Towards Circular Accumulation of Latent Local Knowledge via Automated Interview System: A Preliminary Study

Aoi Sassa*,§, Haru Terashima*, Naoki Tamura*, Kazuyuki Shoji†,
Kenta Urano*, Takuro Yonezawa*, Tadashi Yoshikawa‡, Nobuo Kawaguchi*†
*Graduate School of Engineering, Nagoya University, Japan
†Institutes of Innovation for Future Society, Nagoya University, Japan
‡Toyota Technical Development Corporation, Japan
§aoi@ucl.nuee.nagoya-u.ac.jp

***Abstract*—Local knowledge, such as hidden gems and the perceived character of places, is valuable across numerous regional development domains, including tourism and regional policy, making its systematic accumulation a critical challenge. However, existing approaches structurally miss local knowledge from residents who do not actively share information online, and offer no mechanism to store or reuse it. Therefore, even as the scale of data collection expands, improvements in quality are hindered. To address these challenges, we propose an LLM-based interview system that actively elicits local knowledge, combining profile-adaptive question generation with Cognitive Conflict Theory-based probing. The system is designed to collect residents' place-specific knowledge and accumulate it into a reusable knowledge base that is referenced in subsequent dialogues to support deeper elicitation. In an experiment with 38 residents of a Japanese city, the system collected 74 knowledge units. Based on participants' self-reports, 94% had not posted the content online themselves.**

***Index Terms*—AI interview system, local knowledge accumulation, latent knowledge**

I. Introduction

Local knowledge held by residents such as impressions of places, local customs, and personal experiences is useful in regional development applications, including tourism, regional revitalization, and relocation support. Collecting and sharing such knowledge can attract visitors and prevent population outflow from rural areas. However, this knowledge is not being collected systematically, and building an accumulation infrastructure is an important challenge. In fact, only about 30% of social media users actively post [1]. Therefore, we need collection methods that do not rely on individuals' willingness to share information. In addition, local knowledge changes over time as events, businesses, and community practices evolve, so a one-time collection quickly becomes outdated.

Previous research on the collection of local knowledge ranging from Volunteer Geoinformation (VGI) to social media data mining has demonstrated its usefulness for such applications [2], [3]. However, these methods assume residents voluntarily post information. Residents who do not post may still share local knowledge in everyday conversation with those around them, but this knowledge is missed by existing approaches. The other limitation is that even when local knowledge is collected through surveys or VGI, no mechanism exists to store and reuse it. Without such a platform, each collection starts from scratch and cannot be kept fresh or grow in coverage as it scales.

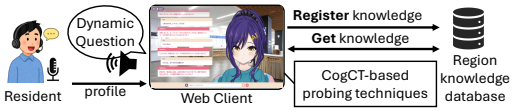


Fig. 1. Overview of the proposed LLM-based interview system for circular accumulation of latent local knowledge

We define this missed knowledge as Latent Local Knowledge, classified into Subjective and Objective Knowledge. The former is particularly valuable, because it conveys the lived regional appeal central to tourism and relocation support, while the latter can often already be obtained from public sources. This separation reflects a fundamental difference in nature. Objective knowledge has a single answer that is invariant across observers, whereas subjective knowledge admits a diversity of valid views. This work focuses on subjective knowledge, with objective knowledge collected as a byproduct.

To address these problems, we propose an LLM-based conversational interview system. Fig. 1 shows an overview of the system. By guiding residents through dialogue, the system collects knowledge without requiring residents to publicly post on social media themselves. It also stores collected knowledge in real time and reuses it in later sessions, forming an accumulation infrastructure that grows in depth and coverage over time. To collect subjective knowledge effectively, the system uses two techniques. First, **Profile-adaptive Question Generation** selects topics based on each resident's background profile to elicit specific local knowledge. Second, **Cognitive Conflict Theory (CogCT)-based Probing** presents opinions from other residents to prompt participants to verbalize previously unconscious knowledge [4].

In an experiment with 38 residents in a Japanese city, the system collected 74 knowledge units (as segmented by the LLM). By participants' self-reports, 94% had not posted it

online themselves, and 72% had shared it verbally with family or neighbors.

This paper makes the following contributions:

- Problem framing and system concept for eliciting non-posted local knowledge
- Prototype design combining profile adaptation, memory reuse, and CogCT-inspired probing
- Preliminary experiment indicating the feasibility of eliciting self-reported non-posted local knowledge

II. Related Work

A. Positioning of This Work in Local Knowledge Collection

Collection methods for local knowledge have evolved over time. Goodchild proposed VGI, enabling low-cost collection through voluntary citizen contributions [2]. More recently, researchers have collected subjective local knowledge by analyzing social media posts and review sites [3]. However, both approaches assume that residents voluntarily post information. Zheng et al. pointed out that "what we can obtain is always a sample of data from partial users, as not everyone shares data." [5] Since platforms cannot control whether users post, the knowledge inevitably reflects only those willing to share. This problem has been recognized since the early days of VGI and remains unsolved. Knowledge held by non-posting residents is structurally missed. The proposed system incorporates mechanisms to facilitate information sharing through dialogue, and reduces bias in the collection of knowledge without requiring residents to publicly post information themselves.

On the other hand, local knowledge also changes over time, so one-time collection cannot maintain freshness. Niedomysl and Kitagawa show that long-term communication of regional appeal to potential residents is important [6], [7], so local knowledge must stay fresh and cumulative. However, existing approaches have no mechanism to store and reuse collected knowledge in future sessions, so each collection starts from scratch. The proposed system is designed to address both issues through a circular accumulation mechanism: knowledge is accumulated in real time and reused in later sessions, keeping the collected data fresh and cumulative.

B. Positioning of Knowledge Elicitation Techniques

In face-to-face interviews, social desirability bias prevents people from sharing honest opinions. Lucas et al. showed that participants disclose more to virtual agents when they believe the agents are computer-controlled rather than human-operated, due to reduced social pressure [8]. The proposed system adopts LLM-based dialogue based on this finding.

Fixed-question surveys cannot adapt to each resident's knowledge domain, yielding only surface-level responses. Cuevas et al. demonstrated that LLMs can act as "adaptive surveyors" by dynamically generating specific follow-up questions [9]. The proposed system applies this as profile-adaptive question generation that infers familiar topics from background profiles to collect more detailed local knowledge.

Subjective knowledge is often held unconsciously and cannot be elicited by simple questions. CogCT [4] shows that exposure to differing views prompts self-reflection and deeper verbalization of one's own knowledge. Yao et al. recently applied CogCT to the construction of data for cultural alignment in LLMs, using inter-cultural divergence to prompt reflection on one's own culture [10]. Jo et al. showed that referencing past dialogue as memory promotes deeper self-disclosure within the same user over time [11], but their memory utilization was restricted to individual users' own past dialogues. The proposed system extends this across different users: knowledge collected through the circular accumulation mechanism is reused as CogCT stimuli, a design intended to support deeper elicitation as more knowledge accumulates.

III. Proposed Method

A. System Overview

We propose LLM-based interview system for circular accumulation of latent local knowledge. The system is a fully automated pipeline that conducts voice-based interviews with local residents. It focuses on subjective knowledge, using CogCT to probe deeper into individual opinions. Objective knowledge that emerges during the conversation is also collected as a byproduct. The system is built on three design principles:

1) Profile-adaptive question generation: Infers familiar topics from the profiles and generates questions accordingly
2) CogCT-based probing: Presents the opinions of other residents to elicit more detailed and specific responses
3) Circular accumulation: Immediately saves collected knowledge to a database, enabling reuse across sessions

B. Interview Flow Design

The interview follows four steps.

1) Profile Collection: At the start of the session, the system collects basic attributes through conversation, including age, gender, family composition, and years of residence. These are used to infer the resident's familiar topics in the next step.

2) Topic Determination: Based on the profile, the LLM infers topics the resident is familiar with and presents them in conversation. For example, playground knowledge is offered to a parent with young children, and restaurant knowledge near the office is offered to an office worker. Once the resident's familiar topic is found, the session moves to probing.

3) CogCT-based Probing: The system conducts deep probing to elicit the resident's subjective knowledge and personal experiences about the topic by applying CogCT [4]. To generate a conflict prompt, the LLM retrieves opinions from other residents with similar backgrounds on the same topic and presents them as: "Some people say [knowledge]—what do you think?" This encourages residents to compare the presented opinion with their own experience and share more specific, personal information.

4) Termination: By default, the session ends once two or more knowledge units have been collected, though this threshold was adjusted during the experiment as needed because of session time constraints. Whether knowledge has been acquired at this point is determined in real time by the LLM during the session. The units of knowledge are also determined

by the LLM. Acquired knowledge is saved to the database in real time to be reused in later sessions.

C. Technical Architecture

The system architecture is shown in Fig. 2. It consists of a frontend that handles user interactions and a backend server that manages the interview process and stores local knowledge.

1) Frontend: The frontend is a web client based on AITuber Kit [12], with minor modifications by the authors. Avatar characters facilitate the conversation, and the chat log panel displays the conversation history. For voice-input/voice-response dialogue, it uses the Web Speech API for speech recognition and VOICEVOX [13] for text-to-speech.

2) Backend: The backend uses Dify [14] to manage interview sessions (Dify Chatflow), and to store local knowledge (Dify knowledge base). Conversation Log Module saves conversation history in real time and provides it to other modules as needed. Interview Flow Module controls the flow of the interview, by switching the prompt text sent to the LLM, and this flow follows the Interview Flow Design described above. User Profile Module stores user attributes collected during the session and adds them to the prompt. LLM Agentic Module takes the conversation log including the current user input, the user profile text, the prompt for the current interview flow, as input to generate a response. LLM Agentic Module can call the Search Function and Record Function as needed. Search Function vectorizes the search query generated by the LLM and retrieves relevant knowledge from the database using cosine similarity and keyword matching. Record Function vectorizes the new knowledge generated by the LLM, structured as "A person who is [profile] thinks [knowledge] about [topic]", and stores it in the database with the speaker's attributes.

The database is Dify's knowledge base. It is designed to store subjective knowledge and objective knowledge separately. However, the current prototype does not include an automatic classification module, and all collected knowledge is stored in a single database with speaker attributes. Automated classification is left as future work.

IV. Experiments

A. Experimental Setup

We conducted an experiment with 38 residents in Shinshiro City, Aichi Prefecture, Japan (February 2026, Shinshiro City West Community Center). Two systems ran in parallel, with each session lasting 5–10 minutes. Post-session questionnaires were completed by 36 participants.

At session start, participants gave on-screen consent to anonymized use of their conversation and questionnaire data for system improvement and academic publication. Internally, profile and utterance data are retained in raw form for matching and retrieval, but profiles are never shown to other participants. Any external release is anonymized.

B. Collection Results

74 knowledge units were collected (LLM-segmented). The collected data included personal stories and specific episodes. For example: "From the perspective of a man in his 40s with children, Shinshiro General Park has well-equipped playground facilities. The long slide is very popular with kids. You can have a picnic with a packed lunch, and children naturally learn to take turns with strangers. On the downside, it gets very hot in summer, so bring drinks."

The authors informally searched each unit online using multiple keyword queries, and some of units could not be found. Examples include: "Regarding a particular cherry blossom spot as the aging trees become more noticeable each year, I hope that young trees will be planted gradually to preserve the landscape," "The Shinshiro Summer Fireworks Festival never gets old because the concept changes every year," and "The cherry blossom spots in Ichikawa are renowned for their stunning beauty; they feel like a smaller-scale version of the Senbon-zakura (Thousand Cherry Trees) at Mount Yoshino."

TABLE I
Questionnaire results on the latent status of collected knowledge ($n = 36$)

Question	Yes	No
Have you shared this with family or people in the community?	72%	28%
Have you posted this online or in any other public forum?	6%	94%

C. Initial Signals on Latent Status

To gather initial signals on the latent status of the collected knowledge, we asked participants about their sharing behavior (Table I). 72% had shared the content in local conversations, and 94% had never posted it online. This suggests that the system collects knowledge that residents share in conversation but do not actively publish online. Note that this reflects each participant's own posting behavior. A more stringent goal in this system is that the collected content should not be found in publicly available online sources. Our informal internet keyword search failed to detect some of the content, providing qualitative evidence supporting this goal. Establishing a reliable latent share metric remains a challenge for future work. Results are based on self-reports, so memory bias and social desirability effects cannot be ruled out. Furthermore, the 38 participants were drawn from a single municipality, so generalizability across regions and community sizes is not established by this preliminary study.

D. CogCT Probing: Illustrative Example

Table II shows a practical example of CogCT, extracted from a conversation about the gohei-mochi tea house (gohei-mochi is a traditional rice cake skewer with miso sauce). Before the LLM presented another person's opinion, the resident gave a brief preference statement. After the presentation, the response expanded to include a personal evaluation and new knowledge (making it at home). This illustrates how the mechanism works in practice, but absent a control condition, does not constitute evidence of its effectiveness. Quantitative comparison against a non-CogCT baseline remains future

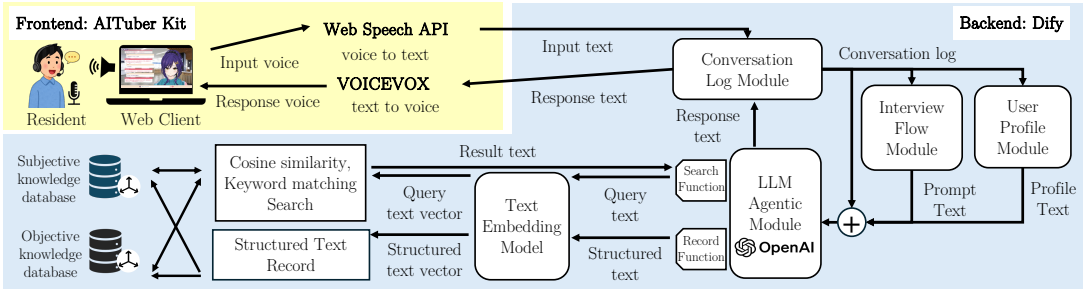


Fig. 2. System architecture showing frontend components (Web Speech API, VOICEVOX, Dify Chatflow) and backend modules for flow control, search, and knowledge recording.

TABLE II
ILLUSTRATIVE EXAMPLE OF THE COGCT-BASED PROBING MECHANISM (CONVERSATION EXCERPT AT A GOHEI-MOCHI TEA HOUSE)

Speaker	Content
Resident	I like gohei-mochi with just miso sauce—not the kind that has walnut paste in it.
LLM (opinion presentation)	So you prefer the simple kind with just miso. By the way, some people say "gohei-mochi has a strong flavor"— how did you find the intensity? Was it just right, or do you prefer a milder taste?
Resident (after deepening)	It was just right for me. I sometimes make it at home, too.

work. So does assessment of user experience and AI-avatar-specific pressures (e.g., recording awareness, accommodation to the agent's expectations).

V. CONCLUSION AND FUTURE WORK

We proposed an LLM-based interview system that combines profile-adaptive question generation, memory reuse, and CogCT-inspired probing to elicit latent local knowledge from residents. In an experiment with 38 residents in a Japanese city, the system collected 74 knowledge units (as segmented by the LLM); based on participants' self-reports, 94% had not posted the content online themselves. The system adopts a circular structure designed to collect, preserve, and reuse knowledge from residents who do not actively post online, to keep the data fresh and growing as sessions accumulate.

Future work spans three layers. (1) Design: We will implement a system that uses NLP technology to automatically classify statements as subjective or objective, and adaptive dialogue termination based on similarity to the accumulated knowledge base. (2) Experiment: Through A/B testing (with and without CogCT), we will quantitatively evaluate the effectiveness of CogCT in terms of factors such as utterance length and the novelty of the information. We will verify the circular structure by conducting experiments with more participants. Furthermore, we aim to maximize the proportion of collected content that is genuinely latent (absent from both participants' own posts and public online sources). We will develop a reliable latent-share metric and iterate on the dialogue design by analyzing cases where the system successfully elicited such content. (3) Application: We will develop a system to contribute to regional development by sharing our accumulated knowledge with residents for supporting those who need it, with local governments for informing policy-making, and with visitors for conveying the region's appeal.

ACKNOWLEDGMENT

This work was partially supported by Cabinet Office's SIP3 (JPJ012495), JST CREST (JPMJCR22M4), and JST RISTEX (JPMJRS23K3).

REFERENCES

[1] Mobile Society Research Institute, NTT DOCOMO, INC., "Mobile society white paper 2025, reference material 2-6," 2025, (In Japanese). [Online]. Available: https://www.moba-ken.jp/whitepaper/wp25.html
[2] M. F. Goodchild, "Citizens as sensors: the world of volunteered geography," *GeoJournal*, vol. 69, no. 4, pp. 211–221, 2007.
[3] V. Taecharungroj *et al.*, "Analysing tripadvisor reviews of tourist attractions in phuket, thailand," *Tour. Manag.*, vol. 75, pp. 550–568, 2019.
[4] R. A. Cosier *et al.*, "Cognitive conflict and goal conflict effects on task performance," *Organizational Behavior and Human Performance*, vol. 19, no. 2, pp. 378–391, 1977.
[5] Y. Zheng *et al.*, "Urban computing: Concepts, methodologies, and applications," *ACM Trans. Intell. Syst. Technol.*, vol. 5, no. 3, Sep. 2014.
[6] F. Kitagawa, "Lessons from the japanese 'regional revitalisation' – impacts of regional industrial policies through decentralisation," *Contemporary Social Science*, vol. 19, no. 4, pp. 602–621, 2024.
[7] T. Niedomysl, "Towards a conceptual framework of place attractiveness: A migration perspective," *Geografiska Annaler: Series B, Human Geography*, vol. 92, no. 1, pp. 97–109, 2010.
[8] G. M. Lucas *et al.*, "It's only a computer: Virtual humans increase willingness to disclose," *Computers in Human Behavior*, vol. 37, pp. 94–100, 2014.
[9] A. Cuevas *et al.*, "Collecting qualitative data at scale with large language models: A case study," arXiv:2309.10187, 2024.
[10] J. Yao *et al.*, "Caredio: Cultural alignment via representativeness and distinctiveness guided data optimization," arXiv:2504.08820, 2026.
[11] E. Jo *et al.*, "Understanding the impact of long-term memory on self-disclosure with large language model-driven chatbots for public health intervention," in *CHI '24*, 2024.
[12] tegnike, "aituber-kit," GitHub repository. [Online]. Available: https://github.com/tegnike/aituber-kit
[13] VOICEVOX, "voicevox," GitHub repository. [Online]. Available: https://github.com/VOICEVOX/voicevox
[14] langgenius, "dify," GitHub repository. [Online]. Available: https://github.com/langgenius/dify