60$, $p < .001$. (Figure 6).

Participants attributed their strong preference for WIKIGAP not only to its ability to surface knowledge disparities, but also to the various design elements (D1-D3) that supported the quick access and ease of reading. Many highlighted how having access to multilingual content directly within the English article – through translated facts in a sidebar (D2, D3) – lowered the barrier to engaging with other language editions. It helped users "absorb information without the language difference obstacle" (P8). Participants appreciated the sidebar's structure and organization by language,

which kept all cross-lingual facts “apparent within a single page” (P5), eliminating the need to navigate away and reducing friction in multilingual exploration. In contrast to ILL’s full-page switch model, participants noted that the **sentence-level fact card interface (D3)** made the information easier to process. As P11 remarked, “[WikiGap] presents sentences instead of long paragraphs from the Wikipedia page.”

Other visual elements also shaped participants’ perceptions of usability. Several users found the **in-text gap markers (D1)** helpful in surfacing areas of missing information in the English article. However, reactions to the color scheme used for language cues were mixed. While some found the colors intuitive, others, like P19, felt that the red-green coding could unintentionally suggest “right or wrong”.

Despite minor confusion about the color scheme, most users described the overall experience as more enjoyable and cognitively less demanding than their usual experience with Wikipedia’s default ILL interface. P9 summarized this sentiment: “*Sometimes even when I open my Wikipedia, I would not read that much because of laziness, but WIKIGAP did a good job extracting information and makes the whole experience much more pleasant.*”

2. Visibility of Cross-Lingual Differences Improves Fact-Finding Performance and Supports Cross-Cultural Learning (D4, D5). To evaluate the impact of WIKIGAP on users’ ability to retrieve and retain cross-lingual information, we measured task performance through a fact-finding quiz, using both accuracy and completion time as metrics. As shown in Figure 7, participants performed significantly better when using WIKIGAP, achieving higher quiz accuracy and completing the task more quickly than with the default Wikipedia interface. A paired-samples *t*-test showed that accuracy was significantly higher in the WIKIGAP condition ($M = 0.91$, $SD = 0.09$) than in the control condition ($M = 0.73$, $SD = 0.19$), $t(20) = 4.75$, $p < .001$. Completion time also improved significantly, with users finishing the quiz faster using WIKIGAP ($M = 12.39$ minutes, $SD = 5.78$) compared to the control ($M = 20.84$ minutes, $SD = 8.68$), $t(20) = -5.52$, $p < .001$.⁵ These results suggest that WIKIGAP not only facilitated access to relevant information but also improved users’ efficiency in locating and interpreting cross-lingual facts. WIKIGAP consistently improved user performance across all five culturally diverse topics. For a topic-by-topic breakdown of quiz accuracy and completion time, see Appendix Figure 8.

These performance gains were reinforced by participants’ qualitative feedback, which highlighted the value of both passive discovery and active exploration features. **The fact search (D5)** enabled users to directly retrieve relevant facts, supporting goal-oriented behaviors and speeding up information retrieval. As P16 explained, “*I don’t need to go to multiple pages. I can search in the WIKIGAP extension and find any information quickly. Using the WIKIGAP extension saved me a lot of time.*” This targeted retrieval helped users answer questions more efficiently, contributing to the improved fact-finding performance observed in the quiz metrics. Complementing this, the **language filter controls (D4)** allowed users to personalize their exploration by prioritizing facts from languages culturally aligned with the topic. Several participants reported that they first filtered to the language they deemed most relevant, as they “*will trust more when the language is directly aligned with the original country of the food.*” (P1). Similarly, P7 remarked, “*I would love to see how the Tagalog page talks about Philippine adobo. In fact, diverse cultural perspectives on food is more linked to the cultural ecosystem where the food is grounded.*”. Together, D4 and D5 supported efficient, goal-driven access to cross-lingual content.

Beyond performance gains, the extension also facilitated passive discovery and cross-cultural learning – a central goal of our design. Several users remarked that they encountered information they would not have actively sought

⁵Faster completion times in the WIKIGAP condition is partly expected, as the quiz questions were constructed from facts extracted by INFOGAP. In real-world use, not all facts would be presented by WIKIGAP, and completion time gains may therefore be more modest.

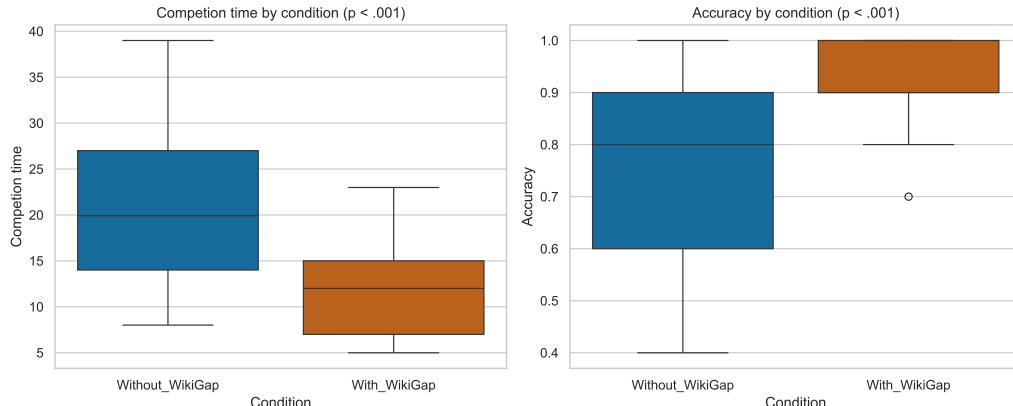


Fig. 7. **Average completion time** (left) and **quiz accuracy** (right) across conditions. Blue represents the control condition (no WIKIGAP), and orange represents the treatment condition (with WIKIGAP). Using WIKIGAP demonstrates a statistically significant improvement in user performance on the quiz.

out on their own. As P1 reflected, “*Without [WIKIGAP], I wouldn’t have thought about looking into Wikipedia pages in different languages.*” Others described moments of unexpected cultural insight, such as P5, who noted, “*I would have completely missed out on the historical aspect if I used Wikipedia normally (only in English).*” Participants also pointed to surprising culture-specific facts, like the culinary use of Sprite in a traditional dish (P13). These examples illustrate that making multilingual differences visible not only improved task efficiency, but also broadened users’ exposure to culturally situated knowledge.

3. WIKIGAP Raises Awareness of Multilingual Gaps and Motivates Exploration. Participants consistently credited WIKIGAP for revealing factual disparities they were previously unaware of. Several users described being surprised that different language versions did not simply mirror each other. As P20 reflected, “*Before WIKIGAP, I always assumed that everything presented in English would also be visible in another language.*” Similarly, P18 noted that they “*didn’t know there could be so many differences in how the content was presented in various languages.*”

This realization often prompted frustration as participants became aware of the limits of English Wikipedia. As P5 expressed, “*I’m frustrated that this doesn’t already exist! I feel like I’ve been missing out on information from other languages because I just assumed other pages would have the same information translated.*” Other participants had deeper reflections on how WIKIGAP changed their perception about knowledge: “*Knowledge is not plain, knowledge has different dimensions waiting to be discovered, and this tool offers this opportune moment.*” (P10).

Importantly, increased awareness also motivated future engagement. Participants described being more inclined to explore content outside their primary language now that differences were visible and easy to access.

“I was never on the pages of Wikipedia where the language is not what I speak, because I often assume that all pages under different languages are the same if they are of the same topic, so sometimes I struggled to find some information in the Wikipedia. But if I have WIKIGAP, I think I will definitely explore more.” — P9

Several emphasized that centralizing multilingual content within a single interface reduced the effort required to engage, making exploration more likely. Now they can “*see different viewpoints from different language pages*” (P17), they are more inclined to explore content outside their primary language in the future. Several users described feeling

more motivated to engage with multilingual content now that they understood its value. As P1 noted, “*If there’s a tool like WIKIGAP that can help me easily get various information from other pages, I would like to explore. If not, I wouldn’t navigate multiple pages myself.*”

4. Respondents’ Views on LLMs and Wikipedia: Perceived Trust and Tensions. Awareness of cross-lingual disparities also led participants to reflect on the trustworthiness of Wikipedia and how it compares to LLM-based knowledge sources. We observed a wide range of perspectives. For some, WIKIGAP surfaced surprising inconsistencies between language editions. Those differences made them question the completeness or neutrality of English Wikipedia. Others interpreted these variations not as flaws, but as a feature of Wikipedia’s multilingualism and editorial diversity. These divergent views led participants to compare Wikipedia to emerging alternatives such as LLMs, positioning WIKIGAP as a catalyst for rethinking how people evaluate and engage with knowledge sources.

A minority (n=2) of the participants reported a decline in trust in Wikipedia after using WIKIGAP. They described how uncovering missing or differing facts between language editions made them feel less confident in the platform’s ability to represent comprehensive or balanced information. This, in turn, increased their inclination to rely on LLMs as knowledge sources:

“I feel that after using WIKIGAP, I would be inclined to use Wikipedia even less. Before this, I assumed that Wikipedia would be a reliable and complete source of knowledge. However, I now know that the English version misses several key details. This leads me to be more reliant on LLM-based tools that have all the information in one place.” — P2

For this participant and others, WIKIGAP altered their perception of Wikipedia as a trustworthy authority. The perceived comprehensiveness of LLMs, combined with their ability to synthesize content across sources, positioned them as a more efficient alternative, even if they lacked transparency.⁶

However, not all participants responded this way. Others voiced strong concerns about the lack of attribution and provenance in LLM-generated responses, and emphasized that WIKIGAP made them appreciate Wikipedia’s editorial structure and citation practices even more. For these users, the ability to trace facts back to their original source, including reading the sentence in the original language edition, was critical to their sense of trust:

“I don’t trust ChatGPT that much because sometimes it just makes things up, and it doesn’t provide sources. With WIKIGAP, I can click on, say, the French Wikipedia page, read it, and see the sources. It feels more reliable. I also appreciate that Wikipedia cites everything, so there’s some quality control, even if it’s not always perfect.” — P13

These contrasting responses underscore an important dynamic in how users engage with knowledge platforms: while LLMs offer convenience and perceived comprehensiveness, they can lack the transparency and editorial accountability that are central to Wikipedia’s epistemic model. For some users, WIKIGAP undermined their confidence in Wikipedia. For others, it reaffirmed the value of its provenance-driven structure, even in the face of inconsistency. We return to the socio-technical implications in Sec. 6.

5. WIKIGAP Suggests Opportunities for Supporting Cross-Lingual Editing. Although WIKIGAP was primarily designed to support readers by surfacing factual cross-lingual content, a few participants, particularly those with prior

⁶In practice, the assumption that multilingual LLMs such as ChatGPT and Gemini can synthesize information learned from web text in different languages is not entirely true; they often fail to retrieve facts learned in one language when prompted in another language [17].

1041 editing experience, reflected on how it could assist with editorial tasks. These insights point to an emergent opportunity
1042 for WIKIGAP to support not only knowledge consumption but also knowledge curation on Wikipedia.
1043

1044 Some users noted that the ability to trace and verify information from other language editions made the tool useful
1045 for improving article quality. As one frequent editor explained, “*Sometimes an article in another language has crucial*
1046 *information that’s missing in English. WIKIGAP helps me find and translate it, then cite it properly.*” (P16). Another
1047 participant elaborated on how the tool fits into their existing editing workflow: “*I’d check the source WIKIGAP shows*
1048 *me, then cite the same source... WIKIGAP helps me discover facts and references.*” (P17). While these reflections were
1049 not prompted by editor-specific questions, they point to an emergent opportunity: making multilingual gaps visible
1050 to readers may also lower barriers to cross-lingual knowledge integration, extending WIKIGAP’s impact beyond
1051 consumption to potential curation.
1052

1053 6 Discussion

1054 We interpret our findings through the lens of *boundary objects*—shared artifacts that enable coordination across hetero-
1055 geneous communities without requiring consensus or uniform interpretation [56, 57]. We argue that WIKIGAP functions
1056 as a boundary object within Wikipedia’s multilingual ecosystem by making cross-lingual knowledge differences legible,
1057 actionable, and traceable within everyday reading practices. In this role, WIKIGAP enables three intertwined forms of
1058 epistemic and coordinative work. First, it reframes readers’ epistemic assumptions by exposing the incompleteness of
1059 English Wikipedia and foregrounding knowledge produced by other language communities (§6.1). Second, it provides a
1060 shared reference point that supports articulation work between readers and editors across language editions, enabling
1061 verification and coordination without enforcing uniformity (§6.2). Third, it gestures toward a broader infrastructural
1062 role by modeling how pluralistic, provenance-aware interfaces can counter knowledge consolidation in emerging
1063 AI-mediated systems (§6.3).
1064

1065 6.1 Promoting Epistemic Equity by Challenging the “English-as-Superset” Assumption

1066 Our user study revisits prior findings that the dominance of English Wikipedia leads readers to treat it as a comprehensive
1067 superset of knowledge across language editions, rather than as a culturally situated and socio-geographically positioned
1068 construction [34]. Participants explicitly articulated this belief. Although more than 60% of participants reported a
1069 non-English primary language, all identified English as their primary source on Wikipedia, largely because English
1070 Wikipedia covers more topics and is updated more frequently. Despite being familiar with Wikipedia’s ILLs, participants
1071 predominantly consumed content in English, reflecting how infrastructural dominance concentrates reader attention
1072 within a single language edition and reinforces the invisibility of knowledge produced by other language communities.
1073 At the interface level, Wikipedia’s current design renders other language editions technically accessible but epistemically
1074 invisible: ILLs signal the existence of alternative versions, yet do not communicate that English content is neither
1075 exhaustive nor neutral. This gap was evident in participants’ assumptions; as one remarked, “*Before WikiGap, I always*
1076 *assumed that everything presented in English would also be visible in another language.*”
1077

1078 Consistent with this perception, participants strongly preferred WIKIGAP over the default ILLs interface and per-
1079 formed significantly better on fact-finding tasks, indicating that access alone—without cues about epistemic differences,
1080 provenance or community context—is insufficient to support meaningful cross-lingual engagement. WIKIGAP challenges
1081 the “English-as-superset” assumption by both raising awareness of cross-linguistic knowledge gaps and redistributing
1082

¹⁰⁹³ reader attention toward knowledge produced by other language communities. Adopting an asset-based design orientation,
¹⁰⁹⁴ the system foregrounds non-English content not as supplementary or corrective, but as a valuable product of
¹⁰⁹⁵ distinct peer production efforts.
¹⁰⁹⁶

¹⁰⁹⁷ By embedding sentence-level complementary facts in context and pairing them with explicit provenance, WIKIGAP
¹⁰⁹⁸ makes the limits of English Wikipedia visible and simultaneously inviting engagement with the work of other language
¹⁰⁹⁹ editions. This design encourages readers to recognize Wikipedia as a collection of interconnected yet distinct knowledge
¹¹⁰⁰ production communities, rather than a single monolithic source of global consensus. In doing so, WIKIGAP promotes
¹¹⁰¹ epistemic equity in a scoped, design-operational sense: by redistributing attention and engagement across Wikipedia's
¹¹⁰² multilingual peer production communities and supporting cross-community interaction as part of everyday readership.
¹¹⁰³
¹¹⁰⁴

¹¹⁰⁵ 6.2 WikiGap as a Boundary Object Supporting Articulation Work Across Reader and Editor Communities

¹¹⁰⁶ Wikipedia's multilingual knowledge production spans heterogeneous communities with distinct linguistic, cultural, and
¹¹⁰⁷ editorial norms. Our findings show that WIKIGAP operates as a boundary object in practice, providing a shared point of
¹¹⁰⁸ reference through which readers and editors from different language communities can interpret, evaluate, and act on
¹¹⁰⁹ cross-lingual knowledge differences without requiring consensus. By surfacing sentence-level complementary facts
¹¹¹⁰ with clear provenance, WIKIGAP offers a way to support shared understanding across communities without requiring
¹¹¹¹ consensus or uniformity. Building on these findings, we interpret WIKIGAP as enabling an asset-based stance: it helps
¹¹¹² readers and multilingual contributors recognize non-English facts as usable resources by making them legible and
¹¹¹³ actionable through provenance and in-context integration.
¹¹¹⁴
¹¹¹⁵

¹¹¹⁶
¹¹¹⁷ *6.2.1 For Readers of English Wikipedia.* WIKIGAP benefits readers by supporting fact-finding, cross-cultural learning,
¹¹¹⁸ and interpretive flexibility. Participants described discovering content they "would have otherwise missed" in English
¹¹¹⁹ and expressed surprise at how much meaningful information existed exclusively in other language editions (Findings 1–
¹¹²⁰ 2). Several noted that the tool increased their curiosity and willingness to visit other language pages, shifting their sense
¹¹²¹ of Wikipedia from a monolithic source of global consensus to a collection of situated, culturally inflected knowledge
¹¹²² communities.
¹¹²³
¹¹²⁴

¹¹²⁵ Through a CSCW lens, this suggests that WIKIGAP enhances the *legibility* of multilingual knowledge work. It
¹¹²⁶ exposes facts that English readers typically never encounter, foregrounds the situated perspectives of marginalized
¹¹²⁷ linguistic communities, and supports interpretive flexibility—one of the core affordances of boundary objects. Rather
¹¹²⁸ than collapsing perspectives into a single authoritative narrative, WIKIGAP allows readers to compare, contextualize,
¹¹²⁹ and reinterpret information based on their own background knowledge and cultural standpoint.
¹¹³⁰
¹¹³¹

¹¹³² *6.2.2 For (Potential) Editors Across Languages.* Although our study focused on readers, participants with editing
¹¹³³ experience highlighted how WIKIGAP could support multilingual editorial workflows, they pointed to using the
¹¹³⁴ provenance link to see the original facts from the other languages to verify and potentially transfer content across
¹¹³⁵ editions. Maintaining consistency across language editions requires editors to reconcile differing source traditions,
¹¹³⁶ cultural emphases, and update rhythms. This process is inherently distributed and involves articulation work [53]:
¹¹³⁷ interpreting discrepancies, assessing their relevance to local communities, coordinating with other editors, and making
¹¹³⁸ decisions about whether and how to incorporate missing knowledge.
¹¹³⁹
¹¹⁴⁰

¹¹⁴¹ Participants noted that WIKIGAP reduces some of this articulation burden by automatically surfacing contextualized
¹¹⁴² discrepancies, highlighting their provenance, and linking directly to original sentences. These cues help editors evaluate
¹¹⁴³ the origins and significance of differences, understand when they reflect meaningful cultural divergence versus simple
¹¹⁴⁴ Manuscript submitted to ACM

1145 incompleteness, and coordinate potential updates. Rather than automating agreement or enforcing uniformity, WIKIGAP
1146 supports the human work of interpretation and negotiation that underlies Wikipedia’s collaborative ecosystem. By
1147 enabling multiple situated viewpoints to coexist while still facilitating coordination across them, the system exemplifies
1148 how boundary objects can scaffold the complex articulation work required for distributed knowledge production.
1149

1150 In this way, WIKIGAP renders the often invisible coordinative labor of multilingual maintenance more visible and
1151 manageable by making cross-lingual differences legible and actionable through provenance and in-context integration,
1152 reducing coordination friction and supporting editor engagement.
1153

1154 6.3 Positionality, LLMs, and the Future Role of WikiGap as Boundary Infrastructure

1155 Across both quantitative preferences and qualitative interviews, participants expressed strong appreciation for WIKI-
1156 GAP’s explicit display of provenance and its framing of multilingual knowledge as situated rather than objective. This
1157 recognition of positionality became particularly salient when contrasted—often spontaneously—against the increasingly
1158 common practice of using Large Language Models (LLMs) as ostensibly objective, consolidated knowledge sources.
1159

1160 A small subset of participants (n=2) expressed a desire for consolidated knowledge and therefore perceived gaps
1161 highlighted by WIKIGAP as evidence of English Wikipedia’s “incompleteness.” One participant noted that they would
1162 prefer to turn to LLMs, which present a single synthesized output. While LLM outputs are often perceived as authoritative
1163 or persuasive [10], they rarely foreground multiple cultural or linguistic perspectives and frequently exhibit English-
1164 centric biases [27, 41, 58]. They also obscure provenance [55], making it difficult for users to evaluate where information
1165 originates. LLMs thus risk presenting a “view from nowhere,” [22] even though their outputs reflect specific training
1166 data and design choices.
1167

1168 Although WIKIGAP and LLMs appear to sit at opposite ends of the spectrum—positionality versus objectivity—they
1169 need not be mutually exclusive. Our findings suggest that asset-based design principles foregrounding positionality
1170 can inform future LLM interfaces by encouraging users to see responses as situated rather than uniformly objective.
1171 Concretely, interfaces could surface multiple culturally grounded viewpoints instead of a single synthesized answer,
1172 along with lightweight provenance cues indicating which language communities or traditions each perspective reflects.
1173 Such designs would counter the tendency of current systems to amplify majority perspectives and obscure minority
1174 ones, helping users explore diverse interpretations more intentionally. They would also make latent cultural knowledge
1175 in LLMs more accessible to researchers and end users alike. In this sense, embedding WIKIGAP-like scaffolding into LLM
1176 interfaces represents a promising direction for foregrounding epistemic positionality within increasingly consolidated
1177 AI ecosystems. This extends our system’s asset-based design principles to cross-lingual and multilingual LLM interfaces,
1178 encouraging users to engage with knowledge as situated, grounded in provenance, and produced by distinct language
1179 communities.
1180

1181 6.4 Limitations and Future Work

1182 *Language Coverage.* Our study focuses on English and three target languages (French, Russian, and Chinese). While
1183 these languages span different linguistic families and large Wikipedia communities, they do not capture the full diversity
1184 of Wikipedia’s language ecosystem. Several participants expressed a desire for broader language coverage, particularly
1185 for pages in their language of origin. For example, one participant noted, “*I would love to see how the Tagalog page could*
1186 *complement it. In fact, diverse cultural perspectives on food are more linked to the cultural ecosystem where the food is*
1187 *grounded*” (P7). Expanding to more languages will enrich cross-cultural perspectives and align with users’ expectations
1188 in future work. The underlying INFOGAP pipeline is language-agnostic: in principle, it can operate on *any* language that
1189

¹¹⁹⁷ GPT-4 supports, though quality may vary, especially for low-resource languages. This is important for future research
¹¹⁹⁸ as smaller language editions remain largely understudied and present additional challenges in their own right [32, 45].
¹¹⁹⁹

¹²⁰⁰ *WIKIGAP content selection.* We selected a random subset of facts to present from other language editions (Sec. 3.3.2
¹²⁰¹ for further detail). This selection could be improved for greater localization through applying a geoprovence classifier
¹²⁰² [54], thereby emphasizing those facts that are sourced from a geographical region relevant to the language edition. This
¹²⁰³ is an important step for future work as our respondents expressed the desire for *localized* Volunteered Geographic
¹²⁰⁴ Information.
¹²⁰⁵

¹²⁰⁶ *Capacity for misinformation contagion.* Language editions vary considerably in their susceptibility to organized
¹²⁰⁷ disinformation campaigns [31]. Thus, there is a risk of spreading misinformation from other language editions into
¹²⁰⁸ English Wikipedia through WIKIGAP. WIKIGAP may thus inadvertently make smaller language editions more attractive
¹²⁰⁹ for actors carrying out disinformation campaigns, as audience size has previously been proposed as a motivating factor
¹²¹⁰ for bad-faith contributions [31]. Care should be thus taken in deploying WIKIGAP, with attention paid to the language
¹²¹¹ editions and topics that are selected for the extension.
¹²¹²

¹²¹³ *System Architecture.* WIKIGAP demonstrates a socio-technical intervention that aims to reshape how people access,
¹²¹⁴ relate to, and interpret multilingual knowledge, which is beyond a purely technical innovation. We have an over-
¹²¹⁵whelmingly positive result from the evaluation. It's important to acknowledge that these results come from a relatively
¹²¹⁶ controlled setting. While the findings are encouraging, future work will need to address the technical and interaction
¹²¹⁷ challenges of deploying a live, scalable version of the system in more variable real-world conditions.
¹²¹⁸

¹²¹⁹ *Future Work.* Building on our discussion of pluralistic LLM interfaces, future work could extend WIKIGAP's design
¹²²⁰ principles to other multilingual information-sharing platforms. Prior work shows that platforms such as Twitter
¹²²¹ and Flickr also exhibit substantial cross-lingual divergence in how events, entities, and cultural topics are discussed
¹²²² [12, 13, 19, 40]. Applying WIKIGAP could help users interpret how narratives differ across languages and communities
¹²²³ beyond Wikipedia. Extending this approach across platforms would also inform the design of more culturally aware
¹²²⁴ LLM interfaces, reinforcing our broader argument that pluralistic system design can mitigate the homogenizing effects
¹²²⁵ of consolidated AI infrastructures.
¹²²⁶

¹²²⁷ 7 Conclusion

¹²²⁸ WIKIGAP contends with longstanding structural inequalities in Wikipedia by enabling readers to access cross-lingual
¹²²⁹ facts directly within the English interface. Through a novel combination of fact-centric text comparison [INFOGAP; 52]
¹²³⁰ based on LLMs [42] and user-centered design, the system unobtrusively displays knowledge missing due to linguistic
¹²³¹ and editorial asymmetries. Our evaluation shows that WIKIGAP not only improves usability and fact-finding efficiency,
¹²³² but also fosters greater awareness of cultural variation and the limitations of assuming English Wikipedia as a default
¹²³³ knowledge base. By surfacing multilingual content in context, WIKIGAP connects readers with perspectives on topics
¹²³⁴ that are situated outside of the Anglosphere, connections that were previously occluded by the hegemonic dominance
¹²³⁵ of English Wikipedia [34, Chapter 3]. This work demonstrates how augmentative tools can shift reader practices and
¹²³⁶ challenge epistemic assumptions, paving the way for more inclusive, transparent, and culturally grounded knowledge
¹²³⁷ infrastructures. This paper examines how interface design can challenge the English-as-superset assumption by making
¹²³⁸ multilingual knowledge differences visible to readers during everyday Wikipedia use.
¹²³⁹

1249 8 Acknowledgments

1250 ChatGPT was utilized to generate code snippets and to assist in refining the grammar and writing style of this work. All
 1251 AI-generated content was reviewed and edited by the authors to ensure accuracy and adherence to academic standards.
 1252

1253 1254 References

- 1255 [1] Eytan Adar, Michael Skinner, and Daniel S Weld. 2009. Information arbitrage across multi-lingual Wikipedia. In *Proceedings of the Second ACM*
1256 International Conference on Web Search and Data Mining. 94–103.
- 1257 [2] Leah Hope Ajmani, Jasmine C. Foriest, Jordan Taylor, Kyle Pittman, Sarah Gilbert, and Michael Ann DeVito. 2024. Whose Knowledge is Valued?
1258 Epistemic Injustice in CSCW Applications. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW2, Article 523 (Nov. 2024), 28 pages. doi:10.1145/3687062
- 1259 [3] Ofer Arazy, Hila Lifshitz-Assaf, Oded Nov, Johannes Daxenberger, Martina Balestra, and Coye Cheshire. 2017. On the "How" and "Why" of Emergent
1260 Role Behaviors in Wikipedia. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (Portland,
1261 Oregon, USA) (CSCW '17). Association for Computing Machinery, New York, NY, USA, 2039–2051. doi:10.1145/2998181.2998317
- 1262 [4] Bunty Avieson and Suruchi Mazumdar. 2025. Wiki warriors: language editors counter knowledge hierarchies during the pandemic. *Asian Journal of
1263 Communication (2025), 1–19.*
- 1264 [5] Aaron Bangor, Philip Kortum, and James Miller. 2009. Determining what individual SUS scores mean: adding an adjective rating scale. *J. Usability
1265 Studies 4, 3 (May 2009), 114–123.*
- 1266 [6] Patti Bao, Brent Hecht, Samuel Carton, Mahmood Quaderi, Michael Horn, and Darren Gergle. 2012. Omnipedia: bridging the wikipedia language
1267 gap. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1075–1084.
- 1268 [7] Shaowen Bardzell. 2010. Feminist HCI: taking stock and outlining an agenda for design. In *Proceedings of the SIGCHI Conference on Human Factors in
1269 Computing Systems (Atlanta, Georgia, USA) (CHI '10)*. Association for Computing Machinery, New York, NY, USA, 1301–1310. doi:10.1145/1753326.
 1753521
- 1270 [8] Ruha Benjamin. 2023. Race after technology. In *Social Theory Re-Wired*. Routledge, 405–415.
- 1271 [9] Ewa S Callahan and Susan C Herring. 2011. Cultural bias in Wikipedia content on famous persons. *Journal of the American society for information
1272 science and technology 62, 10 (2011), 1899–1915.*
- 1273 [10] Shruthi Chockkalingam, Seyed Hossein Alavi, Raymond T. Ng, and Vered Shwartz. 2025. Should I go vegan: Evaluating the Persuasiveness of LLMs
1274 in Persona-Grounded Dialogues. In *Proceedings of the Third Workshop on Social Influence in Conversations (SICON 2025)*, James Hale, Brian Deuksin
 Kwon, and Ritam Dutt (Eds.). Association for Computational Linguistics, Vienna, Austria, 65–72. doi:10.18653/v1/2025.sicon-1.4
- 1275 [11] Noah Cohen. 2008. Open-Source Troubles in Wiki World. *New York Times*. <https://www.nytimes.com/2008/03/17/technology/17wikipedia.html>
- 1276 [12] Wei Dong and Wai-Tat Fu. 2010. Cultural difference in image tagging. In *Proceedings of the SIGCHI Conference on Human Factors in Computing
1277 Systems (Atlanta, Georgia, USA) (CHI '10)*. Association for Computing Machinery, New York, NY, USA, 981–984. doi:10.1145/1753326.1753472
- 1278 [13] Amir Dotan and Panayiotis Zaphiris. 2010. A cross-cultural analysis of Flickr users from Peru, Israel, Iran, Taiwan and the UK. *Int. J. Web Based
1279 Communities 6, 3 (June 2010), 284–302. doi:10.1504/IJWBC.2010.033753*
- 1280 [14] Kevin Duh, Ching-Man Au Yeung, Tomoharu Iwata, and Masaaki Nagata. 2013. Managing information disparity in multilingual document collections.
1281 ACM Transactions on Speech and Language Processing (TSLP) 10, 1 (2013), 1–28.
- 1282 [15] Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2020. Language-agnostic BERT Sentence Embedding. In *Proceedings
*1283 of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)**. Association for Computational Linguistics, 879–894.
<https://arxiv.org/abs/2007.01852>
- 1284 [16] Michel Foucault. 2013. *Archaeology of knowledge*. routledge.
- 1285 [17] Omer Goldman, Uri Shaham, Dan Malkin, Sivan Eiger, Avinatan Hassidim, Yossi Matias, Joshua Maynez, Adi Mayrav Gilady, Jason Riesa, Shruti
 Rijhwani, Laura Rimell, Idan Szpektor, Reut Tsarfaty, and Matan Eyal. 2025. ECLeKTic: a Novel Challenge Set for Evaluation of Cross-Lingual
1286 Knowledge Transfer. arXiv:2502.21228 [cs.CL] <https://arxiv.org/abs/2502.21228>
- 1287 [18] David Columbia. 2009. *The cultural logic of computation*. Harvard University Press.
- 1288 [19] Mark Graham, Scott A. Hale, and Devin Gaffney. 2014. Where in the World Are You? Geolocation and Language Identification in Twitter. *The
1289 Professional Geographer 66, 4 (2014), 568–578. doi:10.1080/00330124.2014.907699 arXiv:<https://doi.org/10.1080/00330124.2014.907699>*
- 1290 [20] Jonathan Grudin. 1988. Why CSCW applications fail: problems in the design and evaluation of organizational interfaces. In *Proceedings of the 1988
1291 ACM Conference on Computer-Supported Cooperative Work (Portland, Oregon, USA) (CSCW '88)*. Association for Computing Machinery, New York,
1292 NY, USA, 85–93. doi:10.1145/62266.62273
- 1293 [21] Scott A Hale. 2014. Multilinguals and Wikipedia editing. In *Proceedings of the 2014 ACM conference on Web science*. 99–108.
- 1294 [22] Donna Haraway. 2013. Situated knowledges: The science question in feminism and the privilege of partial perspective 1. In *Women, science, and
*1295 technology**. Routledge, 455–472.
- 1296 [23] Jean Hardy and Jacob Thebaud-Spieker. 2024. A Turn to Assets in Community-Based Computing Research: Tradeoffs, Deficits, and Neoliberalism in
1297 Technological Development. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 14 (April 2024), 20 pages. doi:10.1145/3637291
- 1298 [24] Brent Hecht and Darren Gergle. 2009. Measuring self-focus bias in community-maintained knowledge repositories. In *Proceedings of the fourth
*1299 international conference on communities and technologies**. 11–20.

[25] Brent Hecht and Darren Gergle. 2010. The tower of Babel meets web 2.0: user-generated content and its applications in a multilingual context. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 291–300.

[26] Brent Jaron Hecht. 2013. *The mining and application of diverse cultural perspectives in user-generated content*. Ph. D. Dissertation. Northwestern University.

[27] Daniel Hershcovitch, 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)*, Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (Eds.). Association for Computational Linguistics, Dublin, Ireland, 6997–7013. doi:10.18653/v1/2022.acl-long.482

[28] Molly G Hickman, Viral Pasad, Harsh Kamalesh Sanghavi, Jacob Thebault-Spieker, and Sang Won Lee. 2021. Understanding wikipedia practices through hindi, urdu, and english takes on an evolving regional conflict. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–31.

[29] Cheryl I Johnson and Richard E Mayer. 2012. An eye movement analysis of the spatial contiguity effect in multimedia learning. *Journal of Experimental Psychology: Applied* 18, 2 (2012), 178.

[30] Isaac Johnson and Emily Lescak. 2022. Considerations for multilingual wikipedia research. *arXiv preprint arXiv:2204.02483* (2022).

[31] Zarine Kharazian, Kate Starbird, and Benjamin Mako Hill. 2024. Governance capture in a self-governing community: A qualitative comparison of the croatian, serbian, bosnian, and serbo-croatian wikipedias. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–26.

[32] Sejal Khatri, Aaron Shaw, Sayamindu Dasgupta, and Benjamin Mako Hill. 2022. The social embeddedness of peer production: A comparative qualitative analysis of three Indian language Wikipedia editions. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–18.

[33] John P. Kretzman and John McKnight. 1993. Building Communities from the Inside Out: A Path Toward Finding and Mobilizing a Community's Assets. <https://api.semanticscholar.org/CorpusID:109453662>

[34] Sangeet Kumar. 2021. *The digital frontier: Infrastructures of control on the global web*. Indiana University Press.

[35] Paul Laufer, Claudia Wagner, Fabian Flöck, and Markus Strohmaier. 2015. Mining cross-cultural relations from Wikipedia: a study of 31 European food cultures. In *Proceedings of the ACM web science conference*. 1–10.

[36] Yiwei Luo, Kristina Gligoric, and Dan Jurafsky. 2023. Othering and Low Status Framing of Immigrant Cuisines in US Restaurant Reviews and Large Language Models. In *International Conference on Web and Social Media*. <https://api.semanticscholar.org/CorpusID:259936860>

[37] Emanuel Maiberg. 2025. Wikipedia Says AI Is Causing a Dangerous Decline in Human Visitors. *404 Media* (2025). <https://www.404media.co/wikipedia-says-ai-is-causing-a-dangerous-decline-in-human-visitors/> Published online 16 October 2025; accessed 2025-12-03.

[38] Catherine C Marshall. 1997. Annotation: from paper books to the digital library. In *Proceedings of the second ACM international conference on Digital libraries*. 131–140.

[39] Paolo Massa and Federico Scrinzi. 2012. Manypedia: Comparing language points of view of Wikipedia communities. In *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration*. 1–9.

[40] Julia Mendelsohn, Sayan Ghosh, David Jurgens, and Ceren Budak. 2023. Bridging nations: quantifying the role of multilinguals in communication on social media. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 17. 626–637.

[41] Rada Mihalcea, Oana Ignat, Longju Bai, Angana Borah, Luis Chiruzzo, Zhiqing Jin, Claude Kwizera, Joan Nwatu, Soujanya Poria, and Thamar Solorio. 2025. Why AI Is WEIRD and Shouldn't Be This Way: Towards AI for Everyone, with Everyone, by Everyone. *Proceedings of the AAAI Conference on Artificial Intelligence* 39, 27 (Apr. 2025), 28657–28670. doi:10.1609/aaai.v39i27.35092

[42] Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. Factscore: Fine-grained atomic evaluation of factual precision in long form text generation. *arXiv preprint arXiv:2305.14251* (2023).

[43] Marc Miquel-Ribé and David Laniado. 2018. Wikipedia culture gap: quantifying content imbalances across 40 language editions. *Frontiers in physics* 6 (2018), 54.

[44] Marc Miquel-Ribé and David Laniado. 2020. The Wikipedia Diversity Observatory: A Project to Identify and Bridge Content Gaps in Wikipedia. In *Proceedings of the 16th International Symposium on Open Collaboration* (Virtual conference, Spain) (*OpenSym '20*). Association for Computing Machinery, New York, NY, USA, Article 14, 4 pages. doi:10.1145/3412569.3412866

[45] Hellina Hailu Nigatu, John Canny, and Sarah E Chasins. 2024. Low-Resourced Languages and Online Knowledge Repositories: A Need-Finding Study.. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–21.

[46] Oded Nov. 2007. What motivates Wikipedians? *Commun. ACM* 50, 11 (Nov. 2007), 60–64. doi:10.1145/1297797.1297798

[47] OpenAI. 2024. GPT-4o mini: Advancing Cost-Efficient Intelligence. <https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence>. Large language model, July 18 version.

[48] OpenAI. 2024. GPT-4o System Card. *arXiv preprint arXiv:2410.21276* (2024). <https://arxiv.org/abs/2410.21276>

[49] Peter Pirolli and Stuart Card. 1999. Information foraging. *Psychological review* 106, 4 (1999), 643.

[50] Dwaipayan Roy, Sumit Bhatia, and Prateek Jain. 2022. Information asymmetry in Wikipedia across different languages: A statistical analysis. *Journal of the Association for Information Science and Technology* 73, 3 (2022), 347–361.

[51] Edward W Said. 1977. Orientalism. *The Georgia Review* 31, 1 (1977), 162–206.

1353 [52] Farhan Samir, Chan Young Park, Anjalie Field, Vered Shwartz, and Yulia Tsvetkov. 2024. Locating Information Gaps and Narrative Inconsistencies
 1354 Across Languages: A Case Study of LGBT People Portrayals on Wikipedia. In *Proceedings of the 2024 Conference on Empirical Methods in Natural*
 1355 *Language Processing*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA,
 1356 6747–6762. doi:10.18653/v1/2024.emnlp-main.384

1357 [53] Kjeld Schmidt and Liam Bannon. 1992. Taking CSCW seriously. 1, 1 (1992), 7–40. doi:10.1007/BF00752449

1358 [54] Shilad W Sen, Heather Ford, David R Musicant, Mark Graham, OS Keyes, and Brent Hecht. 2015. Barriers to the localness of volunteered geographic
 1359 information. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 197–206.

1360 [55] Chirag Shah and Emily M Bender. 2022. Situating search. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval*.
 221–232.

1361 [56] Susan Leigh Star. 2010. This is Not a Boundary Object: Reflections on the Origin of a Concept. *Science, Technology, & Human Values* 35, 5 (2010),
 1362 601–617. doi:10.1177/0162243910377624 arXiv:<https://doi.org/10.1177/0162243910377624>

1363 [57] Susan Leigh Star and James R. Griesemer. 1989. Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals
 1364 in Berkeley's Museum of Vertebrate Zoology, 1907–39. *Social Studies of Science* 19, 3 (1989), 387–420. doi:10.1177/030631289019003001
 1365 arXiv:<https://doi.org/10.1177/030631289019003001>

1366 [58] Yan Tao, Olga Viberg, Ryan S Baker, and René F Kizilcec. 2024. Cultural bias and cultural alignment of large language models. *PNAS nexus* 3, 9
 1367 (2024), pgae346.

1368 [59] Nicholas Vincent and Brent Hecht. 2021. A deeper investigation of the importance of Wikipedia links to search engine results. *Proceedings of the*
 1369 *ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–15.

1370 [60] Wikimedia Foundation. 2011. *Editor Survey Report – April 2011*. Technical Report. Wikimedia Foundation. https://upload.wikimedia.org/wikipedia/commons/7/76/Editor_Survey_Report_-_April_2011.pdf Accessed.

1371 [61] Wikimedia Foundation. 2024. Announcing English Wikipedia's most popular articles of 2024. <https://wikimediafoundation.org/news/2024/12/03/announcing-english-wikipedias-most-popular-articles-of-2024/> Accessed: 2025-09-09.

1372 [62] Genta Indra Winata, Frederikus Hudi, Patrick Amadeus Irawan, David Anugraha, Rifki Afina Putri, Wang Yutong, Adam Nohejl, Ubaidillah Arij
 1373 Prathama, Nedjma Ousidhoum, Afifa Amriani, Anar Rzayev, Anirban Das, Ashmari Pramodya, Aulia Adila, Bryan Wilie, Candy Olivia Mawalim,
 1374 Cheng Ching Lam, Daud Abolade, Emmanuele Chersoni, Enrico Santus, Fariz Ikhwantri, Garry Kuwanto, Hanyang Zhao, Haryo Akbarianto Wibowo,
 1375 Holy Lovenia, Jan Christian Blaise Cruz, Jan Wira Gotama Putra, Junho Myung, Lucky Susanto, Maria Angelica Riera Machin, Marina Zhukova,
 1376 Michael Anugraha, Muhammad Farid Adilazuarda, Natasha Christabelle Santosa, Peerat Limkonchotiwat, Raj Dabre, Rio Alexander Audino, Samuel
 1377 Cahyawijaya, Shi-Xiong Zhang, Stephanie Yulia Salim, Yi Zhou, Yinxuan Gui, David Ifeoluwa Adelani, En-Shiun Annie Lee, Shogo Okada, Ayu
 1378 Purwarianti, Alham Filkri Aji, Taro Watanabe, Derry Tanti Wijaya, Alice Oh, and Chong-Wah Ngo. 2025. WorldCuisines: A Massive-Scale Benchmark
 1379 for Multilingual and Multicultural Visual Question Answering on Global Cuisines. In *Proceedings of the 2025 Conference of the Nations of the Americas*
 1380 *Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Luis Chiruzzo, Alan Ritter, and
 1381 Lu Wang (Eds.). Association for Computational Linguistics, Albuquerque, New Mexico, 3242–3264. <https://aclanthology.org/2025.naacl-long.167/>

1382 [63] Marisol Wong-Villares, Aakash Gautam, Wendy Roldan, Lucy Pei, Jessa Dickinson, Azra Ismail, Betsy DiSalvo, Neha Kumar, Tammy Clegg, Sheena
 1383 Erete, Emily Roden, Nithya Sambasivan, and Jason Yip. 2020. From Needs to Strengths: Operationalizing an Assets-Based Design of Technology.
 1384 In *Companion Publication of the 2020 Conference on Computer Supported Cooperative Work and Social Computing (Virtual Event, USA) (CSCW '20*
 1385 *Companion)*. Association for Computing Machinery, New York, NY, USA, 527–535. doi:10.1145/3406865.3418594

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397 **A Sample Quiz Questions**

1398 We present two example questions from each of the five topics used in our study. These examples illustrate the format
 1399 and cultural specificity of the quiz items. All questions were drawn from the 30 language-exclusive facts surfaced by
 1400 WIKI GAP, and were manually reviewed for clarity and accuracy. The correct answers are bolded in Table 4.

1401

1402

1403

1404

Table 4. Example Quiz Questions Across Topics. Correct answers are bolded.

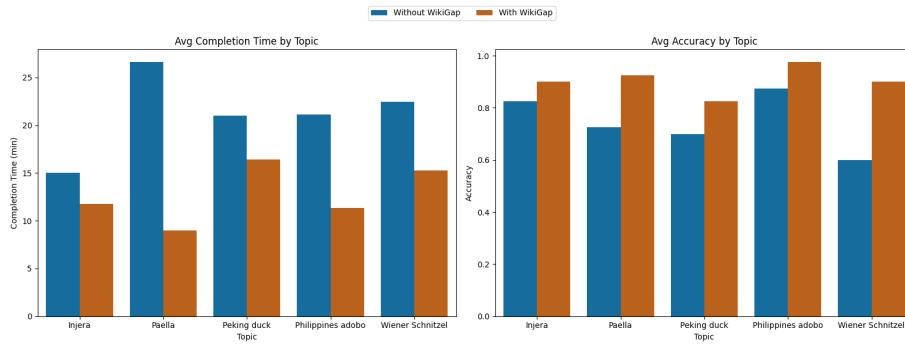
Topic	Question 1	Question 2
Peking Duck	<i>What unique preparation step is done to the duck before roasting it for Peking duck?</i> A. It is marinated in soy sauce B. It is inflated with air under the skin C. It is stuffed with rice D. It is boiled in water	<i>What is served after the meat in a traditional Peking duck meal?</i> A. A fruit platter B. Rice C. Chinese cabbage soup D. A cup of tea
Wiener Schnitzel	<i>During which historical period was Wiener schnitzel brought to Italy and then to Austria?</i> A. World War I B. French Revolution C. Napoleonic Wars D. Renaissance	<i>Which city of China has a type of Western cuisine similar to Vienna schnitzel?</i> A. Beijing B. Nanjing C. Guangdong D. Shanghai
Paella	<i>What festival is Spanish paella associated with?</i> A. La Tomatina B. San Fermín C. Falles D. Semana Santa	<i>How was paella traditionally eaten?</i> A. On plates B. In bowls C. Straight from the cooking pan D. Consumed with bread
Philippine Adobo	<i>Which ingredient is sometimes added to adobo to replace palm or coconut sugar and help tenderize the meat?</i> A. Honey B. Molasses C. Sprite D. Maple syrup	<i>Which of the following is a traditional method used in adobo to keep meat fresh in tropical climates?</i> A. Drying the meat B. Frying with vinegar C. Smoking the meat D. Freezing the meat
Injera	<i>How large is traditional injera typically?</i> A. 30 centimeters B. Half a meter C. About 1 meter D. 2 meters	<i>Injera closely resembles which Middle Eastern pancake variant?</i> A. Naan B. Lahoh C. Pita D. Lavash

B Interview and Questionnaire Protocol

Participants were asked the following questions during either the semi-structured interview or the open-ended questionnaire.

- Describe your experience using WIKIGAP. What stood out to you the most, and why?
- When you enabled WIKIGAP to see multilingual facts, how did you feel about the way information was presented? What aspects of the presentation worked well or could be improved?
- Can you describe a specific fact you discovered through WIKIGAP that surprised you or changed your understanding of the topic? How did this discovery influence your perception or understanding of diverse cultural viewpoints?
- When using WIKIGAP, how did you decide which information to trust? Please explain your reasoning.
- After using WIKIGAP, do you see Wikipedia as more reliable or complete, especially for finding information from other languages? Please explain your answer with examples if possible.
- In what ways, if any, did WIKIGAP change your perception or understanding of diverse cultural viewpoints related to the topics you explored?
- Did using WIKIGAP make you more inclined to explore content outside your primary language in the future? Why or why not?

1457 C Average Completion Time and Accuracy across Different Topics



1472 Fig. 8. Side-by-side bar charts comparing **average completion time** and **accuracy** across different topics, with and without the
 1473 **WikiGap** extension. **Blue** represents the control condition (no **WikiGap**), and **orange** represent the treatment condition (with **WikiGap**).