Systems Friction as a First-Class Design Requirement for Outsiders in Civic Digital Services

Independent research note (portfolio writeup; not peer-reviewed)

Jade Zhao
Informatics student, Luddy School (affiliation for contact only)
Indiana University Bloomington
jlzhao@iu.edu
Independent personal research; not IU-endorsed or lab output

Abstract

Civic and public digital services routinely assume local context, language fluency, and unwritten process knowledge. For outsiders (non-native speakers, new arrivals, and first-time users), that assumption appears as systems friction: stalled forms, PDF-only "language support," blaming error messages, and channels that do not match outsider schedules. This independent research note locks a falsifiable claim: treating outsider friction as a first-class design requirement, rather than a post-launch accessibility patch, is a necessary input for accountable civic digital services and for public-interest automation as AI tools proliferate over years. Method: reflective pattern extraction from a Spring 2026 semester abroad in Madrid (Complutense / IU Overseas Study), recorded as a public-safe pattern library in the madrid-ai-ethics repository, plus a reusable Outsider Impact Review one-pager for builders. Contribution at this stage is the claim, five observed patterns with concrete mechanisms, the review artifact, a grounded discussion of long-horizon AI growth (capability growth is not automatic inclusion; friction compounds when automation lands on a broken access baseline), and explicit limits (single observer, one city, reflective method, no controlled trial, no market forecasts). Adjacent literature on administrative literacy, administrative burden, and WCAG baselines documents related harms; the operational gap targeted here is a builder-facing pre-automation gate. No quantitative validation is claimed. A minimal dual-review kill test is specified for later work. This PDF is a portfolio-style writeup, not a peer-reviewed publication and not IU or clinic research output.

1 Introduction

Public services increasingly run through forms, apps, and portals. When those systems are designed for people who already know the local script (which office comes first, which document format is accepted, which language completes the task), newcomers experience the same interface as a wall. That wall is easy to misread as personal failure. This note treats it as a design signal.

Situation. Spring 2026, the author lived and studied in Madrid through IU Overseas Study at Universidad Complutense. Coursework (POLS-Y 350 Politics of the European Union; HISP-S 100 Elementary Spanish; BUS-J 306 Strategic Management and Leadership; plus the overseas study package) supplied vocabulary for institutional design and multi-level governance. Daily life supplied ground truth for how civic, university, commercial, and transit systems meet someone who does not yet belong. Observation was organized by layer: civic / government-facing; university; commercial apps; transit and utilities—asking who the form is written for and what breaks when the user is new. Parallel clinic observation at ServeIT made the same failure modes visible in nonprofit handoffs: tools built for the maintainer's mental model, not the visitor's first visit [4]. A restaurant-floor parallel remains useful: a system either survives contact with a real customer or it does not; public services are customer-facing systems too, with higher stakes and often worse error UX.

Gap. Accessibility practice (for example WCAG 2.1 baselines used in clinic work) catches many markup and interaction failures [1]. A happy-path or checklist pass can still miss outsider-specific assumptions about local context, sequencing, and language-as-access. Administrative literacy and administrative burden research documents related costs for migrants and claimants in digital welfare systems [2, 3]. What those bodies of work do not fully supply for student and clinic builders is a portable pre-ship gate that can stop automation when baseline access fails—especially as generative tools become easier to bolt onto broken flows.

Claims. Stated as claims, not results:

1. Systems friction in civic digital services systematically disadvantages outsiders; the five patterns in Section 4 are the working library from Madrid reflective notes (suggestive, not causal proof).

2. Treating that friction as a first-class design input, via a reusable Outsider Impact Review, is a practice bridge toward more accountable public digital

services and public-interest AI adoption over time.

3. Long-horizon AI capability growth does not automatically produce inclusion; without access-before-automation discipline, friction and "who the system assumes you are" compound as tools proliferate.

4. The current method is reflective practice with an explicit kill test; no controlled experiment results are claimed here.

Section 2 situates prior themes; Section 3 describes method and design principles; Section 4 reports observed patterns; Section 5 discusses long-term AI growth; Section 6 states limits; Section 7 concludes.

2 Related Work

Three clusters frame the gap. This is not a complete literature survey; deep reading remains owned by the author beyond the sources cited below.

Accessibility and comprehensibility. Clinic practice uses WCAG 2.1 as a baseline for keyboard paths, contrast, structure, and related success criteria [1]. Language access and plain-language defaults sit adjacent to comprehensibility even when checklists focus on markup. Mixed-language households also need structural lang marking rather than a single page default—otherwise assistive technology guesses. That overlap is necessary but not sufficient when the primary task is only completable after hidden local prerequisites.

Administrative literacy, burden, and migrant access. Safarov [2] analyzes Russian-speaking migrants navigating digital public services in Finland and revisits administrative literacy as a combination of bureaucratic literacy, structural knowledge, welfare-system navigation, digital skills, and language skills. The work shows that digitalisation can impose rigid requirements that burden or prevent access to entitlements, and that digital skills alone do not explain exclusion. Moynihan, Herd, and Harvey [3] conceptualize administrative burden as learning, psychological, and compliance costs in citizen–state interactions, and argue that burden levels are often political choices rather than accidents. These sources ground the claim that outsider disadvantage in digital public services is a serious, documented phenomenon. They focus more on claimant resources and state–citizen costs than on a builder-facing requirement that outsider friction gate automation before launch.

Public-interest AI and accountable automation. ServeIT / Serve-AI clinic instincts (human review, honest handoffs, accessibility before automation) rhyme with outsider friction as a ship/no-ship input [5, 4]. Serve-AI is ServeIT Clinic's public-interest AI track (PIT-UN): practical adoption help for community organizations so generative tools do not break trust or accessibility. Policy vocabulary about public-sector automated decision-making (including EU-layer discourse encountered in coursework) is useful for reading headlines; it rarely supplies a one-pager students can run before shipping clinic or civic automation. This note's artifact targets that operational gap. EU AI Act categories and "digital single market" rhetoric are tracked as open public-discourse questions, not as legal conclusions from this author.

3 Approach

Method now: reflective practice and pattern extraction from lived systems use during one semester abroad, recorded as public-safe notes (no private administrative detail). Exploration versus confirmation are kept separate in the research log; confirmation experiments are planned, not run. Primary sources for this writeup are the repository essays on systems friction abroad, language and civic access, AI ethics in public services, and the Madrid-to-ServeIT practice bridge, plus the locked claim package.

The primary artifact is an Outsider Impact Review one-pager. Before shipping a civic flow or adding automation, the reviewer records the system and primary task, names an outsider persona, and marks pass / fail / n/a with a one-line note on six checks: assumed local context; unwritten sequencing; language as access (completable interface, not PDF-only "support"); hours and channels with a usable human path and case identity; recoverable errors that name the fix; automation risk (whether AI would amplify any of the above). Required-before-ship items include plain-language / language-path fixes, sequencing stated in the UI, human override, error copy that names the fix, and deferring automation until the access baseline works. Finished reviews must copy a limitation note: single reviewer; persona-based; not a controlled study.

Practice bridge only: ServeIT / Serve-AI clinic work is a rhyme, not a relicensing of those materials. Serve-AI project materials are CC BY-SA 4.0 with project credit; this Madrid writeup and repository remain independent personal notes under MIT.

3.1 Design Principles

- Plain-language and completable language paths for the primary task (not PDF-only "support").
- Sequencing stated in the interface; no hidden prerequisite steps.
- Recoverable errors that name the fix and preserve state.
- Human escalation with case identity when the digital path fails.
- Transparency over theater: when a decision involves automated scoring, routing, or generation, citizens should know what human override exists.

- No assumed prior belonging: design for the first-time outsider, then add automation only if the access baseline works.

4 Findings

Patterns below appeared repeatedly in lived use of civic, university, commercial, and transit/utility systems in Madrid. They suggest design failure modes; they do not prove causation beyond this observer's semester. Recording method: public-safe reflective notes and essay outlines in madrid-ai-ethics.

1. Assumed local context. Forms and apps expect address history, bank relationships, or national ID workflows a newcomer has never used. Commercial and payment profiles that assume a domestic phone number or local bank relationship fail in ways that look like user error.

2. Unwritten sequencing. The flow works only if step B happened before step A, but the interface does not say so. University and civic offices often have a single office that actually handles a case; discovering it is itself unwritten process knowledge.

3. Language as access control. Spanish-only defaults; "English support" that is a PDF rather than a completable interface; language choice functions like a login wall for task completion. Mixed-language households also need structural lang marking rather than a single page default. Elementary Spanish coursework was not fluency on arrival; it was enough to notice when systems pretended everyone already had native proficiency.

4. Hours and channels that do not match outsider life. In-person only, narrow windows, little or no status tracking after submit, especially hard for students and new arrivals. When phone support becomes the "real" interface, app failure equals access failure—especially for outsiders who cannot fall back on family word-of-mouth.

5. Error messages that blame the user. "Invalid document" instead of naming the fix (for example a different scan format); recoverable error UX is missing when stakes are highest. Wiping form state on one bad field raises the cost further when the user did not choose the workflow.

Civic apps differ from consumer apps in a material way: the user did not choose the workflow and is trying not to miss a deadline. That raises the cost of missing progress visibility, wiping form state on one bad field, or trapping the user in a chatbot loop without a reopen path. Design implications that follow from the language and civic-access notes: progress visibility (where am I; what happens next); recoverable errors that preserve state; a human escalation path with a clear case ID; and honest offline / paper parity where not everyone lives inside a perfect mobile signal.

Same failure modes show up in ServeIT clinic observation: tools built for the maintainer's mental model, not the visitor's first visit. Evaluation bar for later work (not yet run): apply Outsider Impact Review and a parallel happy-path / WCAG-style pass to two public civic flows; if the outsider pass finds nothing the happy-path pass misses, the "first-class requirement" half of the claim is in serious trouble for this artifact. No baselines, sample sizes, or percentage improvements are claimed here.

5 Long-term AI Growth

This section is reflective portfolio analysis, not a longitudinal empirical study and not a market forecast. No adoption percentages, capability timelines, or growth rates are asserted as fact. The question is what the Madrid pattern library implies if AI tools keep proliferating in civic and community systems over years rather than demo week.

Friction compounds as tools proliferate. When automation is added to a flow that already assumes local belonging, the outsider does not get a shortcut; they get a second failure mode—opaque routing, a chatbot loop, or a generated answer that cannot reopen a case. The repository's AI-ethics essay states the spine plainly: access before automation. If the baseline service is already unusable for non-native speakers or first-time users, adding a chatbot does not fix access. Long-horizon capability growth makes that mistake cheaper to ship and harder to reverse once staff and citizens treat the automated path as the default.

Who the system assumes you are. Outsider disadvantage is not only about missing a translation string. It is about encoded assumptions—address history, national ID, bank relationships, language fluency, office hours that match insider schedules. As more decisions involve automated scoring, routing, or generation, those assumptions can harden into proxies (language preference, device type, incomplete profile fields) without anyone naming them as design choices. Transparency without a usable human override is theater; accountability requires logs a human can review and a path to contest the outcome.

Public-interest AI at community scale (Serve-AI rhyme). Serve-AI's public framing is responsible adoption for community organizations: accessibility alongside automation, human review before publish, and trust that survives contact with real clients [5]. Different scale than national government, same ethical spine: who is harmed if this fails; can staff override the model; is the data scope honest. Over years, small organizations will keep meeting vendor defaults and demo-week tools. The durable question is not whether models get more capable; it is whether partners can maintain, disclose, and override AI-touched workflows without a specialist on

permanent retainer.

Growth of capability ≠ automatic inclusion. Accessibility, language access, staff capacity, and recoverable error design remain bottlenecks even if model performance improves. A more fluent generator can still amplify PDF-only "language support," blame-the-user errors, or sequencing that only locals know. Inclusion is a maintenance and governance problem as much as a model problem. Treating outsider friction as a first-class pre-automation gate is one way to keep long-term AI adoption accountable without pretending this note measured adoption curves.

6 Discussion and Limitations

Limitations are part of the claim package. Single observer: selection and confirmation bias are real. One city and one semester (Madrid, Spring 2026): other municipalities and seasons may differ. Reflective method: pattern extraction from lived use, not lab measurement. Public-safe scope omits private details that might strengthen or nuance a finding. Related-work deep reading is still owned by the author beyond the citations above; a seed list in the source repository is not a finished survey. This note does not claim legal or policy expertise (EU coursework supplied vocabulary for reading public discourse only), statistical generalizability, market sizes, percentage growth rates, or that automation always harms outsiders. Poorly scoped automation can amplify gaps; that is a design risk, not a theorem. It is not IRB, clinical, or health-system research. Affiliation lines name where the author studies; they do not mean Indiana University, ServeIT Clinic, or any lab endorses or owns this content. Readers should not treat this PDF as a published conference paper or as official IU research. The long-term AI discussion above is interpretive synthesis from existing public notes, not longitudinal evidence.

7 Conclusion

Outsiders hit civic digital systems where design assumes belonging. This independent research note locks that friction as a first-class design concern, documents five public-safe patterns from Madrid reflective practice with concrete mechanisms, offers an Outsider Impact Review template for builders before automation, and argues that long-horizon AI growth without access-first discipline compounds outsider disadvantage. The next empirical step is the dual-review kill test, not invented metrics. Until that runs, the contribution is the claim, the patterns, the artifact, the AI-growth framing, and the limits.

Acknowledgments

Independent personal research notes. Drafting and structure were AI-assisted; accountability for claims, limits, and publication decisions remains with Jade Zhao. Practice bridge only to ServeIT Clinic and Serve-AI (PIT-UN); Serve-AI project materials are CC BY-SA 4.0 (credit Serve-AI / ServeIT Clinic, IU Luddy). This writeup is not Serve-AI licensed material and is not IU-endorsed research.

References

[1] W3C Web Accessibility Initiative, Web Content Accessibility Guidelines (WCAG) 2.1, W3C Recommendation, 5 June 2018. https://www.w3.org/TR/WCAG21/.

[2] N. Safarov, "Administrative Literacy in the Digital Welfare State: Migrants Navigating Access to Public Services in Finland," Social Policy and Society, 2023. https://doi.org/10.1017/s1474746422000719.

[3] D. P. Moynihan, P. Herd, and H. Harvey, "Administrative Burden: Learning, Psychological, and Compliance Costs in Citizen–State Interactions," Journal of Public Administration Research and Theory, vol. 25, no. 1, pp. 43–69, 2015. https://doi.org/10.1093/jopart/muu009.

[4] ServeIT Clinic, Luddy School of Informatics, Computing, and Engineering, Indiana University. https://serveit.luddy.indiana.edu/.

[5] Serve-AI project page (ServeIT Clinic / PIT-UN track). https://serveit.luddy.indiana.edu/serve-ai/index.html.