PlaceWeave: Understanding Place Through Social Video Narratives and Graph-Enhanced Local Knowledge

Zihan Gao
Information School
University of Wisconsin-Madison
Madison, Wisconsin, USA
zihan.gao@wisc.edu

Jacob Thebault-Spieker
Information School
University of Wisconsin - Madison
Madison, Wisconsin, USA
jacob.thebaultspieker@wisc.edu

Abstract

People visiting or moving to a new city often struggle to understand local vibes and everyday routines. Short-form videos on TikTok capture these local stories, but people still have to jump between chatbots, maps, and apps to turn them into concrete plans. We introduce PlaceWeave, a human-centered trip-planning system that foregrounds a place's "localness". PlaceWeave builds a place knowledge graph from TikTok videos and uses it to ground all AI features: the conversational assistant, localness attributes on the map, and the route planner all draw on graph evidence. The interface combines an interactive map, an evidence-backed Insights Panel, and tools for organizing discoveries and composing itineraries in a single linked workspace. We validate the attributes and run a within-subjects study with 18 participants, comparing PlaceWeave to a baseline using separate chat, map, video, and canvas tools. PlaceWeave helps people create more local-feeling plans, better understand neighborhood character and trade-offs, and avoid fragmented workflows. We show how localness-aware, graph-grounded AI can support more community-sensitive placemaking technologies.

CCS Concepts

• **Human-centered computing → HCI design and evaluation methods**; **User centered design**; • **Computing methodologies → Knowledge representation and reasoning**.

Keywords

Localness, Sensemaking, Knowledge Graphs, Retrieval-Augmented Generation, Provenance, Travel Planning

ACM Reference Format:
Zihan Gao and Jacob Thebault-Spieker. 2026. PlaceWeave: Understanding Place Through Social Video Narratives and Graph-Enhanced Local Knowledge. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 23 pages. https://doi.org/10.1145/3772318.3791894

1 Introduction

When people move or travel, they struggle to understand the everyday texture of a place [28]: Where do locals actually eat? What's the vibe of a neighborhood—family-oriented, artsy, student-heavy? To answer these questions, people are increasingly turning to user-generated content (UGC), especially short videos, for richer insights into local environments [17, 28, 71, 76]. These videos offer immersive, narrative-rich portrayals of place through walking tours, daily routines, or event vlogs that convey experiential cues about how people see, hear, and interact with their surroundings.

This pursuit of authentic, lived-in experience is fundamentally a search for a place's "localness"—a condition that is not merely a matter of physical proximity but reflects how people become situated in a place through cultural practices, social ties, and embodied routines [29, 38, 41, 49, 57]. Localness is related to, but distinct from, concepts such as sense of place, placemaking, or authenticity. Sense of place emphasizes internal emotional bonds and meanings attached to a place [49, 85], while placemaking foregrounds collective efforts to reshape space and build community through design and everyday practice [11]. Authenticity, especially in tourism studies, is often framed as a normative judgment about whether an experience feels "touristy" or "local" [65, 87]. By contrast, we follow Gao et al. [29] in treating localness as a translational construct that focuses on observable manifestations of these bonds: how local knowledge, lived experience, and community relationships are expressed in everyday practices and online social traces.

This turn to short-video UGC is often a workaround: mainstream location tools still struggle to convey a place's "localness," forcing users to navigate a fragmented information landscape. Utility-focused services like Google Maps primarily foreground geometry and efficiency, reducing places to pins, routes, and opening hours [58, 98]. Review platforms like Yelp or TripAdvisor distill complex social atmospheres into star ratings and keyword tags, often privileging tourist-heavy venues and majority preferences [48, 52]. At the same time, the vibrant signals of local life are scattered across unstructured social media, buried in countless ephemeral posts [4, 19, 35]. Newcomers must manually piece together a sense of a neighborhood by bouncing between videos, maps, and notes, making it hard to tell whether their plans reflect local practices or just the most visible options [28].

This disconnect reflects a long-standing challenge that Ackerman [1] terms the *sociotechnical gap*: the qualities that make places meaningful—practice, memory, and social relations—are difficult to formalize computationally [38, 49, 57, 98]. Consequently, mainstream location services and ranking algorithms prioritize what is readily measurable (e.g., coordinates, counts, ratings) over the relational and temporal dimensions through which localness is performed and recognized [20, 29, 37, 48, 52, 58]. In practice, newcomers must shoulder the work of integrating heterogeneous traces

ACM ISBN 979-8-4007-2278-3/26/04
https://doi.org/10.1145/3772318.3791894

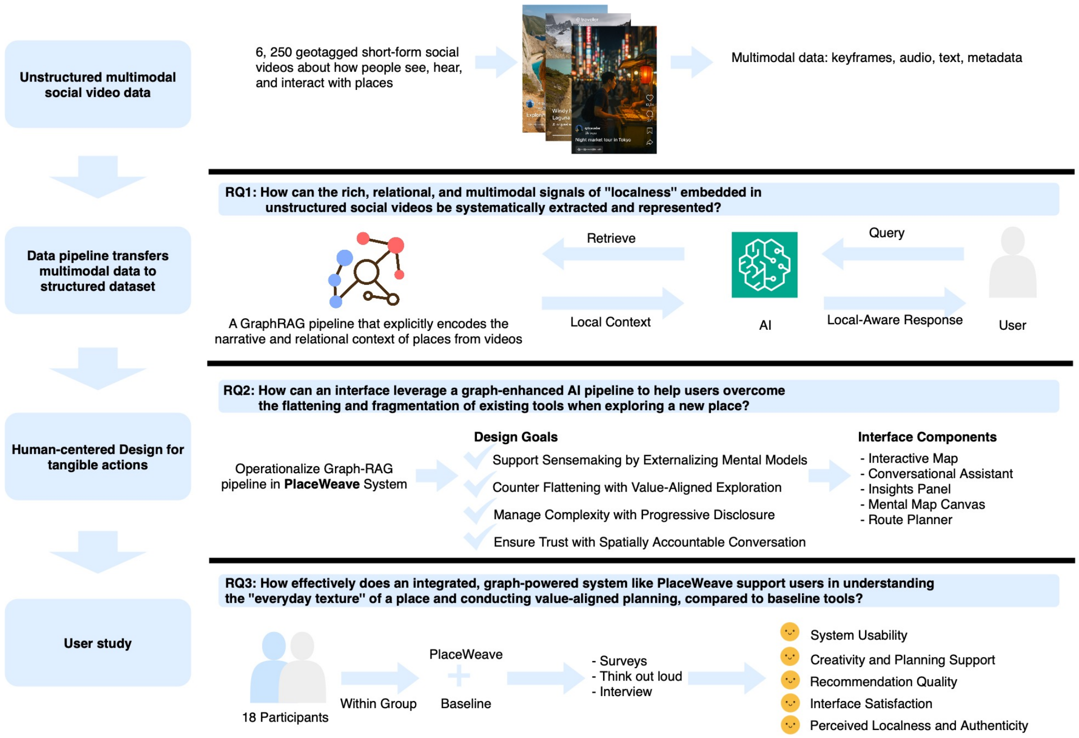

Figure 1: An overview of our research workflow, structured around our three research questions. (RQ1) We start with unstructured social videos and develop a data pipeline to build a Graph-RAG model that represents the relational context of "localness." (RQ2) This pipeline is then operationalized in the PlaceWeave system, a human-centered interface designed with specific goals and components to counter the flattening and fragmentation of existing tools. (RQ3) Finally, we evaluate PlaceWeave's effectiveness through a within-subjects user study, assessing its ability to help users understand a place's "everyday texture" compared to baseline tools.

into decisions [78, 79], often defaulting to options that are visible and safe rather than grounded in local practices [4, 28].

To bridge this gap, we transform narrative social traces into a structured, inspectable representation that supports localness-aware planning. *Our goal in this paper is to build tools that mediate newcomers' relationships to place through a localness lens.* Following Ackerman [1]'s call for "first-order approximations" to complex social phenomena, we treat localness as something we can approximate from digital traces while keeping its limits visible. Concretely, we use TikTok as a primary input because its short-form, narrative-first videos can surface vernacular place narratives, including those from underrepresented communities [4, 13, 17, 24, 69, 110]. We then construct a place knowledge graph that encodes where, when, and how people describe their use of places (e.g., ambience, activities, companions, time-of-day) [42, 62]. Building on graph-enhanced retrieval-augmented generation (Graph-RAG) [46, 77], we pair this graph with LLMs to surface localness attributes (such as typical visitors, rhythms, and community activities) and retrieve supporting evidence. Unlike text-only RAG, which typically retrieves top-ranked text chunks and inserts them into a prompt, all AI features in our system, including the conversational assistant, localness map overlays, and the route planner, query this same shared, structured graph, keeping outputs consistent and traceable to specific entities, relations, and source excerpts.

We operationalize this pipeline in PlaceWeave, a human-centered trip-planning system designed to make localness tangible and actionable. To directly counter the flattening of place into pins and stars, the system's **Insights Panel** surfaces traceable localness attributes with confidence cues and links to source videos, fostering transparency and contestability [14, 15, 26]. To reduce fragmented workflows [28, 79], PlaceWeave brings together an **Interactive Map**, a **Conversational Assistant** for experiential queries, a **Mental Map Canvas** for staging and organizing discoveries, and a **Route Planner** for composing value-aligned itineraries in a single linked workspace. Guided by principles of progressive disclosure [88] and spatial accountability between views [19, 67], all components are tightly coupled: assistant suggestions, insight selections, and itinerary edits are immediately reflected on the map. In the rest of the paper, we examine how this localness-aware, graph-grounded AI pipeline and interface together shape newcomers' planning strategies, sense of place, and trust (Figure 1).

In reporting on PlaceWeave, this paper makes five primary contributions to CHI:

(1) We first demonstrate a framework-informed method for constructing place knowledge graphs from short-form social

video, encoding "localness" as relational, evidence-backed attributes grounded in how people describe and use places.
(2) We then demonstrate a technique for a localness-aware pipeline that adapts existing Graph-RAG ideas to ground conversational assistance, map overlays, and route planning in the same underlying evidence.
(3) Building on these two computational components, we make a systems contribution [102] with PlaceWeave, a human-centered system that brings these capabilities together in an interface that unifies experiential search, spatial exploration, and itinerary composition while keeping provenance and uncertainty visible.
(4) We show that PlaceWeave supports more local-feeling plans, better understanding of neighborhood character and trade-offs, and less fragmented workflows through technical validation and a within-subjects study with 18 participants, comparing PlaceWeave to a realistic multi-tool baseline (ChatGPT + TikTok + Google Maps + Draw.io). We draw out design implications for localness-aware, AI-mediated place technologies.
(5) Finally, while PlaceWeave makes a compelling contribution toward "platial" technologies, aggregating vernacular content into AI-generated plans raises questions about whose stories are surfaced, how attention is redistributed, and how such systems might amplify or mitigate tourism pressures on local communities. We interrogate and unpack these tensions and develop conceptual contributions about designing within such spaces.

2 Related Work

Our work here is contextualized against prior research that focuses on (1) conceptualizations of localness and place, (2) short-form video as a driver of travel inspiration, (3) travel planning and itinerary systems, (4) knowledge graphs and localness-aware reasoning, (5) graph-enhanced retrieval and LLM grounding, and (6) exploratory interfaces for place sensemaking. We build particularly on recent work that conceptualizes localness as a multi-dimensional, socially constructed status that must be demonstrated and recognized, and that organizes it into a practical framework for design and computation [29]. We use this lens to review how existing systems do (and do not) capture localness, and to motivate a system that turns social-video inspiration into grounded, constraint-aware plans through a localness-aware backend and an interface that supports exploratory, evidence-visible sensemaking.

2.1 Conceptualizing Localness

We use "localness" as a lens for understanding how people become situated in a place through everyday practices, knowledge, and relationships. Here, we situate that lens within prior work and explain how we operationalize it.

Research in human geography and environmental psychology has long emphasized that places are not just coordinates but socially produced settings shaped by experience, memory, and practice [38, 57, 98]. Sense of place research focuses on emotional bonds and identity connections to place [41, 49], while placemaking highlights collective efforts to reconfigure spaces and build community. HCI has drawn on these concepts to design technologies that support attachment, reflection, and participation in local environments [11, 28]. However, these constructs are often treated at a high level, making it difficult to translate them into concrete representational choices or annotation schemes.

We build on Gao et al. [29], who propose "localness" as a translational concept that connects these theoretical traditions to observable practice. Their study traces how people evaluate and decide whether someone is local in everyday conversation, showing that localness is a dynamic, socially negotiated status that must be demonstrated and recognized rather than inferred solely from proximity or residence. From these judgments, they derive the *Localness Conceptual Framework*, which organizes localness into three interconnected domains: *Physical* (embodied presence and routines in place), *Cognitive* (local knowledge and cultural understanding), and *Relational* (social connections and commitments). The framework further decomposes each domain into components and example cues, yielding a structured, hierarchical view of how localness is expressed and perceived.

For our purposes, this framework is valuable not only as a definition but as a design and annotation tool. It provides concrete categories that can guide: (1) what we look for when extracting localness signals from social video; (2) how we structure those signals in a place knowledge graph; and (3) how we evaluate whether our system is capturing the aspects of localness that matter. In PlaceWeave, we therefore consistently treat the localness domains as lenses that shape the prompts and outputs of our pipeline, and the organization of our technical validation and user study measures.

2.2 Short-Form Video and Travel Inspiration

Short-form video platforms such as TikTok have become a major channel for travel discovery among Gen Z and Millennials, shaping destination awareness and choice [3, 13, 60, 69, 100, 103, 109, 110]. Their narrative-first format surfaces hyper-local experiences—late-night dumpling spots, quiet reading nooks, festival scenes—through vlogs and walkthroughs that convey ambience and social practice beyond ratings and addresses [89, 110]. Prior work also shows that entertainment, informativeness, emotion-oriented storytelling, and parasocial credibility can increase travel intention, while trends and remixability accelerate the spread of situated narratives [3, 60, 97, 100, 103, 108].

These properties make social video well-suited for operationalizing localness cues because posts encode situated routines—who does what, where, and when—via narratives alongside metadata such as geotags and timestamps. Foundational work on UGC (e.g., check-ins and curated collections such as *Livehoods* and *Curated City*) demonstrates that neighborhood structure and personal guides can be recovered from social traces, supporting "digital geographies" of place [4, 19, 20]. Because short-form video is multimodal, it can additionally surface relational signals—audio-visual ambience, co-mentioned activities and companions—that naturally bind entities such as *place–ambience–role–time*, supporting localness-aware representation and retrieval [4, 19, 20, 89]. At the same time, localness signals in geotagged content are uneven across platforms, demographics, and neighborhoods: posts often overrepresent tourist districts and underrepresent peripheral communities, skewing impressions of what locals actually do [48, 52]. Viral amplification can also make destinations "famous overnight," concentrating attention [101], while inspirational clips may emphasize affect and aesthetics over practical constraints (e.g., hours, rhythms, sequencing), leaving a gap between inspiration and feasible plans [103].

Our work treats short-form video not only as inspiration but as a primary input modality: we extract place-related structure from TikTok into a place knowledge graph that preserves relational and experiential context, enabling retrieval and generation that maintain localness cues and connect serendipitous discovery with grounded, constraint-aware plans [4, 28].

2.3 Travel Planning Tools and Itinerary Systems

HCI and adjacent work have explored travel planning and itinerary generation through interactive collaboration, data-driven sequencing, and constraint-aware optimization. Some systems support shared planning and reuse of peer knowledge [2], while others enable iterative, natural-language planning and replanning through crowd assistance [83]—interaction patterns that motivate our own emphasis on exploratory iteration and user control. A second line of work couples recommendation with optimization to personalize routes and schedules [105], including approaches that model experience flow by mining transitions to sequence POIs coherently [75]. Relatedly, social traces such as check-ins have been used to construct itineraries and recommend collective traveling paths at city scale [21, 43, 44]. More recently, LLM-based planners generate personalized multi-day plans [93] and tool-using agents target scenario-aware itineraries [16], while benchmarks stress-test constraint handling in realistic settings [104].

Despite this progress, gaps remain in turning open-ended, social-video inspiration into grounded plans. Many systems assume travelers begin with a destination or fixed intent rather than serendipitous media [2, 75, 105]. Inputs typically center on POI catalogs, historical trajectories, or textual reviews rather than short-form social video and its experiential cues [21, 43, 75, 105]. Even with LLMs, plan quality hinges on spatiotemporal rationality (ordering, time windows, transit), which remains challenging; recent evaluations underscore the need for retrieval that captures routes and time windows, not just textual POI blurbs [73, 93, 104].

Our system differs in two key ways. First, we shift the starting context from structured POI lists to vernacular social video. Second, we expand the goal from feasibility-focused optimization to *localness-aware* planning. PlaceWeave uses TikTok clips as primary source material, extracts structured signals of local practice, and couples graph-aware retrieval with a human-centered interface that foregrounds evidence, explanation, and user control, to support experiential, not just logistically feasible, itineraries.

2.4 Knowledge Graphs and Localness-Aware Reasoning

Knowledge graphs organize heterogeneous urban and tourism data into entities and relations, enabling context-aware retrieval and reasoning across places, activities, and time [36, 42, 53, 62]. However, most existing urban/tourism graphs remain *entity-centric* and coarse: they integrate POIs, administrative units, and events, but rarely represent experiential or social-practice attributes (e.g., ambience, rhythms, user roles), and they are slow to reflect trend-driven nuance from social media video [36, 42, 53, 62]. This matters because place is often understood as performative and co-constructed through interactions among visitors, locals, and media [11, 72, 87]; without explicit modeling of such facets, retrieval tends to follow visibility signals in social traces, skewing toward tourist-heavy areas and away from peripheral communities [48, 52]. As a result, current KG-backed systems are ill-equipped to answer queries like "study-vibe café locals actually use" or to align recommendations with community practices.

We extend this line of work by explicitly grounding our place knowledge graph in the Localness Conceptual Framework [29]. From short-form social videos, we extract localness-relevant attributes and encode them as nodes and relations. These attributes are linked to places via typed relations (e.g., rhythms, local recommendations, community events). This alignment allows structure-aware queries such as "weekday afternoon study spots near a community hub" and provides a direct bridge between theoretical domains of localness and the evidence that powers our system's outputs. For HCI, this graph is not merely a backend representation but a *sociotechnical* classification and encoding schema [7] that shapes what the interface can show and explain. By encoding different facets of localness, we enable interfaces that surface who a place is for, when it is used, and how it fits into community rhythms—supporting newcomers' sensemaking in ways that traditional POI-centric tools cannot.

2.5 Graph-Enhanced Retrieval and LLM Grounding

Retrieval-augmented generation (RAG) is a common strategy for grounding large language model outputs in external data, but text-only RAG typically retrieves chunks based on surface semantic similarity. For place reasoning, this is often insufficient: users care about how places, ambiences, roles, and times fit together into trajectories and everyday routines, not just whether a given POI is mentioned in a document [25, 45, 46, 77, 106]. This motivates retrieval that can exploit the relational structure encoded in a place knowledge graph, instead of returning isolated matches.

Graph-RAG addresses this limitation by indexing a corpus as a graph of entities and relations, organizing it into communities, and performing graph-guided retrieval prior to generation [25, 45, 46, 77, 106]. Such structure-aware retrieval can improve coverage on complex, multi-hop queries. In the travel-planning domain, benchmarks such as *TP-RAG* highlight that plan quality depends on capturing spatiotemporal rationality, such as ordering, time windows, and route coherence, rather than only matching text snippets about POIs [73]. Systems that integrate spatial structure into LLM-based planning loops (e.g., *ItiNera*) similarly report gains, but still hinge on high-quality grounded retrieval [93].

We follow this arc by adapting Graph-RAG to retrieve over our *localness-grounded* place knowledge graph derived from short-form video. Localness-aware relations guide both retrieval and prompting, allowing the conversational assistant, map overlays, and route planner to draw from the same evidence-linked representation of place. For HCI, our contribution is less about proposing a new retrieval algorithm than demonstrating how graph-structured grounding changes the experience of using AI: it supports experiential queries (e.g., "student-y but not touristy"), explains suggestions in terms of rhythms and co-occurrence, and keeps provenance visible—making Graph-RAG a *design choice* that enables transparency, controllability, and value-aligned recommendations in PlaceWeave.

2.6 Exploratory Interfaces for Planning and Place Sensemaking

Large, heterogeneous information spaces benefit from *progressive disclosure* such as overview, zoom/filter, and details-on-demand to let users move from gist to evidence without overload [88]. For urban exploration, social curation and neighborhood-level representations turn scattered narratives into coherent guides (e.g., *Curated City*) [20], while gamified mechanics (e.g., *MapUncover*) encourage open-ended exploration [86]. Bottom-up city structure inferred from social traces (e.g., *Livehoods*) motivates interfaces that help people reason with *neighborhood-scale patterns* rather than isolated POIs [19].

With LLMs, direct manipulation can make model suggestions visible, reversible, and traceable within the user interface. *DirectGPT* demonstrates translating manipulations into prompts and reflecting effects as manipulable state with undo/redo [67]. Complementing these ideas, research on *mental maps* highlights externalizing subjective place knowledge during transitions, emphasizing tools that preserve evolving spatial understandings and support belonging [28]. Where algorithmic uncertainty is inevitable, *seamful design* argues for surfacing evidence and limitations in situ rather than hiding them behind seamless abstractions [14, 15].

We synthesize these strands in an interface that keeps experiential intent visible (media-first itinerary cards, localness badges, rhythm-aware timelines) and makes conversational references *spatially accountable* (chat-to-map highlights, filters, routes), so users can verify, compare, and reflect on plans rather than optimize for efficiency alone [20, 67, 88].

2.7 Our Work Here

Based on these bodies of prior work, our research here focuses on three key research questions:

- **RQ1** (Representing "Localness"): How can the rich, relational, and multimodal signals of "localness" embedded in unstructured social videos be systematically extracted and represented in a way that instantiates the physical, cognitive, and relational domains of localness?
- **RQ2** (Designing for "Localness"): How can an interface leverage a graph-enhanced AI pipeline to help users overcome the flattening and fragmentation of existing tools when exploring a new place?
- **RQ3** (Evaluating Support for "Localness"): How effectively does an integrated, graph-enhanced system like PlaceWeave support users in understanding the "everyday texture" of a place and conducting value-aligned planning, compared to baseline tools?

3 The PlaceWeave System

In this section, we describe the technical foundation of PlaceWeave: how we transform raw TikTok posts into the localness attributes that appear in the interface. We first show how we build a place knowledge graph from social video (Sec. 3.3), then explain how a localness-aware Graph-RAG engine uses this graph to ground the conversational assistant, map overlays, and route planner, and finally report a technical validation of the inferred localness attributes. Together, these choices define what the system can and cannot say about places and set the stage for our user study (Sec. 4).

3.1 Design Goals

Inspired by Ackerman [1]'s call for "first-order approximations" and viewing localness as a domain instance of the broader sociotechnical gap [1], our design of PlaceWeave is guided by four core principles. The following *Design Goals (DGs)* directly address the challenges of fragmentation, flattening, and cognitive overload by creating a system for meaning-making, not merely for efficiency. Table 1 provides a detailed reference mapping each affordance to its primary design goal and the interface components that manifest it.

3.1.1 DG1: Support sensemaking by externalizing mental models.

To combat the fragmentation and high cognitive load of piecing together a plan from disparate sources, our first goal is to help users build and organize their own understanding of a place. Foundational work in urban cognition shows that people navigate using personal "mental maps" [63]. PlaceWeave supports this process by providing a *Mental Map Canvas*, a flexible workspace for externalizing thoughts and discoveries [28, 54]. This canvas acts as a dedicated space for sensemaking [79], where users can use *Rich-Media Place Cards* to preserve the inspirational "why" behind a choice and create *User-Generated Annotations* to form personal collections (e.g., *"quiet morning spots"*, *"artsy walks"*).

3.1.2 DG2: Counter flattening with value-aligned exploration for localness.

To counter the flattening of place into simplistic ratings and keywords, our second goal is to surface the authentic character of a place, enabling exploration that aligns with user values beyond efficiency [27]. Rather than relying on generic popularity, which often reflects tourist-centric views, we use a *Localness Guided Interpretation* to highlight places with authentic rhythms derived from social video [29]. This focus on values is carried into the planning stage with *Rhythm-Aware Temporal Organization*, which grounds suggestions in the actual pulse of the community, honoring the fact that people often trade efficiency for experiential qualities [80].

3.1.3 DG3: Manage complexity with progressive disclosure.

To make the rich, dense information of a city understandable without overwhelming the user, our third goal is to manage complexity through progressive disclosure. Following the visual information seeking tasks (overview, zoom, filter, details-on-demand, relate, history, and extract) [88], PlaceWeave uses *Dynamic Visual Encodings* on the map. These preattentive cues, such as icons and color provide an at-a-glance overview of localness and popularity [39]. For details, the system surfaces fine-grained, *Multimodal Experiential Attributes* (e.g., warm lighting and laptop activity) extracted from the source videos, allowing users to drill down into the experiential details that matter to them.

Table 1: PlaceWeave design framework: how each DG is realized through user-facing affordances and where they appear in the interface (Figure 2).

Affordance	What this does for users while planning	Where in the interface
DG1 *Support* ***sensemaking*** *by externalizing* ***mental models***		
Rich-media place cards	Keep the **"why" behind each place** visible: video thumbnail, tags, and notes stay attached so users remember why they saved it.	(A) Interactive Map, (C) Insights Panel, (D) Mental Map Canvas
Spatial and thematic grouping	Let users **drag places into their own clusters** (e.g., "quiet morning spots," "evening options") to build personalized mental maps.	(D) Mental Map Canvas
DG2 *Counter flattening with* ***value-aligned exploration*** *for* ***localness***		
Localness interpretation	Surface **localness cues** (hidden gems, regulars vs. tourists) so users can pick places that match their values, not just ratings.	(A) Interactive Map, (C) Insights Panel
Rhythm-aware planning	Suggest **when** to go: time windows that reflect community rhythms (busy vs. quiet times), with assumptions made explicit.	(B) Conversational Assistant, (E) Route Planner
DG3 *Manage* ***complexity*** *with* ***progressive disclosure***		
Dynamic visual encodings	Provide an **at-a-glance overview** of popularity and localness (marker size and color) so users can scan the city before zooming in.	(A) Interactive Map
Layered detail views	Let users move from **map** → **pop-up** → **Insights Panel**, revealing experiential details (e.g., "quiet vibe," "warm lighting") only when needed.	(A) Interactive Map, (C) Insights Panel
DG4 *Ensure* ***trust*** *with* ***spatially accountable*** *conversation*		
Map-grounded AI suggestions	Turn open-ended requests (e.g., "slow Sunday with a bike ride") into **concrete markers and cards on the map**, never free-floating text.	(B) Conversational Assistant, (A) Interactive Map
Synchronized highlighting	**Tie chat, canvas, and itinerary back to the map**: selecting a place in one view highlights it everywhere so references stay unambiguous.	Full UI: (A)–(E)

3.1.4 DG4: Build trust with spatially accountable conversation. Finally, to make the system's AI-powered assistance transparent and trustworthy, our fourth goal is to ensure all conversations are spatially accountable. This is a form of direct manipulation, where actions in one part of the interface have immediate, visible consequences in another [67]. When a user asks the *LLM-Powered Local Reasoning Assistant* for suggestions (e.g., planning a slow Sunday with a bike ride), every place mentioned is unambiguously linked to the map via *Synchronized Cross-Component Highlighting*. This tight coupling grounds the abstract language of conversation in concrete spatial reality, making the system's behavior clear and comprehensible.

3.2 The PlaceWeave System in Action

To operationalize these four design goals, we designed PlaceWeave as an integrated system that guides a user from open-ended discovery to a committed itinerary in a single workspace. Figure 2 labels the five main components and how they work together.

3.2.1 Exploring and asking questions. A user's journey typically begins on the **Interactive Map** (A) and in the **Conversational Assistant** (B) (Figure 2). The map presents an overview of all candidate places for the city, with marker size and color encoding popularity and inferred localness (DG3). Users can pan and zoom to orient themselves, then either click markers directly or ask open-ended questions such as "find a quiet study café with an artsy vibe near downtown." The assistant translates these experiential queries into grounded suggestions and returns them as a ranked list. Each suggestion is simultaneously highlighted on the map so that users can immediately see *where* it is, reinforcing spatial accountability (DG4).

3.2.2 Inspecting places and comparing options. When a user clicks a place either on the map or in the chat, the **Insights Panel** (C) (Figure 2) slides open. This panel summarizes localness attributes derived from social video (e.g., typical visitors, ambience, rhythms), grouped by category with confidence badges and short evidence snippets, and provides expandable links to source clips. The panel follows progressive disclosure (DG3): the top section is a scannable summary, while details and low-confidence inferences are available on demand. Users can quickly compare candidates (e.g., two cafés with different crowd patterns) and decide which ones feel aligned with their values and constraints (DG2).

As promising locations accumulate, users can drag them into the **Mental Map Canvas** (D). The canvas functions as a flexible "staging area" where users cluster places, add text notes (e.g., "evening options"), and sketch lightweight connections and flows. This externalization supports sensemaking by helping users see spatial and

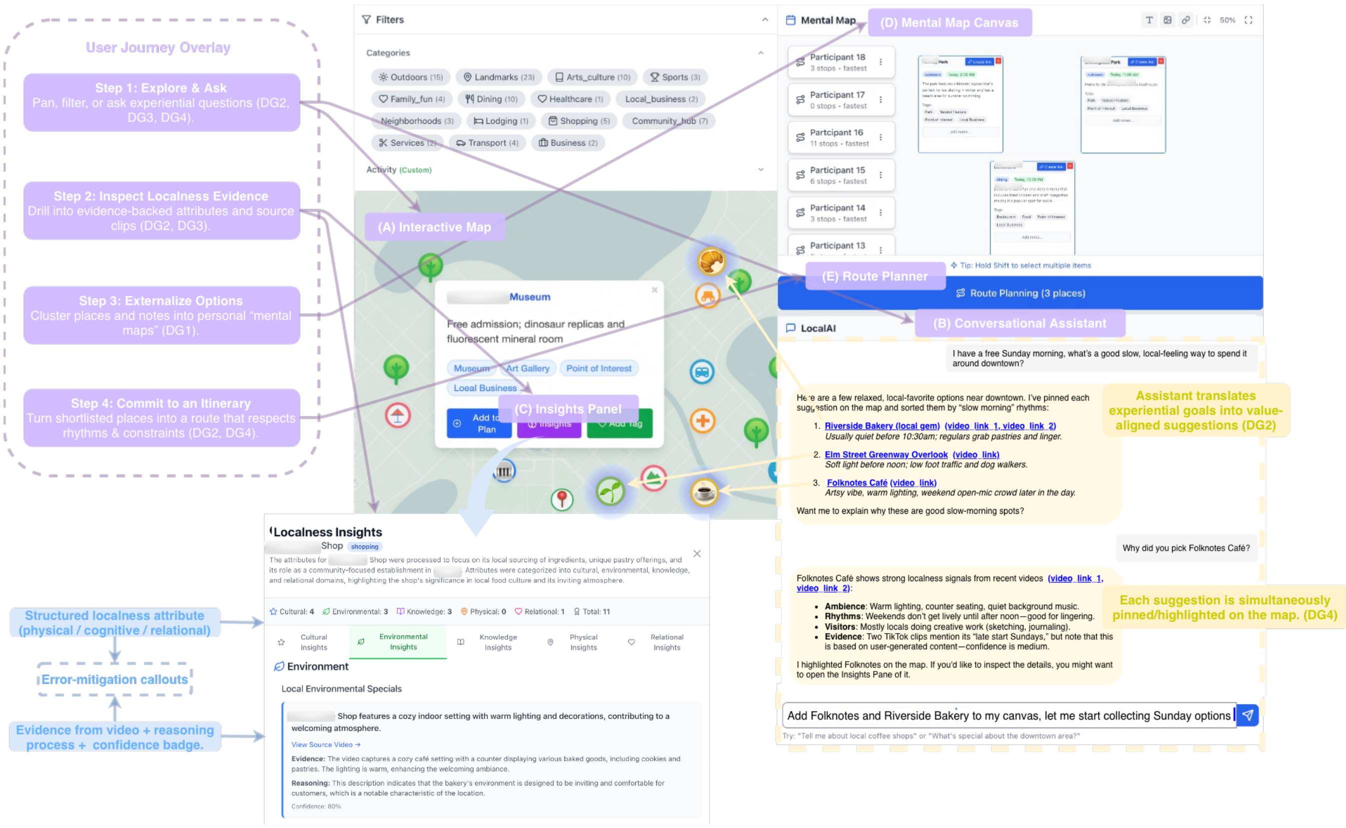

Figure 2: The PlaceWeave interface. (A) The Interactive Map, showing all extracted places with visual encodings for localness and popularity, supports panning, zooming, and selection. (B) The Conversational Assistant supports open-ended, experiential queries; every suggestion is linked to map markers and itinerary entries. (C) The Insights Panel, which opens for a selected place and shows evidence-backed localness attributes, confidence, and links to source videos. (D) The Mental Map Canvas for staging and clustering candidate places, notes, and connections. (E) The Route Planner, which turns a shortlist into a value-aligned, rhythm-aware itinerary. Components are spatially accountable: assistant suggestions, insight selections, and itinerary edits are immediately reflected on the map. In PlaceWeave, users move through four stages: (1) Explore & Ask on the Interactive Map and Conversational Assistant, (2) Inspect Localness Evidence in the Insights Panel, (3) Externalize Options on the Mental Map Canvas, and (4) Commit to an Itinerary with the Route Planner.

thematic relations between options before committing to a plan (DG1). Items on the canvas remain linked to their map markers and insight summaries, so users can move back and forth between overview and details without losing context.

3.2.3 Committing to a plan. Once a user select a rough shortlist in **Mental Map Canvas** (D), they switch to the **Route Planner** (E) by clicking the "Route Planning" button (Figure 3). Here they drag shortlisted places into a sequence and choose planning preferences (e.g., emphasize localness over shortest travel time or keep most stops within one neighborhood). The planner uses the same localness-aware AI engine to suggest a reasonable order, check opening hours, and flag coarse feasibility issues (e.g., long cross-town jumps or visiting a venue before it opens). It relies on lightweight heuristics rather than full combinatorial route optimization: users can always reorder stops, add breaks, or remove places based on their own constraints (Figure 4). The resulting itinerary is synchronized with the map: stops appear as a numbered route, and hovering over a step highlights the corresponding marker and its key localness attributes. In this way, PlaceWeave turns scattered, video-based inspiration into a context-preserving plan while keeping the "why" behind each choice visible from start to finish and still leaving fine-grained logistical trade-offs under user control.

3.3 Data Foundation: Building the Place Knowledge Graph

To realize the design affordances of PlaceWeave, we developed a data processing pipeline that systematically extracts and structures signals of "localness" from multimodal video content. This pipeline, visualized in Figure 5, transforms unstructured social media posts into a queryable knowledge graph that serves as the technical foundation for the entire system.

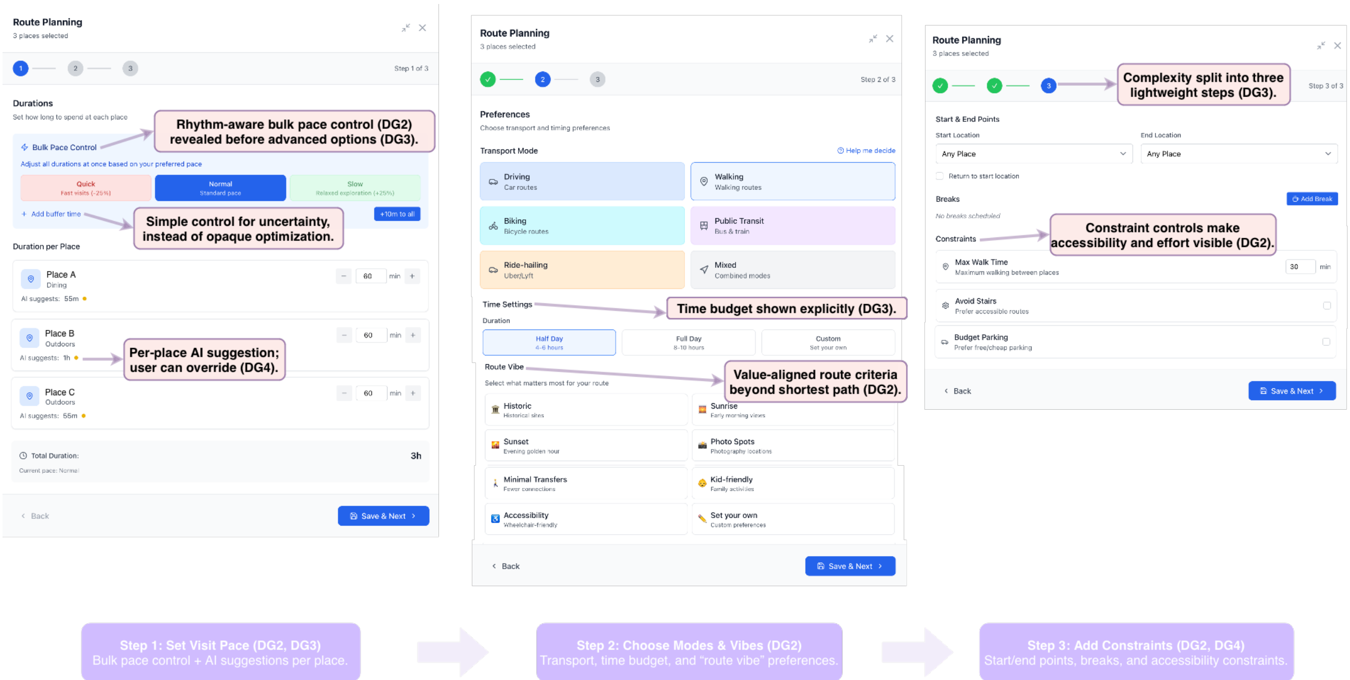

Figure 3: Route Planner configuration wizard. The Route Planner exposes trip construction as a three-step, user-steerable process rather than a single opaque optimization. Step 1 lets users set a global visit pace and adjust AI-suggested durations for each place. Step 2 captures value-aligned preferences such as transport mode, time budget, and "route vibes" (e.g., historic, kid-friendly). Step 3 adds practical constraints like maximum walking time, accessibility needs, and start/end locations. This staged design realizes DG2 (value-aligned exploration) and DG3 (progressive disclosure of planning complexity), with AI suggestions surfaced but always overridable (DG4).

3.3.1 Dataset. Our data collection, conducted from April to August 2025, focused on three locales across the U.S., chosen to represent different place types: a dense urban environment, a sprawling suburban area, and a smaller city serving as a rural hub. Our final dataset comprises 6,250 public videos from 2,960 unique creators. The data distribution reflects real-world content generation patterns, with higher volumes from the more populated urban and suburban sites. This dataset, totaling 78.3 hours of content, serves as the raw material for our pipeline. Complete details of the data distribution, our acquisition queries, filtering, and deduplication methods are available in Appendix A.1.

3.3.2 Stage 1: Multimodal Entity and Relationship Extraction and Alignment. The first stage deconstructs each video post into its constituent informational channels. The goal is to convert implicit, multimodal signals into explicit, machine-readable text and tags. From the visual track, the pipeline first segments each video into shots using a standard boundary detector, then samples both boundary and mid-shot frames. Within each shot, we prioritize frames with higher visual salience (e.g., more faces, objects, or stable views) to better capture the "focus" moments of that segment rather than only hard cuts. On these selected frames, we use established models to identify salient objects, recognize architectural landmarks, and extract on-screen text. Concurrently, we transcribe all spoken narration and tag significant environmental sounds (e.g., *"crowd chatter," "live music"*), keeping their timestamps aligned with the sampled frames. Each extracted feature is time-stamped to ensure temporal alignment across modalities, resulting in a rich set of descriptors that capture what a place looks like, sounds like, and feels like. The specific models and methods used for this process are detailed in Appendix A.2.1.

3.3.3 Stage 2: Geographic Entity Grounding and Verification. With features extracted, the second stage grounds this information by linking it to real-world locations. The process begins by applying named entity recognition (NER) to all textual data to identify potential place names. Recognizing that UGC is often vernacular and ambiguous (e.g., "the Square"), we use GPT-4o to reason over the video's context and filter these mentions. We then employ a robust verification process: canonicalizing entities against the Google Places API and cross-validating them using semantic similarity against Wikidata descriptions. This allows us to reliably map the rich experiential data to a specific latitude and longitude and, critically, to link informal, local aliases identified by the LLM to their formal counterparts. Our specific heuristics and validation outcomes are described in Appendix A.2.2.

3.3.4 Stage 3: Knowledge Graph Construction. The final stage fuses the grounded entities and their extracted attributes into the knowledge graph. This graph is designed to represent not just places, but the rich web of relationships between them. We model geographic

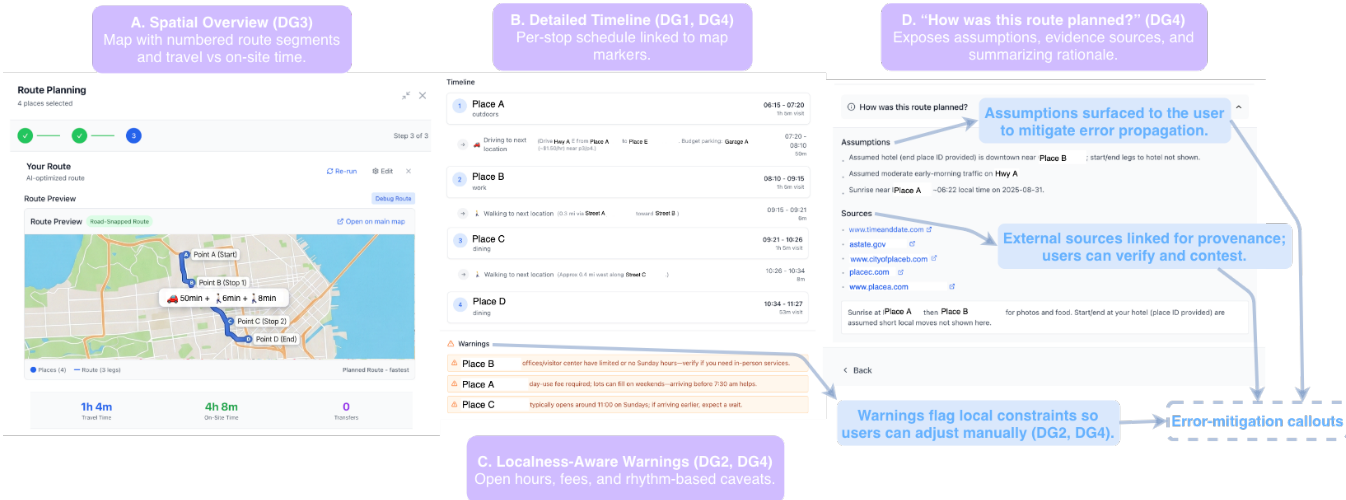

Figure 4: Route Planner outcome view with spatially accountable explanations. The final itinerary combines a map-based overview of the route (A), a step-by-step timeline of stops and travel (B), localness-aware warnings about opening hours and typical rhythms (C), and an explicit "How was this route planned?" panel that lists assumptions and evidence sources (D). Warnings and provenance links make uncertainties and trade-offs visible, supporting value-aligned, trustworthy planning (DG2, DG3, DG4).

locations, experiential attributes (e.g., ambience, activities), and temporal patterns as distinct *nodes*. These are then connected by typed, directional *edges* that describe their relationships, such as `co_occurs(Place, Ambience)` and `rhythm(Place, Time)`. The graph is implemented as a `NetworkX MultiDiGraph` and stored in a PostgreSQL database, enabling both topological and semantic queries. This structured representation allows the system to answer nuanced queries like "find a quiet café (Ambience) for studying (Activity) on a weekday afternoon (Time)," and it instantiates multi-level localness as graph attributes attached to specific places that can be surfaced, combined, and questioned in the interface. The final constructed graph contains 28,341 nodes and 96,507 edges. The complete schema and scale of our constructed graph are detailed in Appendix A.3.

3.4 Localness-Aware AI Engine

Having represented local signals as a place knowledge graph, we next describe how a localness-aware Graph-RAG engine turns this static structure into queryable context for AI assistance. For each conversational turn or planning action, the engine selects relevant nodes and relations from the graph, combines them with semantically similar items, and packages this evidence as input to the LLM. This process drives the generation of localness attributes, map explanations, and itinerary suggestions, tying all AI responses back to concrete places and relationships in the underlying graph.

3.4.1 Hybrid Retrieval and Evidence Packaging. For any AI task, the engine first gathers evidence using a hybrid retrieval strategy. A *graph-based retriever* walks the knowledge graph to collect topologically adjacent nodes that reflect explicit relationships (e.g., a café's links to ambience, activities, and time-of-day, or repeated co-occurrence with a nearby bookstore). In parallel, a *semantic retriever* searches across all nodes for textually similar items (e.g., other places described as a "quiet local spot"). The retrieved items are merged, deduplicated, and formatted as a compact context bundle that is passed to the LLM. This ensures that generation is conditioned on both the structural patterns in the graph and semantically related examples, rather than on isolated text snippets alone [46, 77].

3.4.2 Grounding Localness Attributes and Interactive Features. Using this evidence bundle, the engine produces two kinds of outputs.

First, it generates structured *localness attributes* for the **Insights Panel**. Guided by Gao et al. [29]'s "Localness Conceptual Framework," we prompt the LLM to infer attributes in three domains — *Physical* (direct interaction with a place), *Cognitive* (local knowledge and cultural understanding), and *Relational* (social connections and emotional bonds). For each attribute, the model returns a short description, cited evidence, reasoning process, and a confidence score. These framework-aligned attributes support interpretation of a place's "localness" and provide the basis for our technical validation (Sec. 3.5) and user study measures.

Second, the same engine powers the user-facing features. The **Conversational Assistant** uses graph-guided retrieval to answer experiential queries (e.g., "study-friendly cafés near the river that locals use on weeknights") with suggestions that come with rationales and links back to source videos. The **Route Planner** uses evidence about spatiotemporal rhythms and co-occurrence (e.g., which places are frequently visited together, typical crowd levels by time) to propose value-aligned itineraries and explain trade-offs. All suggestions are returned with coordinates so they can be immediately highlighted on the map, keeping conversational answers spatially accountable.

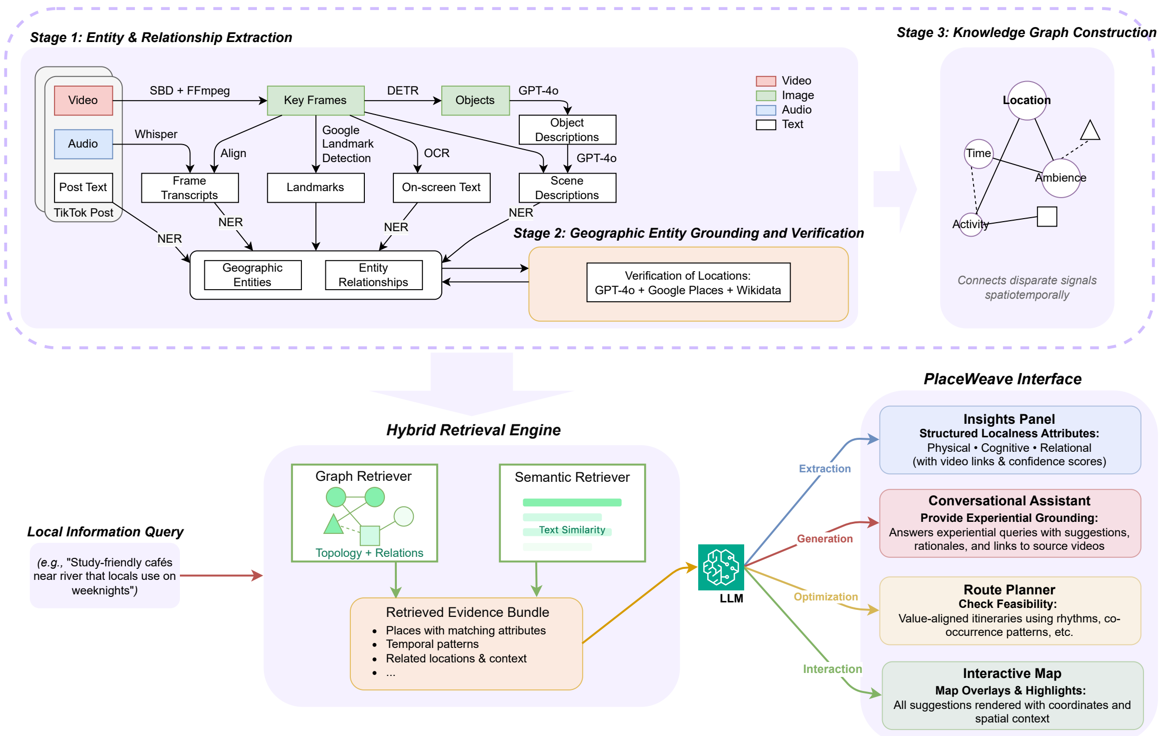

Figure 5: End-to-end data pipeline powering PlaceWeave. Raw social posts undergo (1) multimodal entity extraction and (2) geographic verification before being synthesized into (3) a spatiotemporal knowledge graph. At query time, (4) the Localness-Aware AI Engine employs a Graph-RAG approach, merging topological and semantic retrieval to condition the LLM. The engine's outputs are mapped to specific interface elements: structured localness insights, grounded generation for conversational answers, constraint-based optimization for itinerary planning, and synchronous map interactions.

In principle, PlaceWeave could index transcripts and captions with a standard text-only RAG pipeline. We instead adopt a Graph-RAG approach because many target queries require relational consistency and temporal grounding—that a set of cues describes the same place (and often the same time-of-day/weekday rhythm)—rather than loose keyword co-occurrence. For instance, a text-only index may retrieve separate clips that mention "quiet," "studying," and "weekday evenings," yet still leave it ambiguous whether these cues refer to one café with a consistent rhythm or to different places discussed across creators. By structuring extracted signals as a place knowledge graph, we can retrieve candidates that satisfy conjunctive constraints (e.g., `Place` linked to `Ambience=quiet`, `Activity=studying`, `Time=weekday evening`) and then follow relations across multiple visits and co-mentioned entities to aggregate evidence and provide grounded context. This graph-guided retrieval is what allows the assistant, map overlays, and route planner to answer experiential localness queries while keeping supporting videos and uncertainty visible.

Our current Route Planner focuses on localness-aware feasibility rather than full spatial optimization. Unlike tour-planning systems that formulate itinerary construction as a constrained optimization problem over points of interest and time windows (e.g., Yahi et al. [105]'s interactive tour planner or Tang et al. [93]'s LLM-plus-optimizer framework), PlaceWeave uses heuristic ordering informed by graph structure and opening hours to avoid obvious backtracking and temporal conflicts. We see this as complementary: the Graph-RAG backbone grounds itineraries in community rhythms and local practices. Future work could couple it with more powerful optimization backends for users who prioritize strict efficiency. The detailed prompting frameworks are available online.[1]

3.5 Technical Validation of the Pipeline

Having detailed the pipeline, we next evaluate whether its outputs are accurate and grounded enough to be useful in a user-facing system. This technical validation sits between infrastructure and interaction: we use it to assess how well inferred localness attributes match human judgments and where they fail, in order to ensure that our interface designs can effectively surface uncertainty and provenance rather than hiding them. We compare model-generated attributes to a human-annotated ground truth and analyze hallucination-prone components, then describe UI-level strategies for mitigating their impact on trip planning.

[1]https://osf.io/ymu28?view_only=2a80656d3f224fed9b210c96e02bb55a

Table 2: Technical validation results for LLM-generated localness attributes against a human-annotated ground truth. The pipeline demonstrates high performance on concrete, experiential attributes central to the system's goals, with understandable limitations on more abstract concepts where source data is sparse.

Localness Component	Description	Cov.	Halluc.	Sim.	Items
Localness Domain: Cognitive					
Food Culture	Understanding and appreciation of the place's culinary traditions.	1.00	0.14	0.85	30
Historical Knowledge	Understanding the place's past and how it has evolved over time.	0.94	0.11	0.82	18
Local Recommendations	Having detailed knowledge about local establishments and services.	1.00	0.08	0.89	25
Hidden Gems	Insider knowledge of unofficial landmarks, local secrets, and special tips.	1.00	0.11	0.87	22
Localness Domain: Physical					
Being Born/Native	Claim to localness by birth, with lifelong familiarity and deep historical ties.	1.00	0.05	0.91	40
Long-term Residence	Accumulated experience and commitment from prolonged living in the area.	1.00	0.06	0.90	48
Environmental Experience	A personal, physical connection to natural spaces and local landscapes.	0.92	0.20	0.81	12
Localness Domain: Relational					
Civic Engagement	Participating in local governance and political processes in the area.	0.89	0.43	0.55	35
Sense of Belonging	Feeling comfortable and accepted within the place.	0.97	0.12	0.88	37
Community Investment	Demonstrating care and commitment to the well-being of the area.	1.00	0.04	0.92	21
Personal Relationships	Strong local social networks, friendships, and neighbor recognition.	0.98	0.09	0.86	44

Note: For brevity, only a representative subset of components is shown. Cov. = Coverage, Halluc. = Hallucination Rate, Sim. = Semantic Similarity.

3.5.1 Validation Methodology. For a stratified sample of 150 places from our dataset, we first used our pipeline's hybrid retrieval method to gather all relevant evidence from the knowledge graph. This retrieved evidence was then presented to two trained human coders who, following the same Localness Conceptual Framework [29], independently annotated the localness attributes for each location. The initial inter-rater reliability is $\kappa = 0.74$. After reconciling their annotations, they established a final ground truth dataset comprising 899 distinct attributes. We then compared the LLM-generated attributes against this ground truth using three metrics: (1) **Coverage** (the proportion of human-annotated attributes also identified by the LLM); (2) **Hallucination Rate** (the proportion of LLM-generated attributes not supported by either the ground truth or the retrieved evidence); and (3) **Semantic Similarity** between the descriptions of the LLM and human attributes.

3.5.2 Performance on Experiential Attributes. As shown in Table 2, the pipeline performs strongly on its primary goal of extracting concrete, experiential attributes. For components like *"Local Recommendations"* and *"Hidden Gems"*, hallucination rates were low (≤ 0.11) and semantic similarity was high (≥ 0.87). However, performance degrades on abstract concepts sparsely represented in the data, such as *"Civic Engagement"* (0.43 hallucination rate). To mitigate the risk of misleading users, we surface this uncertainty directly in the UI with "Low Confidence" badges and provide provenance links to the source videos. This empowers users to act as the final arbiter of truth, mitigating the risk of unverified information influencing their plans. This approach directly addresses the challenge of technical validation not always translating to user benefit. By making pipeline limitations transparent and auditable, we treat our pipeline not as the ultimate source of truth, but rather as a sociotechnical source of information designed to foster critical engagement from users.

3.5.3 Design for Mitigating Error Propagation. The multi-stage data processing pipeline is susceptible to cascading errors, where a mistake in an early stage (e.g., OCR) becomes an immutable "fact" in a later one (e.g., the knowledge graph). Rather than treating this solely as a modeling issue, we designed a set of safeguards that combine backend heuristics with frontend transparency to keep users aware of uncertainty.

On the data side, we apply three robustness strategies:

(1) *Multimodal entity verification:* A geographic entity is only promoted to a high-confidence node in our knowledge graph if it is verified through at least two different modalities. For example, a restaurant name must be detected in the audio transcript and either as on-screen text via OCR or as a geotagged entity in the post's metadata.
(2) *Evidence triangulation for attributes:* An experiential attribute (e.g., "lively ambience") requires corroborating evidence from multiple channels or repeated observations before being attached to a place.
(3) *Confidence-based graph pruning:* After construction, we prune our knowledge graph, removing edges and nodes that have only weak or conflicting evidence. This ensures our final knowledge graph—which gets used to inform LLM generation—primarily contains reliable signals of localness.

These backend strategies work in concert with UI-level mitigations, such as confidence badges, provenance links, and exposed rationales, to ensure that the information presented to the user is as reliable and transparently grounded as possible (see Figure 2 and Figure 4). In short, we aim not only to reduce errors, but also to make remaining uncertainty visible so that users can question, contest, and correct the system when needed, following best practices [74].

Table 3: Demographics and employment information for the study cohort (n=18).

Attribute	Distribution (count; % of n=18)
Age	18–25: 5 (27.8%), 26–35: 8 (44.4%), 36–45: 1 (5.6%), 46–55: 4 (22.2%)
Gender	Female: 10 (55.6%), Male: 7 (38.9%), Non-binary: 1 (5.6%)
Education	High school: 1 (5.6%), Some college: 2 (11.1%), Bachelor's: 7 (38.9%), Master's: 6 (33.3%), Doctoral: 2 (11.1%)
Travel-planning experience	Extensive: 12 (66.7%), Some: 6 (33.3%)
Employment / status	Working professionals: 13 (72.2%), Working/Student: 2 (11.1%), Students (only): 3 (16.7%)
Representative roles	Administrative staff, Manufacturing engineer, Geospatial analyst, IT support, Finance manager, Pharmacist, Medical imaging technologist, Consultant, Librarian, Clinical research staff, Media producer, Program administrator

4 User Study

In this section, we examine how PlaceWeave changes people's trip-planning practices compared to a realistic multi-tool workflow. We adopt a within-subject *Comparative Structured Observation* (CSO) design [66], in which participants complete the same structured planning task in two contrasting environments: PlaceWeave and a baseline workflow that combines ChatGPT, TikTok, Google Maps, and a separate canvas tool (Draw.io). CSO is well-suited here because it foregrounds careful observation of how different tools shape behavior under comparable conditions. Our evaluative claims are therefore about workflows, perceived localness and authenticity of the planned itineraries, and trust and confidence in the tools. We next describe our tasks, measures, and analysis methods, focusing on how they let us assess planning strategies, perceived localness, sense of neighborhood character, and trust in the system.

4.1 Study Design, Task, and Conditions

4.1.1 Experimental Design. We designed our study to be a *within-subject, counterbalanced experimental design*, to rigorously evaluate our system against a strong, ecologically valid baseline. Each participant completed a travel planning task under two distinct, time-boxed conditions. This design allows for a direct comparison of the two environments while controlling for individual differences in planning strategies. To mitigate learning or fatigue effects, the order of the conditions was counterbalanced across all participants.

4.1.2 Planning Task and Scenario. The central task was designed to elicit open-ended, exploratory planning behavior. Each participant was randomly assigned two of these three sites. For each assigned site, they were presented with the following scenario: *"A close friend is visiting you for a weekend. They want to avoid the typical tourist spots and have an authentic "local" experience. Your task is to use the provided tools to create a one-day itinerary that captures the unique vibe of the city."* Participants were asked to produce a final itinerary containing at least three distinct places or activities.

4.1.3 Experimental Conditions. Participants performed the planning task in two different environments, each representing a distinct planning paradigm.

PlaceWeave Condition. In this condition, participants used the fully-featured PlaceWeave system. All functionality, such as discovery, spatial mapping, and itinerary construction, was available within a single, unified desktop browser window. This condition represents our proposed integrated environment, where the AI's suggestions are directly grounded in the system's place-based knowledge graph and are immediately actionable on the map and canvas.

Baseline Condition. The baseline was designed to model a fragmented real-world workflow. To create a fair and direct comparison with PlaceWeave's core features, we provided participants with a suite of tools for each function: (1) the web version of ChatGPT (GPT-4o) in a dedicated browser tab for generating ideas and answering queries, (2) the native TikTok mobile app on a provided smartphone for discovery, (3) the web version of Google Maps in a second browser tab for mapping and routing, and (4) the web version of Draw.io in a third browser tab. This tool was specifically chosen to provide a direct, feature-comparable alternative to PlaceWeave's Mental Map Canvas, allowing participants to visually organize their plans if they wished.

4.2 Participants

Following university IRB approval, we recruited 18 participants (10 female, 7 male, 1 non-binary) via campus mailing lists. Participants ranged in age from 18 to 55 (M = 31.4, SD = 9.8). The cohort represented a mix of working professionals and students from various fields, ensuring a range of perspectives. All participants provided informed consent and were compensated for their time through a raffle with a 1-in-5 chance to win a 30 Amazon gift card. Detailed demographic information is provided in Table 3.

4.3 Procedures

Each participant attended a single 60-minute session conducted remotely via video conference. To ensure system stability and consistent interface rendering, participants accessed the system by remotely controlling the researcher's computer. The session was structured into three parts: orientation (5 minutes), planning task interface use (40 minutes), and our post-task evaluation (15 minutes). We began the session by providing participants with a brief orientation, wherein we introduced the study tasks and provided a short tutorial. Participants then performed the planning task for two different sites, using a different experimental condition each time. During both tasks, they followed a think-aloud protocol, verbalizing their strategies and decision-making processes. Upon completing both planning tasks, we asked participants to complete our final usability questionnaire and then conducted a short semi-structured interview comparing their experiences across both conditions.

To mitigate order effects, the assignment of both the tool condition (PlaceWeave first vs. Baseline first) and the site type (e.g., City U with PlaceWeave vs. City U with Baseline) was fully counterbalanced across all 18 participants (Appendix B). Audio from the planning tasks were recorded for later analysis.

4.4 Measures

Our mixed-methods approach focused on participants' experience of planning and their perceptions of the itineraries they produced. Our approach combined quantitative and qualitative data collection:

4.4.1 Quantitative Measures. We administered the standard *System Usability Scale (SUS)* [9]. To capture dimensions of the experience specific to our system's goals, we developed four custom 5-item, 5-point Likert scales. Each item asked participants to reflect on (a) the itinerary they had just created and (b) their subjective experience using the interface in that condition. The scales assessed:

- *Creativity and Planning Support:* The degree to which the interface inspired users and supported their natural planning process (e.g., "I felt inspired while using this interface.").
- *Recommendation Quality:* Users' trust in the system's suggestions and their perceived relevance (e.g., "The recommendations helped me discover places I wouldn't have found otherwise.").
- *Interface Satisfaction:* Overall satisfaction and willingness to use or recommend the system (e.g., "I would choose this interface over other travel planning tools.").
- *Perceived Localness and Authenticity:* The extent to which the suggested places and resulting itinerary *felt* aligned with local practices and neighborhood character (e.g., "The suggested places felt like places locals would actually go.").

The full instruments for all quantitative measures are available online.[2] Taken together, these instruments let us compare how each condition shaped perceived localness, authenticity, and confidence during planning.

4.4.2 Qualitative Data. The primary sources of qualitative data were transcripts from the think-aloud protocol during the planning tasks and the post-task semi-structured interview. The interview guide (see OSF link[2]) prompted reflection on planning strategies, the effectiveness of specific interface components, perceptions of authenticity *of the planned itineraries*, and direct comparisons to the baseline workflow. These accounts allow us to interpret the quantitative measures and situate them in participants' own notions of "local" trips.

4.5 Analysis

Our analysis employed a convergent mixed-methods approach to develop a comprehensive understanding of the user experience. From the numerical evaluation questionnaires, we calculated descriptive statistics (mean, standard deviation, confidence intervals) for the SUS scores and for each of our four custom experience scales, and we compared SUS scores against established standards to interpret usability.

[2]https://osf.io/qhn8u?view_only=2a80656d3f224fed9b210c96e02bb55a

To complement these quantitative results, we analyzed our qualitative data using reflexive thematic analysis following Braun and Clarke [8]. The audio recordings from the think-aloud protocols and semi-structured interviews were transcribed verbatim. Two researchers then independently coded the transcripts, generating initial codes related to planning strategies, trust in the system, interface strengths and limitations, and perceptions of "localness." In line with Braun and Clarke's emphasis on themes as analytic outputs rather than discoveries, we treated theme development as an active, interpretive process: the coding team met to compare codes, discuss divergences, and iteratively construct a shared set of themes that captured patterned meanings across the dataset. The initial inter-rater agreement was $\kappa = 0.71$, after which coders reconciled differences through discussion.

Finally, to triangulate overall takeaways, we compared qualitative and quantitative findings. The statistical results provided a broad measure of the system's success on usability and perceived localness, while the generated themes offered rich, contextual explanations for *why* those results occurred.

5 Findings

We now present our findings from the comparative study, focusing on how PlaceWeave shaped participants' workflows, perceived localness of their plans, and confidence relative to the baseline tools. We first summarize overall reception, then report three themes that we generated through our thematic analysis: (1) how the unified workspace changed planning compared to juggling multiple apps, (2) how localness attributes and provenance reshaped participants' understanding of neighborhood character and trade-offs during planning, and (3) how the system still fell short and surfaced tensions around complexity, serendipity, and platform lenses. Throughout, we relate these themes back to our design goals (DG1–DG4) and research questions (RQ2–RQ3).

5.1 Positive User Reception and High Perceived Authenticity

System Usability Scale (SUS) scores indicated robust usability with a mean of 74.2, 95% CI [69.0, 79.4], SD = 12.8; 76% of participants scored ≥ 68, exceeding the commonly cited acceptability threshold [5, 6, 59]. Across the four constructs (5-point Likert), participants rated the interface well above the neutral midpoint: *Perceived Localness* 4.00 [3.77, 4.25], *Recommendation Quality* 4.22 [4.02, 4.42], *Creativity Support* 4.19 [3.95, 4.41], and *Satisfaction* 4.10 [3.78, 4.39].

These ratings reflect participants' expectations, formed from the on-screen evidence and narratives, about how "local" their itineraries would feel if enacted, as well as their experience using the interface itself. Qualitative feedback echoed these ratings: participants described PlaceWeave as "more convenient" and "visually appealing" than their usual planning workflows, and several explicitly said they would choose it over existing tools for future trips. We unpack these responses in the themes below.

5.2 Theme 1: From Fragmented Toolchain to a Unified Planning Workspace

DG1 aimed to provide a visual planning canvas that integrates multimodal information—interactive maps, "mental map," timelines, and *social media content—to help users externalize their plans and understand a place's spatial and experiential context. DG3 was to create an integrated planning hub combining chat, maps, media, and itinerary management in one interface, thereby streamlining a workflow that is often fragmented across multiple apps.*

Theme 1 draws primarily on DG1 and DG3, examining how the integrated map–timeline–canvas workspace changed participants' planning strategies relative to the baseline toolchain.

5.2.1 A synchronized map–timeline–canvas workspace externalizes spatial reasoning (DG1). We found that PlaceWeave's synchronized map-and-timeline interface functioned as a "thinking tool" for spatial reasoning rather than just a static container. Participants overwhelmingly used the map-based canvas as their primary workspace (e.g., "This map I used most"—P1) and reported that seeing all points of interest laid out geographically and temporally made it easier to "figure out how to navigate the system" and "visualize where I'm going."

Several remarked that having an integrated map with contextual icons and an itinerary timeline helped them form a mental picture of their day's route, which was "visually appealing" compared to their usual planning in spreadsheets or separate apps. P1 noted, "I usually...plan in an Excel document...this is much more convenient," and described the system's suggested ordering and travel-time estimates as "automatic route optimization" that spared her a "very annoying" manual calculation. This indicates that rich, synchronized multimodal representations can ground users' understanding of a new locale more effectively than baseline tools that silo maps, lists, and media.

5.2.2 All-in-one hub reduces fragmentation and cognitive load (DG3). In our study, this all-in-one design proved to be a major advantage of PlaceWeave over the baseline toolkit. Nearly every participant (17 of 18) explicitly contrasted PlaceWeave's unified interface favorably with their usual practice. P10, for instance, listed the four different applications she would normally bounce between to plan a single day (Google Maps for tourist sites, Google Search for timings, back to Maps for routing, TikTok for on-the-ground videos) and then noted that PlaceWeave "can streamline this process by integrating and automating [the] comprehensive planning process" (P10).

Participants could chat with the AI to get ideas, see those ideas on the map instantly, and add them to a growing itinerary timeline all within the same interface. This fluid, uninterrupted workflow is something participants indicated they would choose over their existing tools: in the post-task survey, they agreed that "I would choose this interface over other travel planning tools" (mean 4.50/5) and that they would recommend PlaceWeave to others planning travel (mean 3.84/5). In short, PlaceWeave's tightly coupled, integrated design clearly outperformed baseline toolchains in supporting spatial sensemaking and plan organization by keeping all relevant information and actions together.

5.2.3 Confidence, control, and engagement in a unified workspace (DG3). The impact of the centralized design on user confidence and efficiency was also evident in qualitative behavior. Many participants, after using PlaceWeave, indicated they felt more in control and less overwhelmed by information. P9 remarked that normally she keeps mental notes or separate lists when using Google Maps alongside social media, but with everything embedded in one interface, she could focus on exploring ideas rather than managing windows or context-switching.

This cohesive experience translated into higher *usability* perceptions: participants rated overall satisfaction at 4.10/5 and found the interface engaging to use (4.36/5), suggesting that the hub approach did not just improve functionality but also enjoyment. The tight coupling of the tool's features allowed users to iteratively refine their plan in a mode of interaction that felt natural and empowering. This is supported by the high ratings of "I trust the recommendations provided by this interface" in the Recommendation Quality survey (mean 4.47/5). P13 summed it up: "I feel like I have a digital travel assistant and map all in one. I'm not doing the legwork of cross-checking everything."

5.2.4 Clarity costs: clutter, hierarchy, and focus modes (DG1, DG3). At the same time, bringing everything into one workspace introduced challenges around information density and visual hierarchy. The unified display could feel like "a lot of information...little dots and lines and AI stuff" on the screen at once (P1), and some features like the Mental Map Canvas were underutilized.

Four participants mentioned that the default map view (with all extracted places from social media, shown as icons) could feel overwhelming or cluttered at first glance. Important locations were not visually distinguished from minor ones, leading P7 to comment: "I know that [Location A] is one of the focal points of the city… But [on the map] they have the same size icon as, for instance, a barbershop." In baseline Google Maps, certain landmarks are emphasized by larger labels or pins; in PlaceWeave's rich map, everything was initially uniform, requiring users to click around to figure out significance.

Participants intuitively tried direct manipulation interactions, like dragging and linking items, that were not fully supported, revealing an opportunity to better align the UI with users' mental models. Together, these observations suggest that while an integrated workspace supports experiential planning (DG1, DG3), future designs should better manage information density and provide stronger visual hierarchy and focus modes, in line with the principles of progressive disclosure and details-on-demand [88].

5.3 Theme 2: Localness Attributes, Community Media, and Trust in Plans

DG2 sought to embed community-sourced content, primarily short-form social videos, into the planning process to preserve the rationale behind recommendations and align itineraries with authentic local practices. DG4 centered on a conversational local AI assistant intended to act like a "digital local friend," allowing users to ask natural questions and get personalized, context-rich recommendations.

Theme 2 centers on DG2 and DG4, examining how community video narratives, localness attributes, and provenance influenced perceived authenticity, discovery, and trust in PlaceWeave, compared to baseline tools.

5.3.1 Community video narratives increase perceived authenticity (DG2). PlaceWeave's use of social videos and graph-based place knowledge notably enhanced participants' sense of a plan's authenticity and provided rich explanations *in situ*. Quantitatively, participants rated the system very highly on localness and authenticity of recommendations (e.g., mean 4.40 for "suggested locations seemed authentic and locally oriented" on a 5-point scale). These high *Perceived Localness* scores were directly attributed to the incorporation of social video data.

In interviews, 16 of 18 participants praised the "nuanced, insider knowledge" coming from local social media posts. They found the video-driven narratives more genuine and context-rich than the decontextualized ratings and reviews typically found on Google or Yelp. "I feel like the social media part is a little bit more real… It doesn't feel like it's coming from the tourist board telling me "you should go here because it's the touristy thing to do"" (P6). Instead of generic descriptions, PlaceWeave surfaced details like a place's vibe or insider tips (e.g., "the line is long at this bakery on weekends"—a TikTok-derived insight noted by P9) that participants considered more useful and trustworthy for experiencing the "everyday texture" of the locale.

5.3.2 From "off-the-beaten-path" gems to broadened horizons (DG2). Integrating community videos not only improved the perceived authenticity of individual suggestions but also helped participants discover hidden gems beyond the usual tourist spots. Participants commented that PlaceWeave excelled at unearthing unique places they "would have otherwise missed" (P12), confirming a core advantage of DG2. *Recommendation Quality* was rated highly (participants agreed that "the recommendations helped me discover places I wouldn't have found otherwise," mean 4.18/5).

As one participant put it after seeing a suggested local eatery tucked away in a neighborhood, "I did not know that [place] existed… I like that it identified places I didn't know about" (P1). This sentiment was common—15 of 18 participants felt PlaceWeave broadened their horizons with "off-the-beaten-path" finds. P3, an avid traveler, reflected: "What I would want is someone else to do all of that [research] for me…usually [that] would be AI… I feel like the tools are getting better every day at being able to do that…you can get this AI tool to collect that data for you." Participants thus saw the AI-curated local content as a proxy for having a knowledgeable local friend or concierge scouring social media for under-the-radar experiences.

5.3.3 Conversational local AI as a "digital local friend" (DG4). Our findings show that the conversational agent was indeed beneficial in guiding the planning process and was regarded as a unique strength of PlaceWeave relative to baseline tools. Nearly half of participants (8 of 18) began their planning session by engaging with the local AI chat, using it to brainstorm and generate ideas for their itinerary. Participants found that the AI could handle nuanced, experience-oriented queries that would stump a conventional search engine or map interface. For example, P9 asked for "a fun place for dinner to see the sunset"—a query involving subjective criteria—and the AI successfully suggested a waterfront dinner spot with an explanatory narrative, something she noted "Google Maps can't really do because it doesn't know about "sunset vibes"."

Similarly, P3 used the chat to find activities suitable for children and was impressed that the assistant responded with tailored ideas (e.g., a children's museum and a playground, with context about kid-friendly features). Participants remarked that the AI's narrative style of response (telling a bit of a story about each place, why it's interesting, what locals say about it) made the information "more trustworthy and contextual than just a blurb on a website" (P6). Survey results reflected this: participants agreed that "the interface inspired me" (mean 4.0/5) and "encouraged me to explore new or unexpected places" (mean 4.28/5), indicating that AI-driven discovery sparked curiosity and creativity beyond what baseline tools afforded.

5.3.4 Tight coupling: from ideas to map to itinerary in one move (DG4). The integration of the conversational AI with the rest of the interface further reinforced trust and usefulness. Several participants loved that when the AI mentioned a place, that place would simultaneously be highlighted on the map and could be added to the itinerary with one click. "When the AI offered a compelling idea, [my] immediate impulse was to add it to [the] itinerary on the spot," with multiple users requesting an explicit "Add to Plan" button in the chat for convenience (P9).

This tight coupling meant that the AI truly acted as a partner in the planning process—participants would ask a question, get a suggestion, and then immediately see how it fits into their spatial plan. In baseline workflows, by contrast, acting on a travel suggestion involves several disjoint steps (e.g., read about a place on a blog, locate it on Google Maps, then save it). PlaceWeave collapsed those steps into a smooth dialogue, and participants who heavily used the Local AI tended to give the system top marks on satisfaction.

5.3.5 Transparent rationale and provenance build calibrated trust (DG2, DG4). Crucially, PlaceWeave did not treat social video insights as opaque magic. Localness attributes and itineraries were backed by visible rationale and provenance, which significantly boosted users' confidence while allowing them to stay appropriately skeptical about AI. Participants lauded the route planning component for translating local context into actionable itineraries, with a "rhythm-aware" timeline that annotated plans with venue hours, typical crowd levels, and "warnings" or "assumptions." P5 noted: "I really liked the feature that told me… the closing time of the building, and if there were any warnings and assumptions… that was really nice, because sometimes I'll… look up a place [later] and find out it closed like 2 hours ago."

Participants also appreciated access to "behind-the-scenes" reasoning. P7 highlighted the "How was this route planned?" panel (listing reasons for the chosen route order, drawn from the local knowledge graph): "I really like this portion…It gives the user the behind-the-scenes [info] to make it more trustworthy." P10 valued that the timeline provided "experience-related information [and] experiential descriptions, like what locals normally do on the routes," which "provide more local context" to each day's plan.

At the same time, participants recognized that AI can hallucinate or overstate. P1 said they "trust this AI a bit more [than generic tools] because it's pulling information from humans" but also that they "know AI… hallucinates quite a lot." Here, confidence indicators and source links played an important role. P1 singled out the "little metadata confidence levels" as "fun and actually helpful…they tell me how confident [the AI] is, which can add to my ability to determine how confident I should be [in the information provided]." Participants appreciated being able to click through to source videos to verify surprising claims or treat low-confidence suggestions cautiously. In the baseline workflow, participants opened short-form videos frequently as the main surface for exploration and vibe-checking. In PlaceWeave, participants opened source videos less often overall and typically when they wanted to confirm details: 6 participants clicked through mainly to verify low-confidence suggestions or to inspect the source before adding a stop. This indicates that creator media remained accessible, while the most common reason for opening videos moved toward verification. These patterns suggest that DG2 and DG4 were achieved: localness-aware content, coupled with visible rationale and provenance, fostered strong but calibrated trust.

5.4 Theme 3: Limits of the Lens and Open Design Challenges

Despite the generally positive reception, participants also surfaced important limitations and tensions: information density and control, the platform lens of TikTok, and the need to balance serendipity, logistics, and community concerns.

5.4.1 Information density, affordances, and control. As noted in Theme 1, some participants experienced information overload in the default view, where all extracted places appeared as similar icons. This made it harder to quickly identify what mattered and limited the perceived usefulness of some components (such as the Mental Map Canvas) that were less foregrounded. Participants' instinct to drag, rearrange, and link elements directly on the canvas suggests that future versions should support more fluid, direct manipulation, making it easier to impose personal structure on a dense information space.

These findings highlight a broader design challenge: when AI systems successfully aggregate rich local information, they can produce clutter as well as clarity. Designing for localness thus requires not only better retrieval and representation, but may also necessitate novel interaction mechanisms for users to prune, reweight, and reshape what information the system surfaces, and when the system surfaces it.

5.4.2 Platform lenses and the scope of "local". Participants were also aware that localness inferred from TikTok reflects a particular slice of a community. P7 observed, "You're getting younger people's feelings about the city on TikTok. But you wouldn't get my age group's…TikTok is one platform where like-minded people are grouped. If you go to Facebook…it could be a different map." P3 similarly noted that while the approach was compelling, it likely mirrored the demographics and cultures of a single platform.

Their feedback points to a limitation in the current system, as well as a clear opportunity: our system's success with TikTok validates the value of localness-aware aggregation, but also motivates expansion to multiple sources and explicit communication of provenance. Participants stressed that as more sources are added, the system should continue to label origins (e.g., "From a TikTok video," "From a Reddit thread") so that users can judge context and trustworthiness. This aligns with seamful design arguments that exposing system seams and data origins can enhance understanding and trust [14, 15, 26].

5.4.3 Balancing serendipity, logistics, and community concerns. Consistent with the shift we observed in video use (from baseline exploration to more selective verification in PlaceWeave), participants described a tension between staying open to unexpected finds and efficiently converging on a coherent itinerary. On the one hand, they valued that PlaceWeave "did all the research" and offered coherent itineraries that respected hours and typical rhythms. On the other hand, their concerns about platform lenses and their occasional cross-checking of AI suggestions suggest a desire to remain connected to original sources and to explore beyond what the system proposes.

While our study did not directly measure impacts on local communities, participants' sensitivity to whose voices are represented (e.g., younger TikTok users) raises questions about how such systems might be designed to surface diverse local perspectives and avoid reinforcing narrow or extractive forms of tourism. These open questions point to broader design implications: localness-aware, graph-grounded AI can support more context-aware place-based technologies, but needs to be paired with careful choices about data sources, transparency, and interaction design so that newcomers' plans remain accountable to the communities they are visiting.

6 Discussion

Building on our findings in Sec. 5, we interpret how PlaceWeave supports localness-aware trip planning and what it suggests for place informatics and human–AI interaction. Our claims are about how the system shapes *planning work* and *perceived localness during planning* in a Comparative Structured Observation study, rather than validated on-the-ground outcomes. We focus on three themes: (1) how our pipeline turns social video traces into *inspectable, situated evidence* for place-based reasoning; (2) how the integrated interface supports planning as *interactive sensemaking* rather than one-shot itinerary generation; and (3) what our results imply about limitations, bias, and participatory governance when aggregating community media into "local intelligence."

6.1 From Traces to Situated Evidence: Advancing Place Informatics

6.1.1 Synthesizing local knowledge from multimodal social traces. PlaceWeave broadens the evidentiary basis of place informatics by treating short-form social video as analyzable *social traces*, not merely as inspirational media. Unlike POI databases and text reviews, videos encode multimodal, situated cues—what places *look* and *sound* like, how people move through them, and when particular rhythms (e.g., crowding, lighting, live music) emerge. This enables extraction of experiential signals such as ambience, social roles, and temporal patterns that are difficult to recover from ratings or sparse textual descriptors. Where prior systems often infer place structure from check-ins or other scalar traces (e.g., *Livehoods* [19]) or map affect onto routes [80], our pipeline leverages multimodal evidence to surface vernacular, practice-oriented descriptors that participants associated with "local-feeling" plans. In this sense, our contribution is less a claim of delivering authenticity than a shift in *what counts as evidence* for planning: from decontextualized POI metadata toward situated, community-produced traces of place experience. *This richer evidence base motivates the need for a* *representation that makes localness inspectable and traceable rather than absorbed into opaque summaries.*

6.1.2 Structuring localness with a provenance-linked knowledge graph.

A core contribution of PlaceWeave is a knowledge graph that *structures* multimodal traces into explicit entities and relations (e.g., places, activities, ambience cues, temporal rhythms) while retaining links to supporting evidence. Rather than collapsing localness into a single undifferentiated summary, the graph represents localness as a set of typed attributes and relations whose support can be inspected at the level of nodes, edges, and associated media. This includes explicitly distinguishing between signals that are directly observed (e.g., visual or transcript-derived cues) and those that are inferred or synthesized, making gaps and uncertainty diagnosable. This representation yields benefits beyond richer content. First, it supports *structure-sensitive* retrieval and composition: graph-grounded retrieval improved recall of local attributes and helped answer queries that depend on relationships (e.g., matching vibe, time-of-day, and neighborhood context) rather than isolated place descriptions. Second, it supports *diagnosis and maintenance* of a long pipeline: when outputs are implausible, we can localize which attributes, relations, or sparse evidence regions drove the result, enabling targeted refinement (e.g., schema adjustments, sampling strategies, or source expansion). Finally, the graph provides a substrate for monitoring representational skew. For example, identifying neighborhoods or community perspectives that are undersupported by evidence and should be prioritized for enrichment.

6.1.3 Surfacing provenance for inspection, contestation, and accountability.

Beyond representing provenance, PlaceWeave treats provenance as an interface-level resource for transparency and control. Unlike opaque end-to-end generative models, the system is designed so users can *inspect why* a suggestion was made, navigate supporting media, and calibrate reliance rather than treating outputs as authoritative—aligning with human–AI interaction goals around transparency and seamful design [14, 15]. In our study, participants reported increased confidence when recommendations were accompanied by narrative context and accessible supporting evidence, describing suggestions as "more trustworthy" when they could evaluate their basis. Provenance also enables *contestability* at the level of community representation [28]. Because machine-generated place narratives can shape what visitors seek and what becomes visible, making evidence legible creates a pathway for residents and stakeholders to validate, challenge, or contextualize portrayals of their locale, echoing accountability goals in civic computing [83]. Future iterations could strengthen these affordances with interaction techniques such as lightweight "why am I seeing this?" explanations, uncertainty cues when evidence is thin or skewed, and evidence navigation tools (e.g., interactive evidence maps or fact cards [20]) that make both attribution and limits of coverage explicit.

6.2 Scaffolding Localness Sensemaking in an Integrated Planning Workspace.

6.2.1 Conversation as deliberation: from local preferences to commitments.

Our study suggests that planning with PlaceWeave is better understood as deliberation than query–response: participants iteratively expressed intent, refined constraints, and weighed tradeoffs while negotiating what "local" should mean for *their* trip. Many framed the assistant as a "digital local friend," using dialogue to probe nuance (e.g., vibe, audience, timing) and translate value-laden preferences into actionable candidates. When a suggestion resonated, participants wanted to immediately commit it, underscoring a recurring gap in AI planning tools between *considering* options and *committing* decisions.

A broader implication is that human–AI planning systems should treat *rationales* as first-class planning artifacts, not just explanatory text. Participants used justifications to judge fit, anticipate friction, and compare alternatives without forcing preferences into rigid filters. This shifts the goal from producing a single "best" itinerary toward scaffolding judgment under uncertainty: systems can propose candidates, surface relevant constraints (e.g., opening hours, long jumps across town), and support user-authored tradeoffs while keeping optimization lightweight and optional [93, 105].

6.2.2 Hybrid planning workspaces: coupling dialogue with direct manipulation.

The interface scaffolds deliberation by coupling conversational exploration with a manipulable plan representation, reducing the fragmentation common in travel planning toolchains. Participants moved fluidly from an idea in chat to locating it, checking feasibility, and placing it into a plan—suggesting a general principle for LLM-mediated planning: pair generative dialogue with a workspace that supports externalization, comparison, and incremental commitment. In our case, the planning surface supported a coherent spatial mental model (consistent with imageability [63]), while integrated map–timeline–itinerary views enabled progressive disclosure as users pulled details only when needed to decide [105].

More broadly, this hybrid pattern supports a mixed-initiative model in which systems *propose* and users *compose*: direct manipulation turns AI outputs into provisional materials that can be reordered, pruned, and revised [75]. Because integration can also increase information density, future systems should invest in visual hierarchy and focus modes (e.g., emphasizing landmarks, filtering to "current plan," staging candidate sets) so that evidence-backed planning remains usable as complexity grows. These implications extend beyond travel to other situated planning contexts where users must reconcile preferences, constraints, and socially grounded meanings within a single decision workspace.

6.3 An Agenda for Accountable Place Intelligence: Validity, Reliability, and Community Stewardship

Accountable place intelligence concerns systems that transform fragmented, community-produced traces into actionable *place* narratives for situated decision making—where "place" is treated as lived, socially constructed meaning rather than only coordinates or POI metadata [18, 28, 29, 38, 85, 98]. We use this framing to distinguish our goals from "location intelligence" paradigms centered on organizational geospatial analytics, and from purely metric-driven place analytics; instead, we emphasize evidence legibility, representational responsibility, and governance of downstream impacts

[7, 26, 27, 40]. The limitations below motivate an agenda of coupled research questions: how to represent *coverage* and uncertainty in traces; how to synthesize claims that are inspectable and contestable; how to steward whose narratives become visible; and how to evaluate both traveler experience and community impact at scale.

6.3.1 Reliability under sparse evidence: uncertainty, corroboration, and participatory verification. PlaceWeave is not immune to hallucinations, particularly when social video evidence is sparse or uneven for abstract or civic-oriented queries. This limitation suggests a broader requirement for local-intelligence systems: when evidence is thin, the system should avoid confident synthesis by default and instead make *epistemic status* visible—what is supported, what is inferred, and what is simply unknown.

Addressing this requires interventions across the pipeline and the interface. On the data and modeling side, reliability improves when place claims are *corroborated* across heterogeneous sources (e.g., triangulating short-form video with community forums or local reporting when appropriate) [31, 32, 35, 94], when extraction is robust to noisy transcripts and fleeting visual cues [30], and when synthesis remains *traceable* to specific evidence units rather than collapsing support into a single narrative [33]. Stronger retrieval policies can further reduce overgeneralization by privileging graph-grounded evidence, enforcing coverage thresholds before synthesis, and retrieving diverse, non-redundant support for contentious attributes [25, 48, 49, 73, 77]. In geographically situated domains, this also entails explicitly accounting for non-local or unevenly local contributions in volunteered traces [31, 48, 52].

Equally important are HCI mechanisms that operationalize verification rather than treating it as a backend property. Interfaces can provide lightweight verification widgets (e.g., "verified by evidence," "limited evidence," or "needs confirmation" marks), progressive disclosure of supporting media, and interaction pathways to request corroboration when users sense ambiguity [15, 26, 88]. Because place narratives affect communities, participatory verification should also be treated as a design opportunity: residents and local organizations can flag misrepresentations, supply corrections, or validate contested claims over time, turning reliability into an ongoing sociotechnical process rather than a one-time model improvement [20, 33, 35]. Such seamful, user- and community-in-the-loop designs align with transparency and control goals in human–AI interaction [14, 15, 40].

6.3.2 Representational equity: mitigating bias in social-trace local intelligence. Computational systems built on social media inherit representational skews in who produces content and what becomes visible; PlaceWeave is no exception [95, 96]. Short-form videos often over-index younger, highly online creators and visually striking, already popular venues, while quieter community spaces and perspectives may be underrepresented—a pattern that reflects broader geographies of digital exclusion [37, 47]. Aggregation can then amplify this skew, turning a narrow slice of city life into the apparent whole [4, 96]. Mitigation is therefore sociotechnical: beyond expanding evidence sources, systems should make representational conditions legible (e.g., signaling when recommendations reflect a narrow creator set), support exploration beyond saturated clusters, and enable ongoing stewardship [26]. Crucially, bias is not a one-time "fix": participatory feedback channels that allow residents to flag misrepresentation, suggest additions, and curate overlooked places help keep the system accountable to a broader spectrum of local narratives over time [20, 33, 35].

6.3.3 Cultural and linguistic generalizability: co-defining localness across contexts. Our implementation focused on English-language content in U.S. cities, but "localness" is culturally situated [29, 36]. Generalizing PlaceWeave therefore requires more than multilingual support: it requires collaborative re-specification of what attributes, practices, and values should constitute "local" in a given region [29, 94]. Technically, the system can extend the schema and models to new languages and concepts. Sociotechnically, cross-regional deployments should involve residents in co-designing the localness framework, especially where communities are "data invisible" online and dominant platforms systematically underrepresent them [47, 52, 73, 77, 95].

6.3.4 Aggregation, serendipity, and externalities at scale. Systems like PlaceWeave do more than organize information: by aggregating vernacular content into recommendations, they might reshape what counts as "serendipity" and can concentrate attention on a small set of highly connected places [68]. Participants often experienced this as broadening horizons beyond obvious tourist sites, but this is an infrastructural form of serendipity that can flatten the long tail of local experiences into canonical "must-see" clusters [68, 81]. At the interaction level, our findings suggest multiple pathways to serendipity: open-ended video exploration, but also synthesis-driven suggestions that users can selectively ground by inspecting source media. A key design opportunity is to treat diversity and externalities as first-class objectives. Interfaces can deliberately surface less-saturated options, reveal when recommendations are driven by narrow evidence, and provide controls that steer away from dominant clusters [86]. More broadly, future work should stress-test aggregation effects under scale (e.g., whether "hidden gem" recommendation loops accelerate overcrowding or inequity) and explore mitigations that preserve local community well-being rather than merely optimizing individual discovery [68, 101].

6.3.5 Creator rights and community governance: reuse boundaries and contestation. AI-mediated aggregation also raises questions of creator visibility, consent, and the consequences of reclassifying community traces into decision infrastructure. Content shared for peer-to-peer inspiration may be repurposed for itinerary planning in ways creators did not anticipate, and attribution alone does not guarantee control [89, 110]. Future systems should therefore support clear provenance and credit, but also provide meaningful mechanisms such as opt-out or removal pathways and transparent reuse policies that are legible to both travelers and local communities [26]. Because these issues extend beyond individual UX, we argue for governance-oriented design. Participatory auditing and community challenge mechanisms can surface hidden failure modes and inform constraints on what is recommended and how it is summarized [20, 72, 87]. Longer term, co-design with local organizations can define legitimate operating boundaries (e.g., what sources are acceptable, what inferences are disallowed, and whether automated synthesis is desired at all), turning aggregation from a neutral recommender stance into accountable, community-aligned place intelligence [11, 33].

6.3.6 *Limits of multimodal extraction: what traces fail to capture.*

A separate technical limitation lies in how we sample and interpret video frames for multimodal analysis. Although our pipeline samples boundary and central frames and prioritizes visually salient moments, highly dynamic clips can obscure nuanced gestures, micro-interactions, or fleeting contextual cues [12, 107]. This can yield omissions that are not hallucinations per se, but *coverage failures* that silently narrow what the system can represent. Future work could more tightly couple transcript cues with visual salience and temporal attention to better capture "focus moments", and could expose coverage indicators (e.g., "limited visual evidence") so users understand when localness signals may be incomplete [15, 26].

6.3.7 *External validity: from perceived localness to lived experience.*

Our evaluation was conducted in a lab-like setting in which participants planned hypothetical trips but did not enact them. Accordingly, our results support claims about *planning work*—workflow fit, planning confidence, and *perceived* localness/authenticity based on the evidence and narratives shown in the interface—rather than the lived experiences these itineraries would produce on the ground. Establishing whether PlaceWeave-generated plans lead to richer, community-aligned visits requires in-the-wild, longitudinal study [18, 20]. A next step is a field deployment that combines traveler use during real trips with post-trip diaries and participatory evaluation with residents and local organizations [8, 11, 87], enabling both outcome validation (what experiences occurred) and iterative refinement of what "localness" should mean in that context [28].

6.4 An Agenda Beyond Place: Structured, Traceable Infrastructures for Fragmented Narratives

Although developed for local exploration, PlaceWeave speaks to a broader human–AI interaction challenge: people increasingly rely on fragmented, community-produced traces (e.g., short videos, posts, informal reports) to build *situated* understanding. Beyond place planning, similar fragmentation arises in domains such as crisis informatics, public-health communication, and peer support, where sensemaking depends on assembling many partial accounts into narratives *without* erasing uncertainty or provenance [56, 61, 91, 92].

As LLMs make synthesis easy, the bottleneck often shifts from generation to *infrastructure*: the representational and governance layers that keep synthesis grounded, inspectable, and accountable. We therefore argue for a paradigm of *structured, traceable sociotechnical infrastructures* for fragmented narratives—systems that (i) translate vernacular media into shared, user-legible representations, (ii) maintain traceable provenance from claims back to source traces, and (iii) provide interfaces for contestation and stewardship so users and communities can question, repair, and update how narratives are assembled [23, 90].

We observe a recurring pattern in HCI+AI systems work on accountable synthesis: *ground* system goals in lived perception and interactional expectations [28, 29]; *scaffold* those expectations into user-legible representational substrates that connect human sensemaking to machine inference [40, 79]; *audit* where current AI succeeds and fails, including skews and disparities, using documentation and benchmarks that render limitations inspectable [31, 32, 34, 70, 84]; *steward* accountability through participatory governance with affected stakeholders [33, 50]; and *instantiate* these commitments in domain systems that bind synthesis to accountable assembly rather than treating generation as an endpoint [61, 67, 92].

PlaceWeave instantiates this paradigm for place planning by binding its assistant, map overlays, and route planner to a shared place knowledge graph and provenance-linked evidence, so that "localness" inferences remain auditable and contestable *in the moment of planning*. Graph-RAG is one useful mechanism for retrieval and composition over structured traces, but the broader infrastructural agenda also includes complementary mechanisms directly relevant to this setting: provenance standards for representing how claims were produced [10], documentation and audit artifacts for exposing coverage gaps and bias [34, 70, 84], maintainable shared knowledge substrates [99], and verification workflows that distribute corroboration across people and interfaces rather than outsourcing "truth" to a single model [50]. Framed this way, the opportunity for HCI is to design conceptual architectures and governance interfaces that let users not only consume synthesized narratives, but also *inspect, contest, repair, and steward* how those narratives are assembled [23, 51, 64, 90].

7 Conclusion

This paper introduced PlaceWeave, a system that leverages a Graph-RAG pipeline to extract and structure "platial" knowledge from social video. We demonstrated its effectiveness in improving users' spatial sensemaking and planning confidence. More broadly, our work contributes a sociotechnical approach to the challenge of data migration, showing how structured knowledge graphs can preserve the rich, relational context essential for place-making. We offered design implications for creating more integrated and value-aligned exploratory systems. Finally, through a critical reflection on the politics of data and the ethics of representation, we highlighted the profound responsibility we bear as designers of systems that shape how people understand and interact with the world, arguing for a future of place-based technologies grounded in transparency, provenance, and participatory design.

References

[1] Mark S. Ackerman. 2000. The Intellectual Challenge of CSCW: The Gap Between Social Requirements and Technical Feasibility. *Human–Computer Interaction* 15, 2-3 (2000), 179–203. doi:10.1207/S15327051HCI1523_5

[2] Netta Aizenbud-Reshef, Artem Barger, Ido Guy, Yael Dubinsky, and Shiri Kremer-Davidson. 2012. Bon voyage: social travel planning in the enterprise. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work* (Seattle, Washington, USA) *(CSCW '12)*. Association for Computing Machinery, New York, NY, USA, 819–828. doi:10.1145/2145204.2145326

[3] Wenxi Bai, Timothy J Lee, Fan Wu, and Jose Weng Chou Wong. 2025. How effective are user-generated travel short videos in promoting a destination online? *Journal of Vacation Marketing* 31, 3 (2025), 578–595. doi:10.1177/13567667231221816

[4] Andrea Ballatore and Stefano De Sabbata. 2020. Los Angeles as a digital place: The geographies of user-generated content. *Transactions in GIS* 24, 4 (2020), 880–902. doi:10.1111/tgis.12600

[5] Aaron Bangor, Philip Kortum, and James Miller. 2009. Determining what individual SUS scores mean: adding an adjective rating scale. *J. Usability Studies* 4, 3 (May 2009), 114–123.

[6] Aaron Bangor, Philip T Kortum, and James T Miller. 2008. An Empirical Evaluation of the System Usability Scale. *International Journal of Human–Computer Interaction* 24, 6 (2008), 574–594. doi:10.1080/10447310802205776

[7] Geoffrey C Bowker and Susan Leigh Star. 2000. *Sorting things out: Classification and its consequences.* MIT press.
[8] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101. https://doi.org/10.1191/1478088706qp063oa
[9] John Brooke et al. 1996. SUS-A quick and dirty usability scale. *Usability evaluation in industry* 189, 194 (1996), 4–7.
[10] Peter Buneman, Sanjeev Khanna, and Tan Wang-Chiew. 2001. Why and Where: A Characterization of Data Provenance. In *Database Theory — ICDT 2001*, Jan Van den Bussche and Victor Vianu (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 316–330.
[11] Mads Bødker and David Browning. 2012. Beyond destinations: exploring tourist technology design spaces through local–tourist interactions. *Digital Creativity* 23, 3-4 (2012), 204–224. doi:10.1080/14626268.2012.709939
[12] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-End Object Detection with Transformers. In *Computer Vision – ECCV 2020*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer International Publishing, Cham, 213–229.
[13] Pew Research Center. 2024. How Americans Use Social Media. Pew Research Center Internet & Technology. https://www.pewresearch.org/internet/2024/01/31/americans-social-media-use/ Accessed: 2025-07-09.
[14] Matthew Chalmers. 2003. Seamful design and ubicomp infrastructure. In *Proceedings of Ubicomp 2003 workshop at the crossroads: The interaction of HCI and systems issues in Ubicomp.* 577–584.
[15] Matthew Chalmers and Areti Galani. 2004. Seamful interweaving: heterogeneity in the theory and design of interactive systems. In *Proceedings of the 5th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques* (Cambridge, MA, USA) *(DIS '04)*. Association for Computing Machinery, New York, NY, USA, 243–252. doi:10.1145/1013115.1013149
[16] Aili Chen, Xuyang Ge, Ziquan Fu, Yanghua Xiao, and Jiangjie Chen. 2024. TravelAgent: An AI Assistant for Personalized Travel Planning. arXiv:2409.08069 [cs.AI] https://arxiv.org/abs/2409.08069
[17] Hao Chen, Min Wang, and Zhen Zhang. 2022. Research on Rural Landscape Preference Based on TikTok Short Video Content and User Comments. *International Journal of Environmental Research and Public Health* 19, 16 (2022). doi:10.3390/ijerph191610115
[18] Justin Cranshaw, Andrés Monroy-Hernández, and S.A. Needham. 2016. Journeys & Notes: Designing Social Computing for Non-Places. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 4722–4733. doi:10.1145/2858036.2858573
[19] Justin Cranshaw, Raz Schwartz, Jason Hong, and Norman Sadeh. 2012. The livehoods project: Utilizing social media to understand the dynamics of a city. In *Proceedings of the international AAAI conference on web and social media*, Vol. 6. 58–65. https://doi.org/10.1609/icwsm.v6i1.14278
[20] Justin B. Cranshaw, Kurt Luther, Patrick Gage Kelley, and Norman Sadeh. 2014. Curated city: capturing individual city guides through social curation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) *(CHI '14)*. Association for Computing Machinery, New York, NY, USA, 3249–3258. doi:10.1145/2556288.2557401
[21] Munmun De Choudhury, Moran Feldman, Sihem Amer-Yahia, Nadav Golbandi, Ronny Lempel, and Cong Yu. 2010. Automatic construction of travel itineraries using social breadcrumbs. In *Proceedings of the 21st ACM Conference on Hypertext and Hypermedia* (Toronto, Ontario, Canada) *(HT '10)*. Association for Computing Machinery, New York, NY, USA, 35–44. doi:10.1145/1810617.1810626
[22] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuynck. 2020. ECAPA-TDNN: Emphasized Channel Attention, Propagation and Aggregation in TDNN Based Speaker Verification. In *Interspeech 2020 (interspeech_2020)*. ISCA, 3830–3834. doi:10.21437/interspeech.2020-2650
[23] Nicholas Diakopoulos. 2016. Accountability in algorithmic decision making. *Commun. ACM* 59, 2 (Jan. 2016), 56–62. doi:10.1145/2844110
[24] Shuangyan Du and Cecilia Yin Mei Cheong. 2025. Beyond the scenic view: a multimodal discourse analysis of sustainable tourism imaginaries on TikTok in Anhui, China. *Humanities and Social Sciences Communications* 12, 1 (2025), 690. doi:10.1057/s41599-025-05015-3
[25] Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. 2025. From Local to Global: A Graph RAG Approach to Query-Focused Summarization. arXiv:2404.16130 [cs.CL] https://arxiv.org/abs/2404.16130
[26] Upol Ehsan, Q. Vera Liao, Michael Muller, Mark O. Riedl, and Justin D. Weisz. 2021. Expanding Explainability: Towards Social Transparency in AI systems. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 82, 19 pages. doi:10.1145/3411764.3445188
[27] Batya Friedman, Peter H. Kahn, Alan Borning, and Alina Huldtgren. 2013. *Value Sensitive Design and Information Systems.* Springer Netherlands, Dordrecht, 55–95. doi:10.1007/978-94-007-7844-3_4
[28] Zihan Gao, Justin Cranshaw, and Jacob Thebault-Spieker. 2024. Journeying Through Sense of Place with Mental Maps: Characterizing Changing Spatial Understanding and Sense of Place During Migration for Work. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW2, Article 503 (Nov. 2024), 31 pages. doi:10.1145/3687042
[29] Zihan Gao, Justin Cranshaw, and Jacob Thebault-Spieker. 2025. A Turing Test for "Localness": Conceptualizing, Defining, and Recognizing Localness in People and Machines. arXiv:2505.07282 [cs.HC] https://arxiv.org/abs/2505.07282
[30] Zihan Gao, Jiaying Liu, Yifei Xu, and Jacob Thebault-Spieker. 2025. From Clips to Communities: Fusing Social Video into Knowledge Graphs for Localness-Aware LLMs. In *Companion Publication of the 2025 Conference on Computer-Supported Cooperative Work and Social Computing* (Bergen, Norway) *(CSCW Companion '25)*. Association for Computing Machinery, New York, NY, USA, 497–503. doi:10.1145/3715070.3749277
[31] Zihan Gao and Jacob Thebault-Spieker. 2026. Is Your Chatbot a Tourist or a Townie? Quantifying Geographic and Localness Disparities in LLM Representations of Place. *Proc. ACM Hum.-Comput. Interact.* 10, 2, Article CSCW022 (2026), 45 pages. doi:10.1145/3788058
[32] Zihan Gao, Yifei Xu, and Jacob Thebault-Spieker. 2026. LocalBench: Benchmarking LLMs on County-Level Local Knowledge and Reasoning. *Proceedings of the AAAI Conference on Artificial Intelligence* (2026).
[33] Zihan Gao, Mohsin Y. K. Yousufi, and Jacob Thebault-Spieker. 2025. Collective Narrative Grounding: Community-Coordinated Data Contributions to Improve Local AI Systems. In *NeurIPS 2025 Workshop on Algorithmic Collective Action.* https://openreview.net/forum?id=2ZRwKlGSDa
[34] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for datasets. *Commun. ACM* 64, 12 (Nov. 2021), 86–92. doi:10.1145/3458723
[35] Michael F Goodchild. 2007. Citizens as sensors: the world of volunteered geography. *GeoJournal* 69 (2007), 211–221.
[36] Simon Gottschalk and Elena Demidova. 2018. EventKG: A Multilingual Event-Centric Temporal Knowledge Graph. In *The Semantic Web*, Aldo Gangemi, Roberto Navigli, Maria-Esther Vidal, Pascal Hitzler, Raphaël Troncy, Laura Hollink, Anna Tordai, and Mehwish Alam (Eds.). Springer International Publishing, Cham, 272–287.
[37] Mark Graham and Martin Dittus. 2022. *Geographies of digital exclusion: Data and inequality.* Pluto Books.
[38] Steve Harrison and Paul Dourish. 1996. Re-place-ing space: the roles of place and space in collaborative systems. In *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work* (Boston, Massachusetts, USA) *(CSCW '96)*. Association for Computing Machinery, New York, NY, USA, 67–76. doi:10.1145/240080.240193
[39] Christopher Healey and James Enns. 2012. Attention and Visual Memory in Visualization and Computer Graphics. *IEEE Transactions on Visualization and Computer Graphics* 18, 7 (2012), 1170–1188. doi:10.1109/TVCG.2011.127
[40] Jeffrey Heer. 2019. Agency plus automation: Designing artificial intelligence into interactive systems. *Proceedings of the National Academy of Sciences* 116, 6 (2019), 1844–1850. doi:10.1073/pnas.1807184115
[41] M Carmen Hidalgo and Bernardo Hernandez. 2001. Place attachment: Conceptual and empirical questions. *Journal of environmental psychology* 21, 3 (2001), 273–281.
[42] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia D'amato, Gerard De Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, Axel-Cyrille Ngonga Ngomo, Axel Polleres, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan Sequeda, Steffen Staab, and Antoine Zimmermann. 2021. Knowledge Graphs. *ACM Comput. Surv.* 54, 4, Article 71 (July 2021), 37 pages. doi:10.1145/3447772
[43] Hsun-Ping Hsieh and Cheng-Te Li. 2014. Traveling Path Recommendation Using Temporal Transit Patterns. *Proceedings of the International AAAI Conference on Web and Social Media* 8, 1 (May 2014), 648–649. doi:10.1609/icwsm.v8i1.14495
[44] Hsun-Ping Hsieh and Cheng-Te Li. 2021. Composing Traveling Paths from Location-Based Services. *Proceedings of the International AAAI Conference on Web and Social Media* 6, 1 (Aug. 2021), 618–619. doi:10.1609/icwsm.v6i1.14228
[45] Yuntong Hu, Zhihan Lei, Zhongjie Dai, Allen Zhang, Abhinav Angirekula, Zheng Zhang, and Liang Zhao. 2025. CG-RAG: Research Question Answering by Citation Graph Retrieval-Augmented LLMs. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Padua, Italy) *(SIGIR '25)*. Association for Computing Machinery, New York, NY, USA, 678–687. doi:10.1145/3726302.3729920
[46] Yuntong Hu, Zhihan Lei, Zheng Zhang, Bo Pan, Chen Ling, and Liang Zhao. 2025. GRAG: Graph Retrieval-Augmented Generation. In *Findings of the Association for Computational Linguistics: NAACL 2025*, Luis Chiruzzo, Alan Ritter, and Lu Wang (Eds.). Association for Computational Linguistics, Albuquerque, New Mexico, 4145–4157. doi:10.18653/v1/2025.findings-naacl.232
[47] Isaac L. Johnson, Yilun Lin, Toby Jia-Jun Li, Andrew Hall, Aaron Halfaker, Johannes Schöning, and Brent Hecht. 2016. Not at Home on the Range: Peer Production and the Urban/Rural Divide. In *Proceedings of the 2016 CHI Conference*

on Human Factors in Computing Systems (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 13–25. doi:10.1145/2858036.2858123

[48] Isaac L. Johnson, Subhasree Sengupta, Johannes Schöning, and Brent Hecht. 2016. The Geography and Importance of Localness in Geotagged Social Media. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 515–526. doi:10.1145/2858036.2858122

[49] Bradley S Jorgensen and Richard C Stedman. 2011. Measuring the spatial component of sense of place: a methodology for research on the spatial dynamics of psychological experiences of places. *Environment and Planning B: Planning and Design* 38, 5 (2011), 795–813.

[50] Prerna Juneja and Tanushree Mitra. 2022. Human and Technological Infrastructures of Fact-checking. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 418 (Nov. 2022), 36 pages. doi:10.1145/3555143

[51] Helena Karasti. 2014. Infrastructuring in participatory design. In *Proceedings of the 13th Participatory Design Conference: Research Papers - Volume 1* (Windhoek, Namibia) *(PDC '14)*. Association for Computing Machinery, New York, NY, USA, 141–150. doi:10.1145/2661435.2661450

[52] Ankit Kariryaa, Isaac Johnson, Johannes Schöning, and Brent Hecht. 2018. Defining and Predicting the Localness of Volunteered Geographic Information using Ground Truth Data. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3173574.3173839

[53] Elias Kärle, Umutcan Şimşek, Oleksandra Panasiuk, and Dieter Fensel. 2018. Building an Ecosystem for the Tyrolean Tourism Knowledge Graph. In *Current Trends in Web Engineering*, Cesare Pautasso, Fernando Sánchez-Figueroa, Kari Systä, and Juan Manuel Murillo Rodríguez (Eds.). Springer International Publishing, Cham, 260–267.

[54] David Kirsh. 2010. Thinking with external representations. *AI & SOCIETY* 25, 4 (2010), 441–454. doi:10.1007/s00146-010-0272-8

[55] Qiuqiang Kong, Yin Cao, Turab Iqbal, Yuxuan Wang, Wenwu Wang, and Mark D. Plumbley. 2020. PANNs: Large-Scale Pretrained Audio Neural Networks for Audio Pattern Recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 28 (2020), 2880–2894. doi:10.1109/TASLP.2020.3030497

[56] Yubo Kou, Xinning Gui, Yunan Chen, and Kathleen Pine. 2017. Conspiracy Talk on Social Media: Collective Sensemaking during a Public Health Crisis. *Proc. ACM Hum.-Comput. Interact.* 1, CSCW, Article 61 (Dec. 2017), 21 pages. doi:10.1145/3134696

[57] Charis Lengen and Thomas Kistemann. 2012. Sense of place and place identity: Review of neuroscientific evidence. *Health & Place* 18, 5 (2012), 1162–1171. doi:10.1016/j.healthplace.2012.01.012

[58] Agnieszka Leszczynski. 2014. On the Neo in Neogeography. *Annals of the Association of American Geographers* 104, 1 (2014), 60–79. doi:10.1080/00045608.2013.846159

[59] James R. Lewis and Jeff Sauro. 2018. Item benchmarks for the system usability scale. *J. Usability Studies* 13, 3 (May 2018), 158–167.

[60] Congying Liu, Mingdi Jiang, and Zulqarnain Arshad Muhammad. 2024. The impact of TikTok short video factors on tourists' behavioral intention among Generation Z and Millennials: The role of flow experience. *PLOS ONE* 19, 12 (12 2024), 1–23. doi:10.1371/journal.pone.0315140

[61] Michael Xieyang Liu, Tongshuang Wu, Tianying Chen, Franklin Mingzhe Li, Aniket Kittur, and Brad A Myers. 2024. Selenite: Scaffolding Online Sensemaking with Comprehensive Overviews Elicited from Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 837, 26 pages. doi:10.1145/3613904.3642149

[62] Yu Liu, Jingtao Ding, Yanjie Fu, and Yong Li. 2023. UrbanKG: An Urban Knowledge Graph System. *ACM Trans. Intell. Syst. Technol.* 14, 4, Article 60 (May 2023), 25 pages. doi:10.1145/3588577

[63] Kevin Lynch. 1960. The image of the environment. *The image of the city* 11 (1960), 1–13.

[64] Yao Lyu, Jie Cai, and John M. Carroll. 2025. A Systematic Literature Review of Infrastructure Studies in SIGCHI. *Proc. ACM Hum.-Comput. Interact.* 9, 7, Article CSCW364 (Oct. 2025), 41 pages. doi:10.1145/3757545

[65] Dean MacCannell. 1973. Staged authenticity: Arrangements of social space in tourist settings. *American journal of Sociology* 79, 3 (1973), 589–603.

[66] Wendy E. Mackay and Joanna McGrenere. 2025. Comparative Structured Observation. *ACM Trans. Comput.-Hum. Interact.* 32, 2, Article 14 (April 2025), 27 pages. doi:10.1145/3711838

[67] Damien Masson, Sylvain Malacria, Géry Casiez, and Daniel Vogel. 2024. DirectGPT: A Direct Manipulation Interface to Interact with Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 975, 16 pages. doi:10.1145/3613904.3642462

[68] Giovanni Mauro, Marco Minici, and Luca Pappalardo. 2026. The urban impact of AI: modelling feedback loops in location-based recommender systems. *Machine Learning* 115, 1 (2026), 19. doi:10.1007/s10994-025-06904-z

[69] Darragh McCashin and Colette M Murphy. 2023. Using TikTok for public and youth mental health – A systematic review and content analysis. *Clinical Child Psychology and Psychiatry* 28, 1 (2023), 279–306. doi:10.1177/13591045221106608 PMID: 35689365.

[70] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (Atlanta, GA, USA) *(FAT* '19)*. Association for Computing Machinery, New York, NY, USA, 220–229. doi:10.1145/3287560.3287596

[71] Stein Monteiro. 2024. Searching for Settlement Information on Reddit. *International Migration* 62, 3 (2024), 100–119. doi:10.1111/imig.13261

[72] Barbara Neuhofer, Dimitrios Buhalis, and Adele Ladkin. 2014. A Typology of Technology-Enhanced Tourism Experiences. *International Journal of Tourism Research* 16, 4 (2014), 340–350. doi:10.1002/jtr.1958

[73] Hang Ni, Fan Liu, Xinyu Ma, Lixin Su, Shuaiqiang Wang, Dawei Yin, Hui Xiong, and Hao Liu. 2025. TP-RAG: Benchmarking Retrieval-Augmented Large Language Model Agents for Spatiotemporal-Aware Travel Planning. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 12403–12429. doi:10.18653/v1/2025.emnlp-main.626

[74] Jakob Nielsen. 2005. Ten usability heuristics. (2005).

[75] Masato Nomiyama, Toshiki Takeuchi, Hiroyuki Onimaru, Tomohiro Tanikawa, Takuji Narumi, and Michitaka Hirose. 2018. Xnavi: Travel Planning System Based on Experience Flows. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 1, Article 27 (March 2018), 25 pages. doi:10.1145/3191759

[76] Sangkeun Park, Yongsung Kim, Uichin Lee, and Mark Ackerman. 2014. Understanding localness of knowledge sharing: a study of Naver KiN 'here'. In *Proceedings of the 16th International Conference on Human-Computer Interaction with Mobile Devices & Services* (Toronto, ON, Canada) *(MobileHCI '14)*. Association for Computing Machinery, New York, NY, USA, 13–22. doi:10.1145/2628363.2628407

[77] Boci Peng, Yun Zhu, Yongchao Liu, Xiaohe Bo, Haizhou Shi, Chuntao Hong, Yan Zhang, and Siliang Tang. 2025. Graph Retrieval-Augmented Generation: A Survey. *ACM Trans. Inf. Syst.* 44, 2, Article 35 (Dec. 2025), 52 pages. doi:10.1145/3777378

[78] Peter Pirolli and Stuart Card. 1999. Information foraging. *Psychological review* 106, 4 (1999), 643–675. https://doi.org/10.1037/0033-295X.106.4.643

[79] Peter Pirolli and Stuart Card. 2005. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of international conference on intelligence analysis*, Vol. 5. McLean, VA, USA, 2–4.

[80] Daniele Quercia, Neil Keith O'Hare, and Henriette Cramer. 2014. Aesthetic capital: what makes london look beautiful, quiet, and happy?. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing* (Baltimore, Maryland, USA) *(CSCW '14)*. Association for Computing Machinery, New York, NY, USA, 945–955. doi:10.1145/2531602.2531613

[81] Daniele Quercia, Rossano Schifanella, and Luca Maria Aiello. 2014. The shortest path to happiness: recommending beautiful, quiet, and happy routes in the city. In *Proceedings of the 25th ACM Conference on Hypertext and Social Media* (Santiago, Chile) *(HT '14)*. Association for Computing Machinery, New York, NY, USA, 116–125. doi:10.1145/2631775.2631799

[82] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine Mcleavey, and Ilya Sutskever. 2023. Robust Speech Recognition via Large-Scale Weak Supervision. In *Proceedings of the 40th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 202)*, Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (Eds.). PMLR, 28492–28518. https://proceedings.mlr.press/v202/radford23a.html

[83] Joseph Rafidi. 2013. Real-time trip planning with the crowd. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems* (Paris, France) *(CHI EA '13)*. Association for Computing Machinery, New York, NY, USA, 2743–2748. doi:10.1145/2468356.2479503

[84] Inioluwa Deborah Raji, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) *(FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 33–44. doi:10.1145/3351095.3372873

[85] Edward Relph. 1976. *Place and placelessness*. Vol. 67. Pion London.

[86] Eve Schade, Gian-Luca Savino, Jasmin Niess, and Johannes Schöning. 2023. MapUncover: Fostering Spatial Exploration through Gamification in Mobile Map Apps. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 678, 13 pages. doi:10.1145/3544548.3581428

[87] Chouki Sfandla and Peter Björk. 2013. Tourism Experience Network: Co-creation of Experiences in Interactive Processes. *International Journal of Tourism Research* 15, 5 (2013), 495–506. doi:10.1002/jtr.1892

[88] Ben Shneiderman. 2003. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *The Craft of Information Visualization*, BENJAMIN B. BEDERSON and BEN SHNEIDERMAN (Eds.). Morgan Kaufmann, San Francisco, 364–371. doi:10.1016/B978-155860915-0/50046-9
[89] Ellen Simpson and Bryan Semaan. 2021. For You, or For"You"? Everyday LGBTQ+ Encounters with TikTok. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW3, Article 252 (Jan. 2021), 34 pages. doi:10.1145/3432951
[90] Susan Leigh Star and Karen Ruhleder. 1994. Steps towards an ecology of infrastructure: complex problems in design and access for large-scale collaborative systems. In *Proceedings of the 1994 ACM Conference on Computer Supported Cooperative Work* (Chapel Hill, North Carolina, USA) *(CSCW '94)*. Association for Computing Machinery, New York, NY, USA, 253–264. doi:10.1145/192844.193021
[91] Kate Starbird and Leysia Palen. 2011. "Voluntweeters": self-organizing by digital volunteers in times of crisis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) *(CHI '11)*. Association for Computing Machinery, New York, NY, USA, 1071–1080. doi:10.1145/1978942.1979102
[92] Sangho Suh, Bryan Min, Srishti Palani, and Haijun Xia. 2023. Sensecape: Enabling Multilevel Exploration and Sensemaking with Large Language Models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) *(UIST '23)*. Association for Computing Machinery, New York, NY, USA, Article 1, 18 pages. doi:10.1145/3586183.3606756
[93] Yihong Tang, Zhaokai Wang, Ao Qu, Yihao Yan, Zhaofeng Wu, Dingyi Zhuang, Jushi Kai, Kebing Hou, Xiaotong Guo, Jinhua Zhao, Zhan Zhao, and Wei Ma. 2024. ItiNera: Integrating Spatial Optimization with Large Language Models for Open-domain Urban Itinerary Planning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, Franck Dernoncourt, Daniel Preoţiuc-Pietro, and Anastasia Shimorina (Eds.). Association for Computational Linguistics, Miami, Florida, US, 1413–1432. doi:10.18653/v1/2024.emnlp-industry.104
[94] Jacob Thebault-Spieker, Aaron Halfaker, Loren G. Terveen, and Brent Hecht. 2018. Distance and Attraction: Gravity Models for Geographic Content Production. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) *(CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3173574.3173722
[95] Jacob Thebault-Spieker, Brent Hecht, and Loren Terveen. 2018. Geographic Biases are 'Born, not Made': Exploring Contributors' Spatiotemporal Behavior in OpenStreetMap. In *Proceedings of the 2018 ACM International Conference on Supporting Group Work* (Sanibel Island, Florida, USA) *(GROUP '18)*. Association for Computing Machinery, New York, NY, USA, 71–82. doi:10.1145/3148330.3148350
[96] Jacob Thebault-Spieker, Loren Terveen, and Brent Hecht. 2017. Toward a Geographic Understanding of the Sharing Economy: Systemic Biases in UberX and TaskRabbit. *ACM Trans. Comput.-Hum. Interact.* 24, 3, Article 21 (April 2017), 40 pages. doi:10.1145/3058499
[97] Pier Paolo Tricomi, Saurabh Kumar, Mauro Conti, and V.S. Subrahmanian. 2024. Climbing the Influence Tiers on TikTok: A Multimodal Study. *Proceedings of the International AAAI Conference on Web and Social Media* 18, 1 (May 2024), 1503–1516. doi:10.1609/icwsm.v18i1.31405
[98] Yi-Fu Tuan. 1977. *Space and place: The perspective of experience.* U of Minnesota Press.
[99] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Commun. ACM* 57, 10 (Sept. 2014), 78–85. doi:10.1145/2629489
[100] Cheng Wang, Wenjing Cui, Yating Zhang, and Huawen Shen. 2022. Exploring short video apps users' travel behavior intention: Empirical analysis based on SVA-TAM model. *Frontiers in Psychology* Volume 13 - 2022 (2022). doi:10.3389/fpsyg.2022.912177
[101] Yana Wengel, Ling Ma, Yixiao Ma, Michal Apollo, Kamil Maciuk, and Ann Suwaree Ashton. 2022. The TikTok effect on destination development: Famous overnight, now what? *Journal of Outdoor Recreation and Tourism* 37 (2022), 100458. doi:10.1016/j.jort.2021.100458
[102] Jacob O. Wobbrock and Julie A. Kientz. 2016. Research contributions in human-computer interaction. *Interactions* 23, 3 (April 2016), 38–44. doi:10.1145/2907069
[103] Guihua Wu and Xinyi Ding. 2023. Which type of tourism short video content inspires potential tourists to travel. *Frontiers in Psychology* Volume 14 - 2023 (2023). doi:10.3389/fpsyg.2023.1086516
[104] Jian Xie, Kai Zhang, Jiangjie Chen, Tinghui Zhu, Renze Lou, Yuandong Tian, Yanghua Xiao, and Yu Su. 2024. TravelPlanner: a benchmark for real-world planning with language agents. In *Proceedings of the 41st International Conference on Machine Learning* (Vienna, Austria) *(ICML'24)*. JMLR.org, Article 2246, 24 pages.
[105] Alexandre Yahi, Antoine Chassang, Louis Raynaud, Hugo Duthil, and Duen Horng (Polo) Chau. 2015. Aurigo: an Interactive Tour Planner for Personalized Itineraries. In *Proceedings of the 20th International Conference on Intelligent User Interfaces* (Atlanta, Georgia, USA) *(IUI '15)*. Association for Computing Machinery, New York, NY, USA, 275–285. doi:10.1145/2678025.2701366
[106] Qinggang Zhang, Shengyuan Chen, Yuanchen Bei, Zheng Yuan, Huachi Zhou, Zijin Hong, Hao Chen, Yilin Xiao, Chuang Zhou, Junnan Dong, Yi Chang, and Xiao Huang. 2025. A Survey of Graph Retrieval-Augmented Generation for Customized Large Language Models. arXiv:2501.13958 [cs.CL] https://arxiv.org/abs/2501.13958
[107] Yifu Zhang, Peize Sun, Yi Jiang, Dongdong Yu, Fucheng Weng, Zehuan Yuan, Ping Luo, Wenyu Liu, and Xinggang Wang. 2022. ByteTrack: Multi-object Tracking by Associating Every Detection Box. In *Computer Vision – ECCV 2022*, Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (Eds.). Springer Nature Switzerland, Cham, 1–21.
[108] Chen Zhao, Huawen Shen, and Yating Zhang. 2022. The study on the impact of short video tourism Vloggers at social media platform on online sharing intention. *Frontiers in Psychology* Volume 13 - 2022 (2022). doi:10.3389/fpsyg.2022.905002
[109] Qing Zhou, Marios Sotiriadis, and Shiwei Shen. 2023. Using TikTok in tourism destination choice: A young Chinese tourists' perspective. *Tourism Management Perspectives* 46 (2023), 101101. doi:10.1016/j.tmp.2023.101101
[110] Diana Zulli and David James Zulli. 2022. Extending the Internet meme: Conceptualizing technological mimesis and imitation publics on the TikTok platform. *New Media & Society* 24, 8 (2022), 1872–1890. doi:10.1177/1461444820983603

A Technical Appendix: Dataset, Pipeline, and Schema

A.1 Dataset Acquisition and Curation

Sites. City U (urban), City S (suburban), County R (rural).
Acquisition. We used Apify ([3]) to collect public, geotagged, English-language videos from TikTok. This approach was necessary as the official TikTok Research API lacks the essential geotag data required for our place-based analysis.
Filtering. Exclude duplicates, non-geo, commercial or advertising content;
Counts. City U: 2,650 videos (33.2h), 1,210 creators; City S: 2,900 (36.1h), 1,340 creators; County R: 700 (9.0h), 410 creators; total 6,250 (78.3h), 2,840 creators. Verified places: 1,540 unique.
Language (ASR). EN.

A.2 Pipeline Implementation Details

A.2.1 Stage 1: Multimodal Feature Extraction and Alignment.

Visual Keyframe Extraction. We applied Shot Boundary Detection (SBD) using PySceneDetect[4] to identify major visual transitions and extracted representative frames at these boundaries using FFmpeg. Each keyframe's timestamp was recorded for temporal alignment.

Visual Content Analysis. Within each keyframe, we performed:

- **On-screen Text Extraction:** We applied OCR using PyTesseract[5] to capture embedded textual information.
- **Object Detection:** We used DETR (DEtection TRansformer) [12] to identify and localize salient objects.
- **Landmark Recognition:** We applied the Google Landmark Recognition API[6] to identify notable geographical or cultural landmarks.
- **Scene Description Generation:** We used the GPT-4o vision-language model, providing it with both the raw keyframe and the detected objects, to produce rich, semantic captions.

Audio Content Analysis. Audio tracks were processed as follows:

[3]https://apify.com
[4]https://github.com/Breakthrough/PySceneDetect
[5]https://github.com/madmaze/pytesseract
[6]https://developers.google.com/maps/documentation/landmark

- **Audio Transcription:** We transcribed all spoken narration using OpenAI's Whisper model [82] for its robustness to real-world noise.
- **Environmental Sound and Paralinguistic Analysis:** We employed pre-trained PANNs [55] and YAMNet models to tag environmental sounds (e.g., ambient noise, music genres). Additionally, SpeechBrain's ECAPA-TDNN model [22] was used to generate speaker embeddings to capture sociolinguistic cues.

Temporal Alignment. To capture dynamic activities, we incorporated ByteTrack [107] to track detected objects across frames. All extracted features were then synchronized using their timestamps to form a cohesive, time-aligned representation of each video.

A.2.2 Stage 2: Geographic Entity Grounding and Verification.

Candidate Generation and Resolution. We firstly used spaCy's `en_core_web_trf` model for initial NER on all textual data. To handle vernacular language, GPT-4o was used to reason over the video's context and resolve ambiguous mentions (e.g., "the Square"). These entities were then resolved to canonical locations using the Google Places API and enriched with structured metadata from Wikidata.

Verification Heuristics. A candidate location was considered a verified match if the cosine similarity between its official metadata (from Wikidata) and its contextual description within the video exceeded a threshold of 0.72. Tie-breakers included temporal comention and a neighborhood prior. This process auto-verified 87% of mentions. A manual spot-check of 300 sampled mentions revealed a false-match error rate of approximately 2.6%.

A.2.3 Stage 3: Knowledge Graph Construction.

Implementation. The graph was implemented as a `NetworkX MultiDiGraph` within a PostgreSQL database, using the `pgvector` extension for efficient similarity search on embeddings.

A.3 Knowledge Graph Schema and Scale

Node Types and Attributes.

- **Location Nodes:** Geographic entities with coordinates, standard/vernacular names, types (e.g., park), and engagement metrics (e.g., frequency of appearance).
- **Video Nodes:** Metadata about the source TikTok video, including text, interaction metrics (likes, shares), and hashtags.
- **Keyframe Nodes:** Visual snapshots with associated scene descriptions and temporal positioning.
- **Object Nodes:** Recognized visual elements with labels, confidence scores, and bounding box coordinates.
- **Experiential Nodes:** Abstract concepts extracted from content, including *Ambience* (e.g., quiet, festive), *Activity* (e.g., study, live-music), *Time* (hour-bucket, weekday/weekend), and *UserRole* (e.g., student, tourist).

Edge Types and Semantics.

- **Location–Location:** Co-occurrence, transitions, or cultural links (e.g., `near`, `transitioned_to`).
- **Content–Location:** Links between a video or keyframe and a place it depicts (e.g., `shows`, `mentions`).
- **Content–Experiential:** Links between a piece of media and an abstract concept (e.g., `describes(MediaItem, Ambience)`).
- **Place–Experiential:** The core relational links (for example, `co_occurs(Place, Activity)`).

Scale and Technology. The final graph contains **28,341 nodes** and **96,507 edges**. Embeddings for semantic retrieval were generated using the `all-mpnet-base-v2` model (768d) and indexed using HNSW (M=16, $ef\_construction$=200).

B Participant Tasks Assignments

See Table 4.

Table 4: Counterbalancing scheme for the user study (n=18). This design ensures that the order of conditions and the pairing of sites to conditions were fully balanced across all participants.

Participant(s)	Task 1 Condition	Task 1 Site	Task 2 Condition	Task 2 Site
P1 – P3	PlaceWeave	Urban (U)	Baseline	Suburban (S)
P4 – P6	Baseline	Urban (U)	PlaceWeave	Suburban (S)
P7 – P9	PlaceWeave	Suburban (S)	Baseline	Rural (R)
P10 – P12	Baseline	Suburban (S)	PlaceWeave	Rural (R)
P13 – P15	PlaceWeave	Rural (R)	Baseline	Urban (U)
P16 – P18	Baseline	Rural (R)	PlaceWeave	Urban (U)