INLAY: Preemptive, In-Context Intelligence for Casual Web Browsing

Microsoft Research

CHI'26

aipoc@microsoft.com | Pratyay Suvarnapathaki, Adithya S Kolavi, Harsh Vijay, Mayukh Das, Ajay Manchepalli, Venkata N Padmanabhan

Status-Quo: 'AI Sidebar'

Browsers rely on explicit invocation for AI Assistance, typically in a separate UI surface (the 'sidebar').

Effective for goal-oriented tasks, but doesn't account for '**casual browsing**,' i.e. meandering, serendipitous information seeking where intent is implicit or latent.

- **Initiation Friction** (remembering to invoke AI assistance)
- **Cognitive Translation Friction** (formalizing vague thoughts into a prompt)

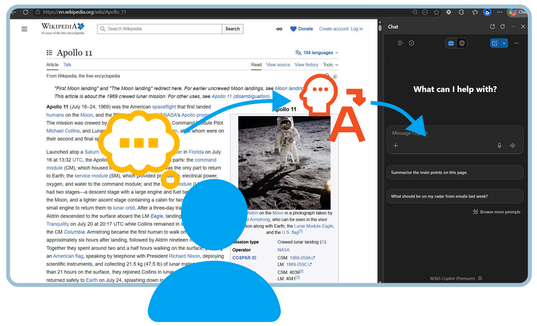

Theoretical Grounding

Information Foraging Theory (IFT) models casual browsing as a cost-reward optimization.

- Users follow **'information scents'** (cues like headings, links, or images) to decide whether to **explore** further or **exploit** the current page.
- Diminishing returns lead to **Premature Satisficing** (settling for subpar results to avoid context-switching)

Missed opportunities for deeper engagement and serendipitous discovery.

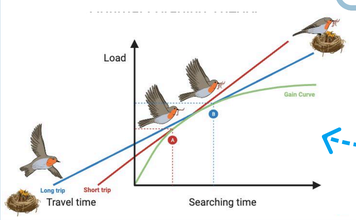

Inlay: Core Principles and System Architecture

We introduce Inlay as an early provocation, positioned **diametrically opposite to the status-quo**:

Instead of awaiting invocation, Inlay **proactively embeds AI-driven insights and cues directly within webpages**, meeting users exactly where their attention already is ← Zero Initiation Friction

This **amplifies existing 'information scents'** by shifting the UX from high-effort query formulation to **low-effort intent confirmation**, drastically reducing foraging costs ← Low Translation Friction

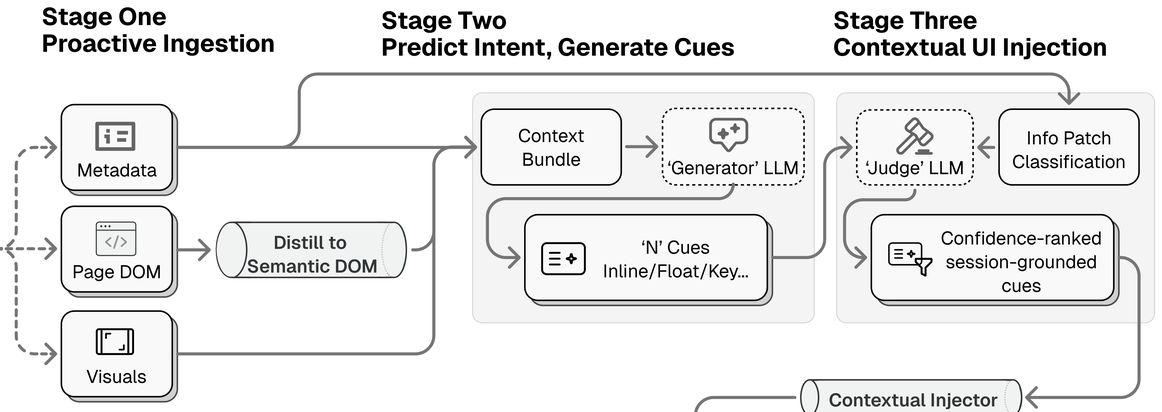

Design Choices

Mixed-Initiatives Design

The tedious task of scanning the page is automated away; the user remains in the driver's seat for serendipitous casual browsing.

'Delegate to Copilot?'

Inlay augments existing AI browsers by externalizing latent user intent and delegating it to downstream agents for execution.

Preserving User Agency

Inlay's suggestions are surfaced as low-friction cues that are easy to '**delegate**' or simply '**dismiss**'

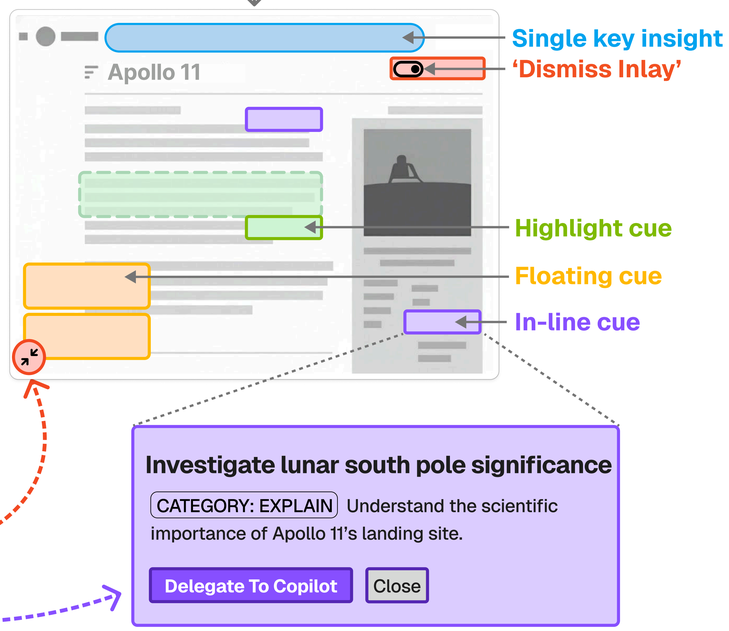

Exploratory User Study

N=12, 30-minute freeform browsing session. Quantitative survey, semi-structured interview.

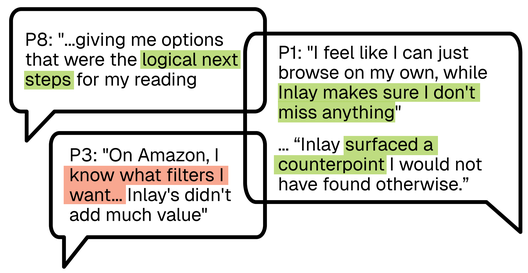

Quantitative Results (Telemetry + Likert)

High utility signal (~**5:1 delegate:dismiss** ratio)

Habituation over session duration. User activity peaks in the third quarter, stabilises in the fourth.

- Q4 highest-rated → Inlay's proactivity did not compromise users' perceived agency.
- Strong **preference for in-line** and in-context highlight cues over AI sidebars
- Q5 → Cues redundant in familiar domains, but highly valued for **uncovering 'unknown unknowns'** in unfamiliar ones

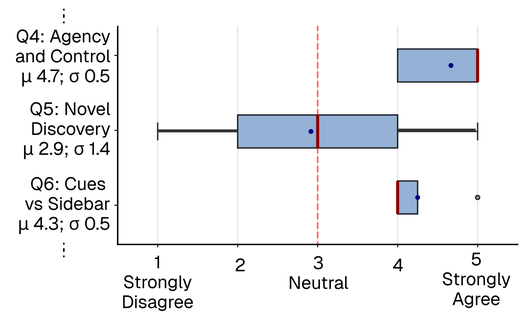

Future Work

For practical consumer deployment, need to bridge casual and task-oriented browsing by integrating intent clarification with robust agentic action-taking.

Subsequent studies needed to model adaptive intervention timing based on real-time intent and assess longitudinal persistence.